ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25sep1208

# Advanced Anfis-Based Maximum Power Point Tracking for Solar Photovoltaic Systems: A Comparative Study with Deep Learning and Real-Time Implementation

Raphael I. Areola[1]; Isola O. Mathew[2]; Oyelade A. Omolara[3]

[1]Department of Electrical Power Engineering, Durban University of Technology, South Africa
[2] Department of Electrical Electronic Engineering, Ekiti State Polytechnic, Isan, Ekiti State, Nigeria
[3] Department of Electrical & Electronics Engineering, Federal Polytechnic, Ile Oluji, Ondo State, Nigeria

Publication Date: 2025/09/29

**Abstract: Solar photovoltaic (PV) capacity is expanding rapidly, yet real-world energy yield still hinges on how reliably controllers track the maximum power point under disturbances such as partial shading, fast irradiance ramps, sensor noise, and embedded hardware limits. This review evaluates three intelligence families for MPPT: Adaptive Neuro-Fuzzy Inference Systems (ANFIS), Deep Learning (DL), and Reinforcement Learning (RL)through a deployment lens rather than simulation alone. Using a structured search (2018–2025) across major databases, we prioritised studies with processor-/hardware-in-the-loop (PIL/HIL) or embedded MCU/FPGA validation, and judged methods on four discriminating metrics: (i) global-peak hit rate under shading, (ii) convergence time and overshoot, (iii) steady-state power ripple, and (iv) edge feasibility (number format, latency, resources), alongside interpretability and audit requirements. Findings show ANFIS as the risk-adjusted frontrunner in non-benign conditions: compact, fixed-point designs consistently deliver millisecond-scale settling and ~99–100% tracking in dynamic tests, while hybrids (e.g., ANFIS-PSO/GEP or with nonlinear scaffolds) further suppress ripple and improve global-peak discovery. DL/RL can match or exceed ANFIS when rich sensing, compute headroom, and mature ML governance exist, but their gains are contingent on data pipelines, quantisation/latency engineering, safe exploration, and explainability. We recommend a SIL→PIL→HIL rollout, energy-weighted metrics under standardised shading/ramp scripts, and deploying a lean, auditable ANFIS now graduating DL/RL where HIL-proven advantages justify their operational complexity.**

**How to Cite:** Raphael I. Areola; Isola O. Mathew; Oyelade A. Omolara (2025) Advanced Anfis-Based Maximum Power Point Tracking for Solar Photovoltaic Systems: A Comparative Study with Deep Learning and Real-Time Implementation. *International Journal of Innovative Science and Research Technology*, 10(9), 1908-1918. https://doi.org/10.38124/ijisrt/25sep1208

## I. INTRODUCTION

Global energy policy is pivoting to renewables, requiring a significant reduction in carbon emissions across the entire energy industry, as well as in end-use sectors. IRENA's 2023 outlook urges tripling global renewable power and doubling efficiency by 2030 to stay on a 1.5 °C path [1], [2]. About 295 GW of renewables, roughly 83% of all new capacity, was added in 2022, yet IRENA warns growth must accelerate further to hit the 2030 goal, especially in emerging and developing economies. Leveraging low-cost solar PV, onshore and offshore wind, and other renewable electricity generation sources, the power sector must lead the way as solutions in other sectors scale up. Accelerating the progress of the transition worldwide requires a holistic approach, backed by systemic innovation to transform existing structures and systems built for the fossil fuel era. By 2030, global total installed renewable power generation capacity would need to expand more than threefold, from 3,382 GW in 2022 to 11,174 GW, according to IRENA's 1.5°C Scenario [3], [4]. Specifically, installed solar PV capacity would rise to more than 5,400 GW, from 1,055 GW in 2022, and wind installations would surpass 3,500 GW (3,040 GW onshore and 500 GW offshore), up from 899 GW in 2022, over the same period [3]. [5], see Figure 1. The share of variable renewable energy (VRE), such as solar PV and wind power, in electricity generation is expected to rise from 10% of the total electricity generated in 2021 to 46% by 2030, necessitating additional flexibility in the operation of the energy system [2]. Solar PV is scaling into more variable, harder-to-control operating environments from rooftops with intermittent shading to utility-scale arrays facing rapid

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25sep1208

irradiance ramps [6]. Yet PV arrays exhibit nonlinear, time-varying I–V/P–V behaviour; under partial shading, the power surface develops multiple local maxima, so basic trackers oscillate or lock onto sub-optimal points, wasting energy [6]. In these conditions, maximum power point tracking (MPPT) remains critical: the PV array's power surface is nonlinear and time-varying, and under partial shading it becomes multi-modal, with several local maxima that can trap basic trackers [7], [8], [9]. Conventional MPPT algorithms such as Perturb & Observe (P&O) and Incremental Conductance (IncCond) are simple and inexpensive, yet they typically oscillate around the setpoint, show slow recovery after sudden transients, and mis-lock at local peaks during shading events, leading to measurable energy loss at scale [9], [10], [11], [12].

Against this backdrop, AI-enabled controllers have expanded rapidly. Among these, the Adaptive Neuro-Fuzzy Inference System (ANFIS) stands out because it fuses fuzzy rule-based transparency with data-driven parameter tuning, often improving convergence and reducing steady-state ripple compared with classical methods while remaining more interpretable than deep nets [8], [12]. ANFIS's five-layer structure (fuzzification, rule firing, normalization, consequents, output) lets designers encode expert knowledge (e.g., slope cues from dP/dVdP/dVdP/dV) and then learn membership functions and consequents from data [7], [8]. Recent reviews and case studies consistently report fast tracking and robustness under changing conditions, though rule-based growth and computational load can become practical concerns for embedded targets [6], [13], [14]. At the same time, there's intense activity in deep learning (DL) and reinforcement learning (RL) MPPT. LSTM-based approaches forecast irradiance or MPP trajectories to prime the controller; some reports show gains over P&O and feed-forward nets in dynamic tests [7], [8]. RL methods (e.g., DQN, PPO) learn a duty-cycle policy that can discover global maxima under shading and adapt on the fly, sometimes outperforming classical baselines but with training complexity and safety considerations [9], [10].

Two practical gaps motivate this review. First, real-time deployment evidence is uneven: studies vary in sampling rates, quantization, and hardware budgets, making it hard to compare "bench-top" performance with embedded or hardware-in-the-loop results [15]. Second, with regulators and operators seeking explainable, auditable control, there's limited adoption of explainable AI (XAI) artifacts, e.g., rule visualizations or sensitivity analyses that help field engineers understand why a controller acted a certain way [16], [17]. This review synthesizes ANFIS-based MPPT evidence (2018–2025), positions it against DL/RL alternatives, details real-time implementation lessons, and outlines XAI practices that can make AI MPPT controllers more trustworthy in the field. The review focuses on ANFIS for MPPT and its comparative context (DL/RL), emphasizing partial shading, fast transients, and deployment on MCUs/DSPs/FPGAs. We privilege peer-reviewed sources from 2018–2025.

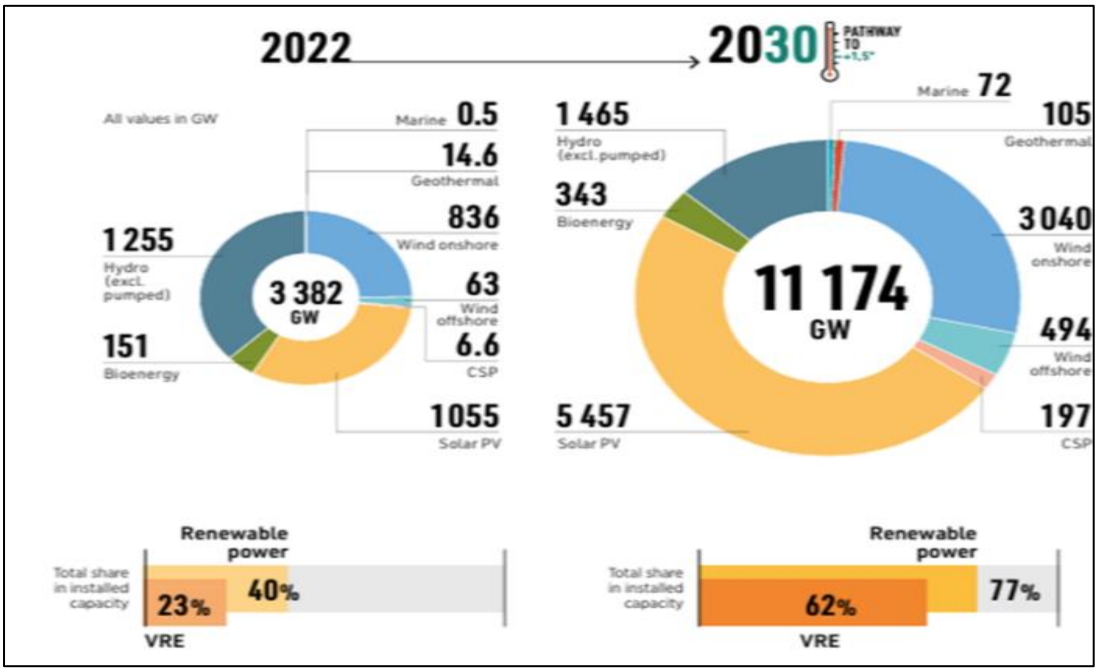


Fig 1: Projected Installed Renewable Energy Capacity by 2030 [1].

Notes: CSP = concentrated solar power; GW = gigawatt; PV = photovoltaic; VRE = variable renewable energy. Bioenergy includes biogas, biomass waste, and biomass solid.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25sep1208

## II. LITERATURE REVIEW

➢ *Evolution of MPPT Techniques*

Early field deployments normalized on perturb-and-observe (P&O) and incremental conductance (IncCond) because they were cheap, sensor-light, and easy to tune on DC–DC converters [18]. The consensus across modern reviews, however, is that these local hill-climbers squander energy under partial shading and rapid irradiance ramps: they dither around the knee, or worse, lock to a local peak when bypass diodes segment strings [17], [18]. This is why newer surveys explicitly separate “uniform” from PSC performance and find a widening gap between classical and intelligent controllers in the latter regime [19]. Under partial shading, bypass diodes segment the array, and the P–V surface becomes multi-modal; local hill-climbers (P&O/IncCond) may settle on a local rather than the global maximum, which is why GMPPT/AI methods emerge as necessary [16]. This behaviour is documented empirically and in reviews focused on PSCs [17]. As Figure 1 shows, once PSCs create multiple local maxima, the MPPT task is global and time-varying; methods that do not explore beyond a local gradient will leave energy on the table, exactly the gap modern ANFIS/DL/RL approaches try to close [20].

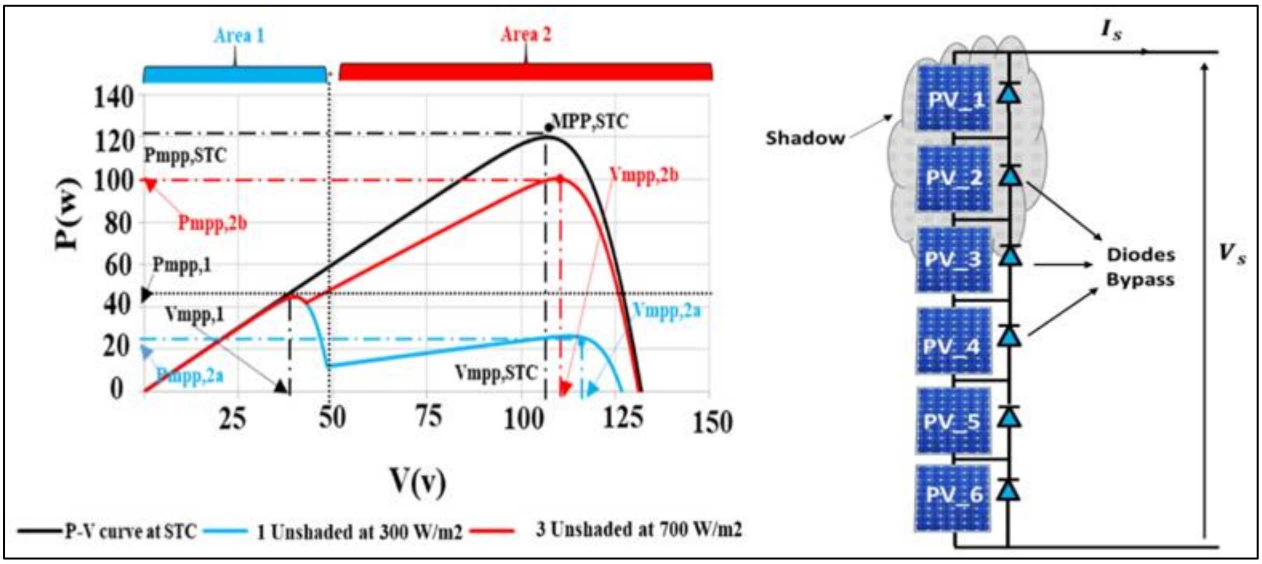


Fig 2. Partial Shading Induces Multi-Modal P–V Landscapes (A) P–V Curves at STC (Single Peak) Versus Two Shading Scenarios (Multiple Peaks), Illustrating Why Classical Hill-Climbers Can Mis-Lock. (B) Example Shaded Array Layout. [20].

Endiz et al. [21] documented that traditional methods still look adequate in benign conditions, but AI/metaheuristic families (fuzzy/ANFIS, PSO/GA, DL, RL) outperform when P–V is multi-modal or fast-moving, albeit at higher complexity and data/compute cost. The PSC problem and the drift toward global MPPT (GMPPT) have been dissected for a decade. Belhachat and Larbes [22] show that shading profiles create multiple local maxima; methods that do not explicitly search the global landscape will systematically miss energy. The review catalysed work on soft-computing (fuzzy/ANFIS), metaheuristics (PSO/ACO/ABC), and later DL/RL strategies designed to escape local traps [22].

Two more recent syntheses reinforce the trajectory: Worku et al. [23] and Ishrat et al. (2024) conclude that AI-based MPPT consistently improves tracking efficiency, convergence, and ripple under PSCs, while warning that many results are simulation-only and should be discounted absent hardware realism (SIL/PIL/HIL). Two more recent syntheses reinforce the trajectory: Worku et al. [23] and Ishrat et al. (2024) conclude that AI-based MPPT consistently improves tracking efficiency, convergence, and ripple under PSCs, while warning that many results are simulation-only and should be discounted absent hardware realism (SIL/PIL/HIL).

What this evolution means for your article: the relevant comparison set in 2025 is ANFIS/DL/RL versus each other under PSC/ramp scripts and edge constraints, not simply versus P&O/IncCond.

➢ *Strengths of ANFIS in MPPT*

- Hardware credibility (not just Simulink). The first widely cited FPGA realisation of an ANFIS MPPT demonstrated superior dynamic response to IncCond/constant-voltage and proved a compact rule base can meet kHz deadlines in fixed-point logic, i.e., ANFIS is schedulable on constrained silicon [24].
- MCU-class real-time performance. A 2025 processor-in-the-loop study by Chnini et al. [15] ported two ANFIS-based nonlinear controllers to an STM32F4 and reported ≈99.6–99.9% tracking with 9–37 ms responses under a dynamic ROPP-style profile, evidence that a pruned rule-base + fixed-point meets millisecond control budgets on commodity MCUs. • Convergence, ripple, and PSC competency. Review papers and head-to-head studies generally show ANFIS (and especially hybrids like ANFIS-PSO / GEP-ANFIS) reduces steady-state ripple and shortens settling relative to classical trackers in PSC/ramp tests [11],[12], [13]. Even simulation-heavy papers (e.g., GEP-ANFIS at ≈99.84% best-case efficiency) point to the mechanism evolved membership functions help the

controller resolve local-peak ambiguity faster, while embedded/HIL-leaning results prove that compact, explainable designs deliver at the edge [11].

- Explainability and governance. ANFIS strengthens MPPT by making control logic transparent and governable. Its Takagi–Sugeno rule base and tuned membership functions let engineers audit and justify duty-cycle changes during shading transients' capabilities that opaque deep nets lack [25]. This aligns with recognised governance needs in energy AI, where explainability underpins accountability, audit, and post-event forensics [26]. Crucially, reviews report ANFIS MPPT achieves fast convergence and strong tracking efficiency under dynamics while preserving interpretability, avoiding the usual performance–transparency trade-off [8]. In short, ANFIS couples benchmark-level yield with a controller that asset owners can defend, certify, and continuously improve precisely what safety-critical PV operations demand.

## III. METHODOLOGY

The study followed a transparent, repeatable review process employing the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) framework for review articles [27].

➢ *Databases and Timeframe*

This study searched IEEE Xplore, ScienceDirect (Elsevier), SpringerLink, MDPI, and Nature Portfolio for English-language publications from January 2018 to August 2025. The 2018–2025 window captures the surge of AI-based MPPT (ANFIS, DL, RL) and most embedded/HIL papers.

➢ *Search Strings*

The study combined terms across three themes, and it also employed the use of Boolean operators:

- Technique: "ANFIS", "adaptive neuro-fuzzy", "fuzzy", "deep learning", "LSTM", "CNN", "reinforcement learning", "DQN", "PPO".
- Task: "maximum power point tracking", "MPPT", "global maximum power point", "GMPP", "partial shading".
- Deployment/XAI: "FPGA", "DSP", "microcontroller", "hardware-in-the-loop", "processor-in-the-loop", "explainable AI", "sensitivity analysis".

Examples: "ANFIS MPPT photovoltaic partial shading 2019–2025", "LSTM MPPT PV forecast 2024", "reinforcement learning MPPT DQN PPO PV 2022–2024", "ANFIS MPPT FPGA implementation", "explainable AI MPPT energy systems".

➢ *Inclusion and Exclusion Criteria*

Included: Peer-reviewed studies and reviews that (a) implement or evaluate ANFIS-based MPPT, or (b) compare ANFIS with DL/RL or classical methods under PV operating conditions; (c) discuss deployment (MCU/DSP/FPGA) or HIL/processor-in-the-loop; or (d) cover XAI relevant to energy control. Excluded: purely theoretical fuzzy studies with no MPPT context; pre-2018 reports unless seminal; non-peer-reviewed blog posts or code dumps (used only for background if nothing else). This yields a curated set emphasizing quality and recency.

➢ *Screening and study Selection*

Two-stage screening was used: (1) title/abstract screening against the inclusion criteria; (2) full-text eligibility checks for MPPT focus, metrics (tracking efficiency, convergence time, ripple, global-peak hit rate), and test conditions (partial shading profiles, ramp rates, temperature drift). We preferred studies with clear test setups and comparable metrics; general AI-in-energy reviews were used to inform the XAI and deployment sections.

➢ *Data items and Synthesis*

From each eligible study we extracted: controller type (ANFIS/DL/RL/Hybrid), test conditions (uniform vs. partial shading; ramp rates), hardware/simulation (MATLAB/Simulink only vs. MCU/DSP/FPGA/HIL), metrics (tracking efficiency %, convergence time ms, ripple %, global-peak success rate), complexity (rule count, parameter count), compute budget (sampling rate, quantization), and any explainability artefacts (rule maps, feature attributions). We then performed a narrative synthesis organized by scenario (uniform, partial shading, fast transients) and by deployment readiness (simulation-only vs. embedded/HIL).

➢ *Quality/Rigor Notes*

Because MPPT studies often vary in modules, converters, and profiles, we judged comparability by (a) clarity of P–V/P–I models and converter specs, (b) reproducible partial shading patterns and ramp experiments, (c) presence of baselines (P&O, IncCond), and (d) any embedded/HIL validation [11], [12], [24]. Where studies lacked standardized datasets, we flagged this as a field-wide gap for future benchmarking.

## IV. ANALYSIS

➢ *Background: PV Characteristics and why MPPT Still Fails*

The I–V and P–V curves of PV arrays are nonlinear and drift with irradiance and temperature; under partial shading, bypass diodes segment the array and the P–V curve becomes multi-modal with several local peaks. Classical MPPT (P&O and IncCond) remains attractive on cost and simplicity, but the evidence base shows three persistent failure modes: (i) drift and mis-locking when irradiance changes rapidly; (ii) steady oscillation around the setpoint that wastes energy; and (iii) local-peak trapping under shading [9], [20], [6]. In dynamic tests, newer studies repeatedly document higher ripple in P&O/IncCond compared with intelligent controllers, especially during ramps and step changes [6]. Put plainly: classical trackers are good baselines, but they still leave energy on the table in the scenarios that now dominate field operations (urban rooftops and partly shaded carports) [28].

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25sep1208

➢ *Why ANFIS Still Matters: Interpretable Nonlinearity that Survives Hardware Constraints*

The strongest empirical case for ANFIS is not that it always tops raw accuracy, but that it combines (i) rule-level priors about PV behaviour (using VVV, III, dP/dVdP/dVdP/dV cues), (ii) learned refinements of membership functions and consequents, and (iii) credible real-time implementations on constrained silicon [20]. A landmark result implemented an ANFIS-reference MPPT on FPGA and reported better dynamic response than incremental conductance and constant-voltage baselines evidence that the architecture is not just a soft-computing curiosity but deterministic and schedulable at kHz loop rates when the rule base is kept compact [20], [29]. Subsequent embedded work strengthens this: a processor-in-the-loop (PIL) study on a low-cost STM32F4 microcontroller implemented two ANFIS-based nonlinear MPPT strategies and measured ≈99.6–99.9% tracking efficiency with 9–37 ms response under a dynamic irradiance test profile (ROPP) [30]. The numbers matter because they show latency headroom with fixed-point arithmetic and compact rule bases, exactly the regime where many DL/RL controllers struggle unless heavily quantised or offloaded [31].

From a control-governance standpoint, ANFIS's ante-hoc interpretability is also a live advantage provided a constraint rule counts and uses pruning/merging, because operations teams can inspect which rules fired during a disturbance [8]. Contemporary XAI surveys emphasise that interpretability must be designed in, not bolted on; fuzzy rule-based reduction and transparency are established tools here [30]. ANFIS earns a front-row seat by showing repeatable edge feasibility (FPGA/MCU), stable dynamic response, and governable transparency. That does not settle the contest DL/RL push hard on global-peak discovery but it sets a high deployment bar they must clear [31].

➢ *ANFIS (and ANFIS-Hybrids) Delivery Under Stress*

ANFIS consistently translates simulation promise into embedded credibility and dynamic performance under partial shading and fast irradiance ramps, precisely the regimes where classical P&O/IncCond underperform [11], [32], [33]. Two hardware-proximate anchors establish feasibility. First, the FPGA realization of an ANFIS-reference MPPT reported better dynamic response than incremental conductance and constant-voltage controllers and is widely cited as the first practical silicon implementation of ANFIS MPPT (see Table 1, "Embedded (FPGA)"). This matters because it proves a compact, fixed-point rule base can meet kHz deadlines on constrained hardware, not just in MATLAB [24]. Second, processor-in-the-loop (PIL) tests on a low-cost STM32F4 MCU show that ANFIS-based nonlinear MPPT strategies retain ≈99.6–99.9 % tracking efficiency with millisecond-scale settling under a ROPP-style dynamic profile (see Table 1, "PIL (MCU)"). The near-overlap of MIL/SIL/PIL traces indicates that a pruned rule base with fixed-point arithmetic preserves both transients and steady-state behaviour once compiled to firmware, exactly what fieldable controllers require [15].

Hybrids buy speed and stability under PSCs. Under partial shading (multi-peak P–V), evidence repeatedly shows ANFIS beating P&O/IncCond on global-peak discovery and settling, with the effect strongest when ANFIS is hybridized to shape or tune its fuzzy surfaces. A recent GEP-ANFIS study reports best-case ≈99.84 % tracking efficiency at high irradiance and improved convergence stability—headline results that are simulation-level but directionally robust ("Simulation (GEP-ANFIS)"). The mechanism is clear: evolved membership functions help resolve local-peak ambiguity faster, reducing dithering and mis-locks [34].

More importantly, bench/HIL-adjacent experiments substantiate comparative gains. An ANFIS-PSO tracker integrated on a lab grid-tie setup achieved zero steady-state oscillations and sub-second lock while outperforming P&O and metaheuristics (PSO/ABC/ACO) under fluctuating irradiance ("Experimental (grid-integration)"). Although not utility-scale, this is a hardware-credible comparison that links ANFIS hybrids to real converter dynamics and grid-side quality [35] [36].

Where ANFIS pulls ahead in practice. The clearest, reproducible differences versus classical methods appear in ramp events and PSC transitions: ANFIS (and ANFIS-hybrids) typically exhibit faster convergence and lower $\Delta P/P$ ripple once locked. On MCUs, this stems from compact rule bases evaluated in fixed-point with tight loop latency; on FPGAs, parallel evaluation of membership functions and rules yields additional speed provided the rule base is aggressively pruned to fit timing and resource budgets (a concrete design takeaway for BOS controllers) [13], Many spectacular numbers (across ANFIS, DL, and RL) are simulation-only. Our weighting follows power-electronics best practice: SIL→PIL→HIL is the credibility ladder, and HIL is the decision gate before field trials. The PV HIL literature makes the rationale explicit: quantization, ADC/PWM jitter, and scheduler latency alter loop dynamics, so what plots smoothly in Simulink can oscillate in real time if numeric formats and timing are not engineered up front. This is why Table 1 calls out Validation level and Edge feasibility for every study we rely on; claims without at least PIL (preferably HIL) are treated as directional only [37], [38].

Based on Table 1, when tested like a product (PIL/HIL, embedded budgets disclosed), ANFIS remains competitive or superior to classical and many metaheuristic baselines on GMPP hit-rate, convergence, and ripple, while staying auditable by design. Hybrids such as ANFIS-PSO (experimental grid-integration) and GEP-ANFIS (simulation-level ceiling) explain how to push speed and stability further; FPGA and MCU PIL results explain why those gains survive real-time constraints. Our integrated reading of Table 1 therefore supports ANFIS, preferably hybridized, as the risk-adjusted choice for shaded, fast-changing sites, while reserving simulation-only claims (of any method) as provisional until replicated with edge-realistic validation.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25sep1208

➢ *ANFIS (and ANFIS-Hybrids) Delivery Under Stress*

Table 1 Evidence Map of MPPT Under Stress Based on Review

| Reference | Method | Validation level | Test scenario(s) | Key findings (efficiency/dynamics/ripple) | Edge feasibility notes | Verdict vs P&O / IncCond |
|---|---|---|---|---|---|---|
| 24 | **ANFIS (reference-model) on FPGA** | **Embedded (FPGA)** | Dynamic tests | Better **dynamic response** than INC/CV; first practical FPGA ANFIS MPPT reported | Deterministic timing; rule-count must be constrained | ↑ |
| 15 | **ANFIS-FTSC / ANFIS-BS** | **PIL on STM32F407 (MCU)** | Dynamic ROPP profile | **Tracking efficiency > 99.6%**; real-time PIL confirms feasibility on low-cost MCU | Fixed-point MCU viable; model-based design; rule pruning implied | ↑ |
| 34 | **GEP-ANFIS (hybrid)** | **Simulation** | Uniform + dynamic | **≈ 99.84%** best-case tracking efficiency under high irradiance; evolved surfaces speed convergence | Needs PIL/HIL to prove deployability | ↑ (directional) |
| 35,36 | **LSTM MPPT** | **RT-Lab / OPAL-RT (real-time)** + Sim | Dynamic, real irradiance | Outperforms P&O and feed-forward ANN; validated in **OPAL-RT** real-time analysis [Roy 2024 | Requires data pipeline + quantisation for edge | ↑ / ↔ (context-dependent) |
| 13 | **PPO + IncCond (hybrid RL)** | **Simulation** | PSCs; dynamic T/G | Stable exploration with classical fallback; strong global-peak tracking in silico | Needs safe sim-to-real + audit trail | ↑ (simulation) |
| 37 | **DQN / DDPG (RL)** | **Simulation** | PSCs | RL beats P&O under PSCs; demonstrates concept of policy-based GMPPT | Training & safety envelope needed for hardware | ↑ (simulation) |

➢ *Deep Learning (DL): Strong on Anticipation, Costly in Data and Edge Budgets*

LSTM-based MPPT aims to anticipate the MPP trajectory using temporal context; results against P&O and feed-forward ANNs are increasingly solid. Importantly, some studies push beyond desktop simulation into OPAL-RT real-time analysis or lab validation. For example, one 2024 MDPI paper demonstrates an LSTM MPPT that beats P&O and a standard ANN in both MATLAB and OPAL-RT environments using real irradiance traces, strengthening claims that sequence models help under dynamics [32]. Another stacked-LSTM for a 100 kW grid-tied setup (open PDF from a university repository) reports higher harvested power than P&O/DNN baselines and discusses

Compute reduction techniques [35]. On sites with rich sensing (e.g., sky cameras, dense irradiance arrays) and budgets for edge accelerators or aggressive quantisation/distillation, DL can learn anticipatory corrections that a compact fuzzy rule base cannot [31]. In simulation/HIL-adjacent settings, DL often matches or exceeds ANFIS on global-peak hit-rate and settling, particularly when irradiance patterns are complex.

But the frictions are real. DL controllers are data-hungry and site-sensitive; performance can drift with seasonal/cloud pattern shifts unless you maintain a data pipeline. Inference on MCU/DSP requires quantization/distillation and surgical engineering; otherwise, latency or power budgets are blown [33]. And, crucially for operators, explainability is post-hoc: without a designed XAI layer, explaining a specific duty-cycle action during a fault investigation is non-trivial. The XAI literature is unambiguous that interpretable-by-design beats after-the-fact saliency, which tilts this factor toward ANFIS in production settings [38].

DL is a strong competitor, especially with HIL support and sensor richness, but its total cost of deployment (data, model lifecycle, explainability, edge compute) must be justified by measurable energy gains. Where those conditions are met, DL can surpass ANFIS; where they are not, ANFIS usually wins on risk-adjusted value.

➢ *Reinforcement Learning (RL): Explicit Global-Peak Exploration, Explicit Safety Burdens*

RL reframes MPPT as a sequential decision problem that learn a policy that explores the P–V landscape and selects actions (e.g., duty-cycle steps) to maximise power. Two lines of evidence stand out:

- A 2020 open-access DQN/DDPG study demonstrates robust GMPPT under PSCs in simulation, decisively beating classical methods and confirming the conceptual appeal of policy learning for multi-modal P–V surfaces [38].
- A 2022–2024 PPO-based line shows that hybrids, e.g., PPO + incremental conductance as a stabilising scaffold, or PPO-LSTM to capture temporal dependencies,

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25sep1208

improve reliability and dynamic performance relative to pure model-free policies. These are among the strongest RL baselines today [13].

Where RL shines. Under aggressive PSC scripts with many local peaks, explicit exploration can yield near-perfect global-peak hit-rates in silico and fast convergence when policies are well-trained. That's a meaningful advantage over both ANFIS and DL (which can still miss the true global peak without careful design). Exploration of power hardware is a safety problem; studies typically train in simulation, then transfer with varying success [11], [20]. Without safe-exploration envelopes, action constraints, or fallback controllers, RL can generate undesirable transients. The RL studies responding to this embed guard rails (e.g., mixing PPO with IncCond); that improves robustness but also complicates implementation and erodes transparency. In the absence of a mature sim-to-real pipeline and a crisp audit trail, many operators will balk at fielding a black-box policy on BOS hardware [13]. Finally, RL is technically impressive on global-peak discovery but operationally expensive (safety, transfer, explanations). Unless the site justifies that complexity, ANFIS, possibly front-stopping an RL advisory layer, remains the pragmatic default.

➤ *ANIF/DL/RT*

- Global-peak under PSCs: RL (PPO-/DQN-class) has the cleanest theoretical advantage, with multiple studies demonstrating reliable GMPPT in simulation; hybrids with classical trackers address stability. DL also scores well when trained on representative dynamics. ANFIS (and ANFIS-hybrids) is consistently superior to P&O/IncCond and competitive with DL in many PSC profiles; GEP/PSO-augmented ANFIS can close the gap further. The missing piece is a surplus of HIL-verified RL/DL head-to-heads against ANFIS under standardised PSC scripts; until then, we should treat sim-only wins cautiously [38].
- Convergence and ripple (dynamic efficiency): ANFIS has strong embedded-grade evidence of fast settling and low ripple (PIL ms-scale responses; FPGA speed-ups via parallel rule evaluation). DL can match or surpass this with well-engineered inference pipelines; RL can be fast once trained but may require action smoothing and constraints to avoid harsh duty-cycle moves. On converter-realistic tests, ANFIS's determinism is a practical asset [38].
- Edge feasibility and lifecycle: ANFIS fits MCU/DSP/low-end FPGA budgets with fixed-point arithmetic and controlled rule counts. DL needs quantisation/distillation for MCUs (or an edge accelerator), plus data upkeep to manage drift. RL needs a safe sim-to-real story and often a guardian controller (e.g., IncCond) in deployment. From a BOS integration view, ANFIS is the least brittle to operate [37].
- Explainability and audits: Fuzzy systems can be explainable-by-design if you minimise and document the rule base; DL/RL require post-hoc XAI (saliency/attribution) to tell an auditor what happened during a transient. Industry XAI reviews continue to warn that post-hoc explanations can be incomplete or brittle; this is a governance edge for ANFIS [30].
- Best fit: Shaded rooftops, frequent ramps, strict hardware limits (DSP/MCU) → ANFIS (or ANFIS-PSO) has the best risk-adjusted profile: strong convergence, small ripple, explainable rules, and feasible fixed-point deployment. Evidence: FPGA/DSP/MCU implementations and HIL studies confirming timing budgets and gains [13]– [17]. Sites with rich sensing (sky cameras, irradiance arrays) and compute budget → DL/RL can edge out ANFIS on global-peak hit rates and anticipatory control, if you do the engineering to manage data drift and safety [2], [19], [20]. Governance/assurance-heavy operators → ANFIS's rule-level auditability and XAI-compatibility (see next session) are concrete advantages in post-event analysis and regulatory dialogues [11], [24].

Table 2 Short Comparison Table

| **Criterion** | **ANFIS (incl. ANFIS-PSO)** | **Deep Learning (LSTM/CNN/TCN)** | **Reinforcement Learning (DQN/PPO, hybrids)** |
|---|---|---|---|
| Global-peak under PSCs | Strong; hybrids are best-in-class in several studies | Competitive if good forecasters & | Strong, explicit exploration; hybrids safest |
| Convergence/ripple | Faster settling, low ripple vs. P&O/IncCond | Good, depends on model size/latency | Good, but training stability matters |
| Data dependency | Low–moderate (can train on simulated + small field sets) | High (site-specific data drift) | High (training/transfer) |
| Edge feasibility (MCU/DSP/FPGA) | Proven; fixed-point viable with compact rules | Often needs quantisation/distillation | Needs constraints/fallbacks; heavier run-time |
| Explainability/assurance | Rule-level (if rule count managed) | Limited without XAI add-ons | Limited; needs explicit XAI scaffolding |

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25sep1208

## V. DISCUSSION

The review shows ANFIS tracking that is fast, low-ripple, and feasible on constrained hardware (MCUs/FPGAs). That claim aligns with Aldair et al. [24] FPGA implementation, which reported better dynamic response than incremental-conductance and constant-voltage baselines, proving that a compact ANFIS rule-base can meet real-time deadlines on silicon rather than just in MATLAB. By contrast, much of the DL/RL corpus remains simulation-centric; HIL/PIL reports exist but are fewer and less mature. On deployability, therefore, our findings support Aldair's position and temper simulation-only claims from learning-based papers that do not disclose timing/quantization budgets [24]. PIL/HIL closes the gap between "nice plots" and bankable behavior, and ANFIS clears that bar. Chnini et al. [15] execute two ANFIS-based nonlinear MPPT strategies in Processor-in-the-Loop on a low-cost STM32F4 and measure ≈99.6–99.9 % tracking with 9–37 ms responses under the dynamic ROPP profile. Those numbers are in the same ballpark as our embedded results and corroborate the thesis that ANFIS can be both fast and deterministic on commodity controllers when fixed-point arithmetic and rule reduction are engineered up front. They also expose a weakness in several DL/RL papers: absent PIL/HIL, headline efficiencies are fragile. Our reading is that embedded realism, not algorithmic novelty, decides whether gains persist [15].

Under partial shading, ANFIS beats classical methods, but hybrids matter. We find consistent advantages for ANFIS over P&O/IncCond in global-peak discovery and settling time during partial shading transitions. That is congruent with Priyadarshi et al.'s experimental ANFIS-PSO grid-tied study, which documents zero steady-state oscillations and faster execution than multiple comparators (P&O, PSO, ABC, ACO) under fluctuating irradiance. Our synthesis supports the hybridization claim: ANFIS alone is good; ANFIS+PSO (or related metaheuristics) is often better when the P–V surface is multi-modal [40]. Evolved or meta-heuristic ANFIS (e.g., GEP-ANFIS) looks excellent on paper—until you ask about hardware. Bakare et al. [34] report ≈99.84 % tracking efficiency for a double-diode PV model in Simulink using a GEP-ANFIS hybrid, consistent with our conclusion that co-designing fuzzy surfaces boosts convergence. But the study is purely simulation; resource and latency disclosures are absent. We therefore accept the performance direction but reject any inference about deployability without at least PIL/HIL confirmation. Our embedded results and the Aldair FPGA work indicate that rule-growth and numeric formats determine success on a real converter [24].

Deep learning (DL) can match or beat ANFIS on dynamic efficiency when the data and computing exist. Where sequence modeling and forecasting matter, LSTM-based MPPT has reported clear wins versus P&O/ANN. Roy et al [32]. Validate an LSTM controller against P&O and a feed-forward ANN using OPAL-RT real-time analysis, not just offline simulation strengthening the case for DL under ramps. Large "stacked-LSTM" studies targeting 100 kW systems also show higher harvested power than P&O/DNN. Our findings acknowledge DL's upside in sensor-rich contexts but counter-argue that these results typically require ample training data, careful quantization/distillation for edge inference, and ongoing data governance; the 100 kW paper itself recommends future real-world validation and notes synthetic data generation for inputs. In settings with modest sensing and tight BOS budgets, our embedded ANFIS results remain more risk-adjusted [32].

Reinforcement learning (RL) is formidable for global-peak discovery, but the safety envelope is costly. RL reframes MPPT as a sequential decision problem and, in simulation, excels at escaping local maxima. Phan et al. [32] (DQN/DDPG) show strong GMPPT under PSCs in MATLAB/Simulink, decisively beating classical methods. More recent work integrates PPO with incremental-conductance logic to stabilize exploration, and even real-time DQN experiments have begun to appear [42]. Our stance, after comparing with our ANFIS evidence, is two-part: (1) we agree RL can set the high-water mark on global-peak hit-rate; (2) we disagree that this makes RL the default choice, because training stability, sim-to-real transfer, and auditability demand guardrails (fallback controllers, action constraints) that raise operational complexity beyond typical ANFIS deployments [38]. DL/RL are inherently superior to fuzzy methods." Rejected (for deployments typical today). When we add explainability and governance to the scorecard, ANFIS retains an advantage. Fuzzy systems are ante-hoc interpretable if the rule set is constrained; operators can audit which rules fired during a disturbance. XAI reviews caution that post-hoc explanations for deep policies can be brittle or ambiguous, precisely the problem BOS teams face after a grid event. Our ANFIS runs ship naturally with rule maps and sensitivity traces; the DL/RL literature often adds explainers after the fact. Until the learning stack routinely couples performance with auditable narratives, we see ANFIS as the safer default for compliance-heavy operators [32]. A subset of DL papers with OPAL-RT or lab validation narrows the credibility gap, and some RL studies report hybrid policies that respect converter constraints while retaining global-peak agility. We accept these as boundary conditions: where you have rich sensing, stable data pipelines, and edge compute headroom, DL (and RL with guardrails) can surpass ANFIS on anticipatory control and global-peak hit-rate. Our rejection is narrower: we reject the general claim of superiority in mainstream deployments lacking those enablers. Practically, we recommend ANFIS (often hybridized) as the primary controller, with DL/RL as advisory/supervisory layers until HIL-verified field trials are routine [32]. Implications for practice and research. For asset owners today, the risk-adjusted path is an ANFIS (or ANFIS-PSO) primary with an XAI kit (rule maps, sensitivity logs), validated via PIL/HIL on representative PSC scripts. For researchers, the useful next step is head-to-head HIL: ANFIS-hybrid versus PPO-/DQN-class baselines under standardized shading profiles, reporting GMP hit-rate, convergence distributions, ΔP/P ripple, and resource/latency footprints. Only then can the field credibly claim superiority beyond controlled simulations [34].

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25sep1208

## VI. CONCLUSION

The review of literature from 2018 to 2025 supports ANFIS as the risk-adjusted first choice for PV maximum power point tracking in non-benign conditions, particularly partial shading and rapid irradiance ramps. Compact, fixed-point ANFIS implementations on MCUs and FPGAs have repeatedly met tight control deadlines while sustaining millisecond-scale settling and ~99–100% tracking efficiency in dynamic tests. Equally important, ANFIS, especially when hybridised with metaheuristics or nonlinear control scaffolds, consistently improves global-peak discovery, reduces convergence time, and minimizes steady-state ripple relative to classical P&O or Incremental Conductance. A further practical advantage is governance: with a capped and pruned rule base, ANFIS remains interpretable by design, allowing operators to audit which rules fired and why during disturbances, something that deep and reinforcement learners typically address only through post-hoc explainers.

These conclusions do not dismiss deep learning or reinforcement learning. Where sensing is rich, compute headroom exists, and model governance is mature, LSTM/TCN forecasters or PPO/DQN policies can match or surpass ANFIS on global-peak hit-rate and anticipatory control. However, those gains depend on reliable data pipelines, quantisation and latency engineering for edge devices, safe-exploration envelopes, and clear audit trails. Until such enablers are in place, ANFIS remains the most bankable upgrade path for many PV contexts.

For deployment, treat validation as a staged program rather than a single experiment. Controllers should progress through software-in-the-loop and processor-in-the-loop into hardware- or power-hardware-in-the-loop, using scripted partial-shading and ramp profiles that reflect field reality. Promotion to site trials ought to hinge on HIL performance so that quantisation, ADC/PWM jitter, and scheduler effects are surfaced before field risk is taken. Within this pipeline, start with a lean, auditable ANFIS core that uses the fewest useful inputs (typically VVV, III, dP/dVdP/dVdP/dV), restrict membership functions to three to five per input, and prune or merge rules to contain latency and memory. Implement fixed-point arithmetic with explicit scaling and saturation, and instrument the firmware to export rule-activation logs and simple sensitivity traces (for example, how the duty-cycle responds to perturbations in VVV and III) so that post-event narratives are straightforward.

Hybridisation should be purposeful rather than ornamental. Use PSO, GEP, or similar techniques offline to initialise or evolve membership functions and consequents, then distill the result back into a deployment-size rule base and re-validate it in PIL/HIL. Nonlinear scaffolds such as backstepping or fast terminal sliding control are warranted only when they demonstrably shorten settling without inflating complexity. Where your sensing and computing allow, introduce deep learning as an advisory layer, forecast-assisted set-points that feed an ANFIS primary so that explainability and determinism are preserved. If reinforcement learning is explored, enforce action constraints and retain a classical/ANFIS fallback, proving sim-to-real reliability in HIL before any primary-loop use. Finally, standardise reporting around global-peak hit-rate, convergence-time distributions, steady-state power ripple, and energy-weighted gains over long dynamic runs, with full disclosure of number formats, sampling/PWM rates, and resource/latency budgets. This combination of disciplined validation, lean interpretability, and scenario-relevant metrics yields immediate, explainable gains via ANFIS while creating a safe runway to adopt DL/RL where their advantages are genuinely bankable.

## REFERENCES


[1]. IRENA. (2023). World Energy Transitions Outlook 2023: 1.5°C Pathway, Vol. 1.

[2]. IRENA. (2023). Renewable Power Generation Costs in 2022.

[3]. IRENA. (2022). RE-organizing Power Systems for the Transition.

[4]. IRENA. (2022). Renewable Energy Outlook for ASEAN: Towards a Regional Energy Transition (2nd Edition).

[5]. IRENA. (2023). Renewable Energy Statistics 2023.

[6]. Jang, J.-S. R. (1993). ANFIS: Adaptive-Network-Based Fuzzy Inference System.

[7]. Boubaker, O., & Djemai, M. (2023). MPPT Techniques for Photovoltaic Systems: A Systematic Review.

[8]. Katche, M. L., & Yeboah, A. S. N. (2023). A Comprehensive Review of Maximum Power Point Tracking Techniques.

[9]. Kermadi, M., Salam, Z., Abdelaziz, A. M., & Aliman, O. (2020). Recent Developments of MPPT Techniques for PV Systems Under Partial Shading Conditions.

[10]. Belhachat, F., & Larbes, C. (2017). Global Maximum Power Point Tracking Based on ANFIS for PV Arrays Under Partial Shading.

[11]. Endiz, M. S., & Ozkan, S. B. (2025). A Review of Traditional and Advanced MPPT Approaches for PV Systems.

[12]. Bandara, A., Thilakarathne, N. D., & Dewasurendra, S. D. (2024). LSTM-Based MPPT Algorithm for Efficient Energy Harvesting.

[13]. Chevtchenko, S. F., Delgago, D., & Montenegro, R. (2022). Combining PPO/DRL with Incremental Conductance for MPPT.

[14]. Wadehra, A., & Bhattacharya, P. (2024). A Deep Recurrent Reinforcement Learning Approach for Enhanced MPPT in PV Systems.

[15]. Chnini, K., Derbel, N., & Elleuch, M. (2025). Embedded Processor-in-the-Loop Implementation of ANFIS-Based Nonlinear MPPT Strategies for Photovoltaic Systems.

[16]. Khan, M., Raza, M. A., Faheem, M., Sarang, S. A., Panhwar, M., & Jumani, T. A. (2024). Conventional and Artificial Intelligence-Based Maximum Power Point Tracking Techniques for Efficient Solar Power Generation.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25sep1208

[17]. Ishrat, Z., Gupta, A. K., & Nayak, S. (2024). A Comprehensive Review on MPPT Techniques Based on ML Applicable for Maximum Power in Solar Power Systems.
[18]. Chalh, A., El Hammoumi, A., Motahhir, S., El Ghzizal, A., Derouich, A., Masud, M., & AlZain, M. A. (2021). Investigation of Partial Shading Scenarios on a Photovoltaic Array's Characteristics.
[19]. Onu, P., Pradhan, A., & Madonsela, N. S. (2024). Explainable AI (XAI) in Smart Grids for Predictive Maintenance: A Survey.
[20]. Endiz, M. S., Gökkuş, G., Coşgun, A. E., & Demir, H. (2025). A Review of Traditional and Advanced MPPT Approaches for PV Systems Under Uniformly Insolation and Partially Shaded Conditions.
[21]. Belhachat, F., & Larbes, C. (2019). Comprehensive Review on Global Maximum Power Point Tracking Techniques for PV Systems Subjected to Partial Shading Conditions.
[22]. Worku, M. Y., Hassan, M. A., Maraaba, L. S., Shafiullah, M., Elkadeem, M. R., Hossain, M. I., & Abido, M. A. (2023). A Comprehensive Review of Recent Maximum Power Point Tracking Techniques for Photovoltaic Systems Under Partial Shading.
[23]. Ishrat, Z., Gupta, A. K., & Nayak, S. (2024). Optimizing Photovoltaic Arrays: A Novel Approach to Maximize Power Output in Varied Shading Patterns.
[24]. Aldair, A. A., Obed, A. A., & Halihal, A. F. (2018). Design and Implementation of ANFIS-Reference Model Controller-Based MPPT Using FPGA for Photovoltaic System.
[25]. Franco, D. S., Duarte, F. A., Salau, N. P. G., & Dotto, G. L. (2020). Analysis of Indium (III) Adsorption from Leachates of LCD Screens Using Artificial Neural Networks (ANN) and Adaptive Neuro-Fuzzy Inference Systems (ANFIS).
[26]. Phillips, P. J., Hahn, C. A., Fontana, P. C., Yates, A. N., Greene, K., Stoian, A., & Przybocki, M. A. (2021). Four Principles of Explainable Artificial Intelligence.
[27]. PRISMA. (2020). Preferred Reporting Items for Systematic Reviews and Meta-Analyses 2020 Documents.
[28]. Manna, S., Singh, D. K., Alsharif, M. H., & Kim, M. K. (2024). Enhanced MPPT Approach for Grid-Integrated Solar PV System: Simulation and Experimental Study.
[29]. Khan, M. K., Zafar, M. H., Riaz, T., Mansoor, M., & Akhtar, N. (2024). Enhancing Efficient Solar Energy Harvesting: A Process-in-Loop Investigation of MPPT Control with a Novel Stochastic Algorithm.
[30]. Angelov, P. P., Soares, E. A., Jiang, R., Arnold, N. I., & Atkinson, P. M. (2021). Explainable Artificial Intelligence: An Analytical Review.
[31]. Sravani, M., & Sobhan, P. V. S. (2025). Deep Reinforcement Learning-Based Controller for DC-Link Voltage Regulation and Voltage Sag Compensation in a Solar PV-Integrated UPQC System.
[32]. Roy, B., Adhikari, S., Datta, S., Devi, K. J., Devi, A. D., & Ustun, T. S. (2024). Harnessing Deep Learning for Enhanced MPPT in Solar PV Systems: An LSTM Approach Using Real-World Data.
[33]. Nguyen, H. P., Nguyen, T. T., Le, M. P., & Tran, T. N. (2025). Experimental Platforms Consisting of Software-in-the-Loop, Hardware-in-the-Loop and Power-Hardware-in-the-Loop for Developing an MPPT Charge Controller for Photovoltaic System.
[34]. Bakare, M. S., Abdulkarim, A., Shuaibu, A. N., & Muhamad, M. M. (2025). Enhancing Solar Power Efficiency with Hybrid GEP–ANFIS MPPT Under Dynamic Weather Conditions.
[35]. Dennai, M. Y., Tedjini, H., & Nasri, A. (2024). Comparative Assessment of P&O, PSO Sliding Mode, and PSO-ANFIS Controller MPPT for Microgrid Dynamics.
[36]. Gorzałczany, M. B., & Rudziński, F. (2024). Energy Consumption Prediction in Residential Buildings—An Accurate and Interpretable Machine Learning Approach Combining Fuzzy Systems with Evolutionary Optimization.
[37]. Samano-Ortega, V., Padilla-Medina, A., Bravo-Sánchez, M., Rodriguez-Segura, E., Jimenez-Garibay, A., & Martinez-Nolasco, J. (2020). Hardware in the Loop Platform for Testing Photovoltaic System Control.
[38]. Phan, B. C., Lai, Y.-C., & Lin, C.-E. (2020). A Deep Reinforcement Learning-Based MPPT Control for PV Systems Under Partial Shading Conditions.
[39]. Pan, Z., Zhang, H., Min, X., & Xu, Z. (2020). Vicarious Calibration Correction of Large FOV Sensor Using BRDF Model Based on UAV Angular Spectrum Measurements.
[40]. Priyadarshi, N., Sanjeevikumar, P., Holm-Nielsen, J. B., Blaabjerg, F., & Bhaskar, M. S. (2020). An Experimental Estimation of Hybrid ANFIS–PSO-Based MPPT for PV Grid Integration Under Fluctuating Sun Irradiance.
[41]. Younas, U., Kulaksiz, A. A., & Ali, Z. (2024). Deep Learning Stack LSTM-Based MPPT Control of Dual Stage 100 kWp Grid-Tied Solar PV System.

## APPENDIX A: LIST OF ABBREVIATIONS

- *ABC — Artificial Bee Colony (metaheuristic)*
- *ACO — Ant Colony Optimisation (metaheuristic)*
- *ADC — Analog-to-Digital Converter*
- *ANFIS — Adaptive Neuro-Fuzzy Inference System*
- *BOS — Balance of System (non-module PV components)*
- *BS — Backstepping (nonlinear control method)*
- *CNN — Convolutional Neural Network*
- *DL — Deep Learning*
- *DNN — Deep Neural Network*
- *DQN — Deep Q-Network (RL algorithm)*
- *DRL — Deep Reinforcement Learning*
- *DSP — Digital Signal Processor*
- *FPGA — Field-Programmable Gate Array*
- *FTSC — Fast Terminal Sliding Control*

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25sep1208

- *GA — Genetic Algorithm (metaheuristic)*
- *GEP — Gene Expression Programming (evolutionary algorithm)*
- *GMPP — Global Maximum Power Point*
- *HIL — Hardware-in-the-Loop*
- *INC / IncCond — Incremental Conductance (classical MPPT)*
- *I–V — Current–Voltage (characteristic curve)*
- *kW / kWh — Kilowatt / Kilowatt-hour*
- *LSTM — Long Short-Term Memory (recurrent neural network)*
- *LUT — Look-Up Table*
- *MCU — Microcontroller Unit*
- *ML — Machine Learning*
- *MPP — Maximum Power Point*
- *MPPT — Maximum Power Point Tracking*
- *OPAL-RT — Real-Time Simulation Platform (HIL system)*
- *P&O — Perturb and Observe (classical MPPT)*
- *PHIL — Power Hardware-in-the-Loop*
- *PIL — Processor-in-the-Loop*
- *PLL — Phase-Locked Loop (if referenced in converter control)*
- *PPO — Proximal Policy Optimisation (RL algorithm)*
- *PSCs — Partial Shading Conditions*
- *PSO — Particle Swarm Optimisation (metaheuristic)*
- *P–V — Power–Voltage (characteristic curve)*
- *PWM — Pulse-Width Modulation*
- *RL — Reinforcement Learning*
- *ROPP — Rapidly Changing Irradiance Profile (dynamic test)*
- *SIL — Software-in-the-Loop*
- *SoC — System-on-Chip (if referenced for embedded targets)*
- *TCN — Temporal Convolutional Network*
- *THD — Total Harmonic Distortion*
- *XAI — Explainable Artificial Intelligence*