Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Comparative Analysis of Rule-Based, Fuzzy Logic and PSO-Tuned Energy Management Strategies for a Standalone Hybrid Renewable Microgrid: A Simulation Study

Timileyin Opeyemi Akande

akande.timileyin@tech-u.edu.ng

Abiola Ajimobi Technical University Ibadan

Claret Nnaemeka Ibeawuchi

Federal University of Technology Owerri

Monsuru Olalekan Abdullahi

University of Lagos

Osagie Gideon Agbonigiarhuoyi

University of Benin

Edwin Nduka Nwabueze

Nnamdi Azikiwe University

Oluwapelumi Abdulaziz Akinola

University of Ibadan

Aliyu Salisu Nura

Bayero University Kano

Research Article

Keywords: hybrid renewable energy systems, energy management system, Mamdani fuzzy logic, particle swarm optimization, battery degradation, levelized cost of energy

Posted Date: November 18th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11389471/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Comparative Analysis of Rule-Based, Fuzzy Logic and PSO-Tuned Energy Management Strategies for a Standalone Hybrid Renewable Microgrid: A Simulation Study

Abstract

Energy management systems (EMS) determine how much diesel, battery wear and capital an off-grid hybrid renewable microgrid consumes, yet rule-based, fuzzy and metaheuristic strategies are rarely compared on one plant with common safeguards, aging-aware economics and statistical testing. This simulation study compares an expert-set rule-based controller, a 27-rule Mamdani fuzzy logic controller (FLC) in untuned and particle swarm optimization (PSO)-tuned form, and a PSO-tuned version of the rule-based threshold policy on a photovoltaic (PV)–wind–battery–diesel microgrid, with equal tuning budgets for the tuned strategies. Hourly weather and load are synthetic, generated from assumed West-African coastal climatology; no field data were used. A Python framework couples component models, diesel forced outages, rainflow- and temperature-aware battery aging, and net present cost (NPC) and levelized cost of energy (LCOE) accounting. Strategies are tested out of sample on 100 simulated years, 12 PV–storage sizings, three stress scenarios and three economic sensitivities. In the baseline plant (100 kWp PV, 400 kWh battery), PSO tuning lowered mean LCOE from 0.467 to 0.429 $/kWh (8.2%) and diesel consumption by 28% relative to the expert-set controller; the PSO-tuned FLC reached 0.431 $/kWh (7.7% lower) with less unserved energy, whereas the untuned FLC cost 1.8% more than the baseline. The PSO saving did not depend on the number of training years (1–10). Savings grew from 0.8% for generously sized plants to 16.1% for undersized ones. Results are conditional on the assumed climate, cost and aging parameters and require validation with measured data.

Keywords: hybrid renewable energy systems, energy management system, Mamdani fuzzy logic, particle swarm optimization, battery degradation, levelized cost of energy

1 Introduction

1.1 Background and motivation

Solar photovoltaic (PV) and wind generation are expanding rapidly as part of global decarbonization efforts, and hybrid renewable energy systems (HRES) that combine several sources with storage and, in isolated sites, a diesel backup are now a standard option for rural electrification and island microgrids (Roy et al. 2022; Talaat et al. 2025; Ukoba et al. 2024). Because irradiance and wind speed fluctuate on time scales from seconds to seasons, the value of an HRES depends heavily on its energy management system (EMS), the control layer that decides every hour how generation, storage and the backup generator are dispatched to serve the load (Ahmad et al. 2023; Eyimaya and Altin 2024; Shafiullah et al. 2022).

In off-grid tropical settings the EMS governs reliability (the battery state of charge (SoC) held in reserve determines how well the plant survives generator outages and low-resource periods), operating cost (remote diesel is expensive and generator efficiency depends strongly on loading) and battery lifetime (cycling depth, throughput and high ambient temperature accelerate capacity fade and drive replacement cost) (Dufek et al. 2022; Kafando et al. 2024; Khan et al. 2022). These effects interact, so an EMS that looks attractive on one indicator can be poor overall.

1.2 Related work and research gap

The EMS literature is commonly organized into deterministic (rule-based), intelligent (fuzzy, neural, learning-based) and optimization-based (metaheuristic, mathematical programming) families (Battula et al. 2021; Elmouatamid et al. 2021; Salehi et al. 2022). Rule-based control is valued for its transparency and negligible computational burden. Fuzzy logic controllers (FLC) encode expert knowledge through linguistic rules and have been applied to battery, demand-side and stand-alone hybrid-system management (Al Sumarmad et al. 2022; Ibrahim et al. 2023; Kechida et al. 2024; Magsumbol et al. 2025). Particle swarm optimization (PSO) and related metaheuristics are widely used for sizing and controller tuning (Bouaouda and Sayouti 2022; Giedraityte et al. 2025; Zaini et al. 2023), while reinforcement learning is emerging (Ramesh et al. 2025). Reviews report trade-offs between optimality, complexity and robustness (Ejiyi et al. 2025; Khan et al. 2022; Shafiullah et al. 2022).

Despite this breadth, three limitations recur. (i) Studies usually evaluate one or two strategies on different plants, load profiles and economic assumptions, so reported improvements cannot be compared directly (Kumar et al. 2025; Menzri et al. 2024; Riaz et al. 2024). (ii) Battery aging is often reduced to a throughput proxy or ignored, although cycling depth and temperature are central in hot climates (Dufek et al. 2022). (iii) Results are typically shown for one weather and load realization without out-of-sample testing or significance analysis, even though load-profile choice alone can change optimization outcomes materially (Jurasz et al. 2022). A like-for-like, statistically supported benchmark of rule-based, fuzzy and PSO-based EMS for a diesel-backed tropical microgrid is therefore still missing. The benchmark separates the effect of tuning from that of controller structure by tuning both the threshold policy and the FLC with the same PSO budget (Section 4.10). This paper addresses that gap with: a unified benchmark of three EMS archetypes, with the FLC also in a PSO-tuned form, on one standalone PV–wind–battery–diesel plant with shared component models, protection logic and economics; aging-aware economics (rainflow counting, a depth-of-discharge life curve and an Arrhenius temperature factor) feeding the net present cost (NPC); out-of-sample testing on 100 stochastic years with paired non-parametric tests and bootstrap confidence intervals; a regime map from a 12-point sizing sweep, three stress scenarios and three economic sensitivities; and separate accounting of online and offline computational cost, with a reproducible Python implementation.

2 Literature review

2.1 Architecture of hybrid renewable energy systems

A standalone HRES combines PV arrays, wind turbines, a battery energy storage system (BESS) and a dispatchable diesel generator to supply an isolated load (Kafando et al. 2024; Roy et al. 2022). PV is interfaced through maximum-power-point-tracking (MPPT) converters, wind turbines through AC or rectified DC links, and lithium-ion batteries are now favored over lead–acid for their efficiency and cycle life (Sayed et al. 2023). Buses may be DC-coupled, AC-coupled or hybrid-coupled; AC coupling, adopted here, gives modularity and lets a single bidirectional inverter handle the battery (Bihari et al. 2021; Salehi et al. 2022). Fig. 1 shows the configuration studied.

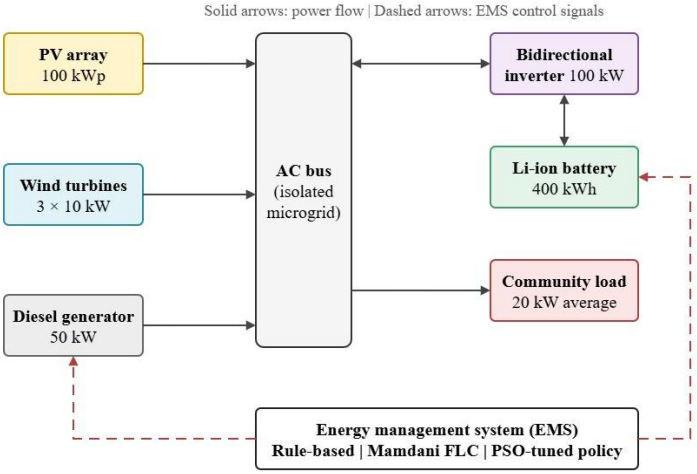


Fig. 1 Configuration of the standalone PV–wind–battery–diesel microgrid. Solid arrows: power flow; dashed arrows: EMS control signals

2.2 Energy management strategies

EMS strategies may be grouped as deterministic, heuristic/metaheuristic and artificial-intelligence-based (Ahmad et al. 2023; Eyimaya and Altin 2024; Shafiullah et al. 2022). Deterministic strategies apply predefined if–then rules on SoC, load and generation; heuristic strategies such as PSO search a decision space for cost-optimal setpoints or parameters; intelligent strategies such as fuzzy inference handle imprecise or nonlinear inputs through linguistic rules. Reinforcement learning needs large training datasets and careful validation (Ramesh et al. 2025; Talaat et al. 2025). Practical implementations include supervisory and non-integer-order control of DC microgrids (Aziz et al. 2021; Soliman et al. 2021), supervisory control and data acquisition (SCADA)-based energy management in a real microgrid (Kermani et al. 2021) and PV control for PV/wind/fuel-cell systems with storage (Gulzar et al. 2023).

2.3 Rule-based, fuzzy and PSO-based EMS

Rule-based EMS dispatch sources in a fixed priority order: renewables serve the load first, surplus charges the battery up to an upper SoC limit, deficits are met by battery discharge down to a lower limit, and the generator starts when the battery cannot cover the deficit (Eyimaya and Altin 2024; Salehi et al. 2022). Thresholds are set by experience rather than by an objective, so these controllers can run the generator at poor loading or cycle the battery deeply (Ahmad et al. 2023), yet they remain the benchmark for intelligent strategies (Shafiullah et al. 2022). Fuzzy EMS map crisp measurements (SoC, load, renewable power) to fuzzy sets, evaluate a rule base and defuzzify to a dispatch command; the Mamdani architecture, with fuzzy consequents and centroid defuzzification, is favored for interpretability, whereas Takagi–Sugeno variants suit integration with optimization (Al Sumarmad et al. 2022; Ibrahim et al. 2023). PSO is a population-based metaheuristic in which each particle updates its velocity from its own best position and the swarm's best position; in HRES it has been applied to sizing, parameter identification and dispatch (Adenuga and Krishnamurthy 2023; Bouaouda and Sayouti 2022; Singh et al. 2022), and its use in demand-side management is reviewed by Zaini et al. (2023). Its weaknesses are sensitivity to settings and the cost of repeated simulation, so practical schemes often tune controller parameters offline (Giedraityte et al. 2025; Zaini et al. 2023).

2.4 Comparative studies and research gap

Most of the studies cited above examine one strategy family, or compare a small number of variants, on plants and metrics that differ from study to study; they are therefore complementary to, rather than replacements for, a common-plant benchmark.

The reviewed work shows what each family can do but not how they rank when plant, protection logic, economics and weather are held constant. Additionally, the benefit of a more sophisticated EMS is expected to depend on how tightly the plant is sized, a dependence that is rarely mapped. The present study therefore asks: (RQ1) How do rule-based, fuzzy and PSO-tuned EMS rank on cost, reliability, diesel use and battery life under identical plant and weather conditions, and how much of the difference is due to tuning rather than controller structure? (RQ2) Are the differences consistent across independent simulated years? (RQ3) How do they change with plant sizing, economic assumptions and stress conditions? (RQ4) What are the online and offline computational costs?

3 Materials and methods

3.1 Research framework

Fig. 2 outlines the workflow. Hourly weather and load years are generated, converted to PV, wind and load power by the component models, and dispatched by each strategy through a shared plant layer that enforces SoC limits, power balance and a generator minimum-load protection rule. The simulated state-of-charge trace is post-processed into battery lifetime, which enters the economic model. All steps are implemented in Python (NumPy, SciPy, Matplotlib).

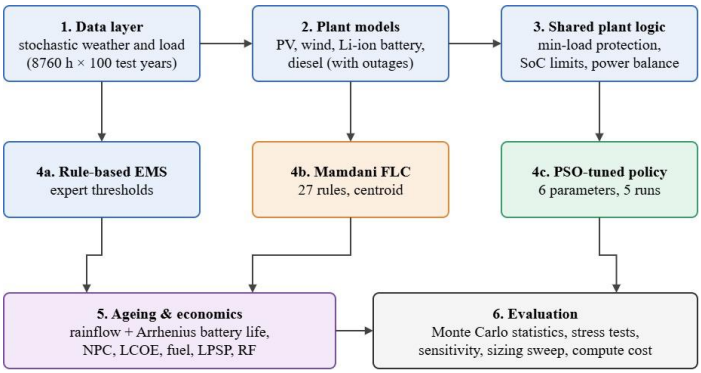


Fig. 2 Research framework of the benchmarking study

3.2 Weather, load and outage data

Data provenance. No measured site series were used in this study. Hourly weather and load years are generated by a documented stochastic generator calibrated to assumed monthly climatology typical of the Gulf-of-Guinea coast (Table 1); the monthly values are modeling assumptions, not measurements, so the results describe a representative synthetic site rather than a specific one. Synthetic data were chosen because they provide many independent years, controlled stress scenarios and paired comparisons with common random numbers, which a single measured year cannot; the price is that realism is limited by the generator assumptions. A public hourly dataset can be substituted with a one-line change, which is the first recommended step before field deployment (Section 4.12).

The generator produces (i) clear-sky global horizontal irradiance (GHI) from solar geometry at 6.5° N, modulated by a daily clearness state with first-order autoregressive persistence ($\rho = 0.65$) and hourly cloud transients, then rescaled month by month to the target climatology with a ±4% inter-annual perturbation; (ii) ambient temperature from the monthly mean, a diurnal cosine and a cloudiness-dependent amplitude; (iii) 10 m wind speed from a Weibull distribution (shape $k = 2$) with a seasonal scale, an afternoon sea-breeze factor and hourly persistence through a Gaussian copula ($\rho = 0.96$); (iv) a community load with morning and evening peaks, a weekend factor, daily and hourly noise and a temperature sensitivity of 1.2% per °C, scaled to a 20 kW mean; and (v) diesel-generator forced outages from a two-state Markov chain (mean time between failures, MTBF, 400 h; mean time to repair, MTTR, 12 h; availability ≈ 97%). An example generated year is shown in Fig. 3.

Table 1 Assumed monthly climatology used to calibrate the stochastic generator

Month	J	F	M	A	M	J	J	A	S	O	N	D
GHI (kWh/m²/day)	5.0	5.4	5.5	5.3	4.9	4.2	3.8	3.9	4.1	4.6	5.0	4.9
Ambient temp. (°C)	28.0	29.0	29.5	29.0	28.0	26.5	25.8	25.5	26.2	27.0	28.0	27.8
Wind at 10 m (m/s)	3.6	3.8	4.0	3.8	4.0	4.4	4.6	4.6	4.2	3.8	3.4	3.4

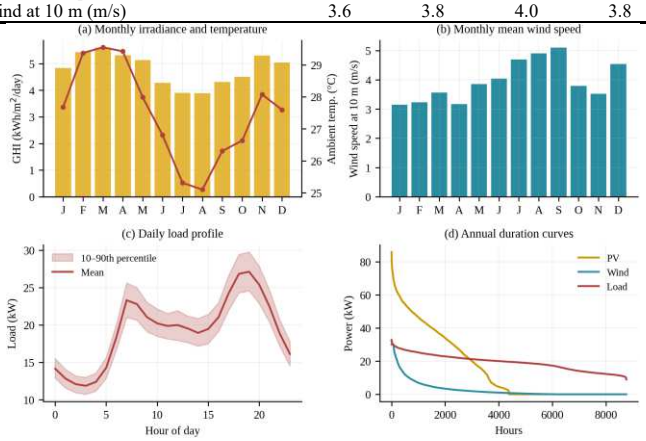


Fig. 3 Example generated year: (a) monthly irradiance and temperature, (b) monthly wind speed, (c) daily load profile with 10–90th percentile band, (d) annual duration curves of PV, wind and load

3.3 Component models

Photovoltaic array. The AC-equivalent PV output follows the irradiance–temperature model (Eq. 1):

$$P_{PV}(t) = P_{PV}^{r}\,\eta_{der}\,\frac{G(t)}{G_{ref}}\left[1+\alpha_T\left(T_c(t)-T_{ref}\right)\right] \tag{1}$$

where P_{PV}^{r} is the rated power, η_{der} = 0.86 a derating factor for soiling and wiring, G_{ref} = 1000 W/m², α_T = −0.4%/°C, T_{ref} = 25 °C and the cell temperature (Eq. 2) is

$$T_c(t) = T_a(t) + \frac{\mathrm{NOCT}-20}{800}\,G(t) \tag{2}$$

with a nominal operating cell temperature (NOCT) of 45 °C and ambient temperature T_a.

Wind turbines. Wind speed is extrapolated to hub height with the power law (Eq. 3),

$$v(t) = v_{10}(t)\left(\frac{h_{hub}}{h_{ref}}\right)^{\beta} \tag{3}$$

with β = 0.14, h_{hub} = 24 m and h_{ref} = 10 m. Turbine power (Eq. 4) is

$$P_{\mathrm{WT}}(t) = N_{\mathrm{WT}}P_{\mathrm{WT}}^{r}\begin{cases}0, & v < v_{\mathrm{ci}} \text{ or } v \ge v_{\mathrm{co}}\\ \frac{v^3-v_{\mathrm{ci}}^3}{v_r^3-v_{\mathrm{ci}}^3}, & v_{\mathrm{ci}} \le v < v_r\\ 1, & v_r \le v < v_{\mathrm{co}}\end{cases} \tag{4}$$

with cut-in, rated and cut-out speeds of 3, 12 and 25 m/s and N_{WT} = 3 turbines of 10 kW.

Battery. The state of charge evolves according to Eq. 5:

$$\mathrm{SoC}(t+1) = (1-\sigma)\mathrm{SoC}(t) + \frac{\eta_c P_{\mathrm{ch}}(t) - P_{\mathrm{dis}}(t)/\eta_d}{E_{\mathrm{bat}}}\Delta t \tag{5}$$

where σ is the hourly self-discharge, E_{bat} the nominal capacity, and η_c and η_d the AC-to-stored and stored-to-AC efficiencies (each the product of a 0.96 cell efficiency and the 0.95 inverter efficiency, giving an AC-to-AC round-trip efficiency of 83%). Operating limits (Eq. 6) are

$$\mathrm{SoC}_{\min} \le \mathrm{SoC}(t) \le \mathrm{SoC}_{\max} \tag{6a}$$

$$0 \le P_{\mathrm{ch}}, P_{\mathrm{dis}} \le P_{\mathrm{bat}}^{\max} \tag{6b}$$

with $\mathrm{SoC}_{\min}$ = 0.20, $\mathrm{SoC}_{\max}$ = 0.90 and $P_{\mathrm{bat}}^{\max}$ = 100 kW.

Diesel generator. Fuel use is linear in output, and operation is restricted to the minimum-load band (Eq. 7):

$$F_{\mathrm{DG}}(t) = a\,P_{\mathrm{DG}}(t) + b\,P_{\mathrm{DG}}^{r} \tag{7a}$$

$$P_{\mathrm{DG}}(t) \in \{0\} \cup [0.3P_{\mathrm{DG}}^{r},\, P_{\mathrm{DG}}^{r}] \tag{7b}$$

with a = 0.246 L/kWh, b = 0.084 L/kWh and P_{DG}^{r} = 50 kW. The hourly power balance (Eq. 8), including any unserved load P_{un} and dumped energy P_{dump}, is

$$P_{\mathrm{PV}}(t) + P_{\mathrm{WT}}(t) + P_{\mathrm{DG}}(t) + P_{\mathrm{dis}}(t) + P_{\mathrm{un}}(t) = P_{\mathrm{load}}(t) + P_{\mathrm{ch}}(t) + P_{\mathrm{dump}}(t) \tag{8}$$

Table 2 lists the baseline plant configuration and model parameters.

Table 2 Baseline plant configuration and model parameters

Component	Parameter	Value
PV array	Rated power / derating / temp. coefficient	100 kWp / 0.86 / −0.4 %/°C
Wind	Turbines / rated power each	3 / 10 kW (cut-in 3, rated 12, cut-out 25 m/s)
Battery	Nominal capacity / usable SoC window	400 kWh / 20–90 %
	Cell efficiency / inverter efficiency / max power	0.96 / 0.95 / 100 kW
	Self-discharge	0.003 % per hour
Diesel generator	Rated power / minimum load	50 kW / 30 %
	Fuel curve (a, b)	0.246 L/kWh, 0.084 L/kWh
	Forced outages (MTBF / MTTR)	400 h / 12 h
Load	Mean / peak (typical year)	20 kW / ≈ 34 kW

3.4 Energy management strategies

All strategies see the same plant layer. Two protections apply regardless of strategy: (i) the generator is never started for a residual deficit below its minimum load when the battery can still cover it above the hard SoC floor, and (ii) during a generator outage the battery may discharge to the hard floor. Surplus renewable energy that the battery cannot absorb is dumped.

Design of the comparison. The rule-based controller uses expert thresholds, the FLC is an expert-designed rule base, and the PSO-tuned policy optimizes the six threshold parameters that define the rule-based controller. Because tuning and controller structure are confounded in these three strategies, a fourth strategy, the FLC tuned with the same PSO settings, training years and fitness (Section 3.4.4), is added so that the effect of tuning can be separated from the effect of structure (Section 4.10).

3.4.1 Rule-based strategy and the parametric policy

The threshold policy has six parameters $\theta = [\mathrm{SoC}_{\mathrm{lo}}, \Delta_{\mathrm{on}}, g_{\mathrm{off}}, \kappa_{\mathrm{ch}}, \kappa_{\mathrm{DG}}, \mathrm{SoC}_{\mathrm{hi}}]$. The battery discharges only while SoC exceeds $\mathrm{SoC}_{\mathrm{lo}}$ and charges from surplus at no more than $\kappa_{\mathrm{ch}}P_{\mathrm{bat}}^{\max}$ while SoC is below $\mathrm{SoC}_{\mathrm{hi}}$. The generator turns on when a deficit exists and SoC falls to the start level and turns off at the stop level (Eq. 9):

$$\mathrm{SoC}_{\mathrm{on}} = \mathrm{SoC}_{\mathrm{lo}} + \Delta_{\mathrm{on}} \tag{9a}$$

$$\mathrm{SoC}_{\mathrm{off}} = \min(\mathrm{SoC}_{\max}, \mathrm{SoC}_{\mathrm{on}} + g_{\mathrm{off}}) \tag{9b}$$

While running, the generator supplies the deficit plus an optional battery-charging allowance $\kappa_{\mathrm{DG}}P_{\mathrm{bat}}^{\max}$, bounded by its rating and minimum load. The rule-based EMS is the expert setting θ_{RB} = [0.25,0,0,1,0,0.90], i.e. load-following diesel operation without cycle charging. It is deliberately a conventional, untuned baseline.

3.4.2 Mamdani fuzzy logic strategy

The FLC has three inputs normalized to [0, 1]: battery SoC (scaled over the usable window), load and renewable power (both divided by a 60 kW reference). Each input has three overlapping sets (Low, Medium, High) with trapezoidal or triangular

membership functions $\mu(\cdot)$ (Fig. 4a). The output is a battery command $u \in [-1,1]$ with five sets (Full discharge, Discharge, Hold, Charge, Full charge; Fig. 4b). The 27 rules are generated systematically (Table 3): the imbalance class $d = s_R - s_L$ (renewable level minus load level, each coded 0–2) selects the base action, which is shifted toward charging when SoC is low and toward discharging when SoC is high. Rule firing strength, aggregation and centroid defuzzification are given by Eqs. 10 and 11:

$$\alpha_r = \min(\mu^r_{\mathrm{SoC}}, \mu^r_L, \mu^r_R) \tag{10a}$$

$$\mu_{\mathrm{agg}}(u) = \max_r \min\left(\alpha_r, \mu_{C_r}(u)\right) \tag{10b}$$

$$u^* = \frac{\int u\, \mu_{\mathrm{agg}}(u)\, du}{\int \mu_{\mathrm{agg}}(u)\, du} \tag{11}$$

The crisp command limits battery power rather than prescribing it, so that power balance is always preserved by the plant layer (Eq. 12):

$$P_{\mathrm{ch}} \le \max(0, u^*) P^{\max}_{\mathrm{bat}} \tag{12a}$$

$$P_{\mathrm{dis}} \le [0.3 + 0.7\max(0, -u^*)] P^{\max}_{\mathrm{bat}} \tag{12b}$$

The fuzzy controller of this section is not tuned; its design is a plausible expert baseline, and a PSO-tuned version is described in Section 3.4.4. Fig. 4 and Fig. 5 show the membership functions and the resulting control surface.

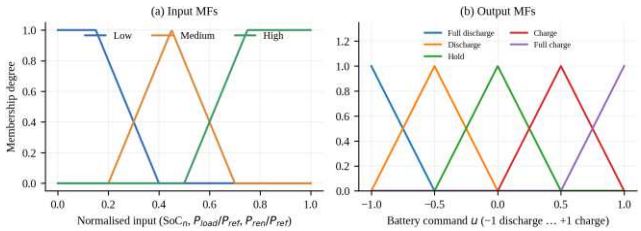


Fig. 4 Membership functions of the Mamdani FLC: (a) inputs, (b) output

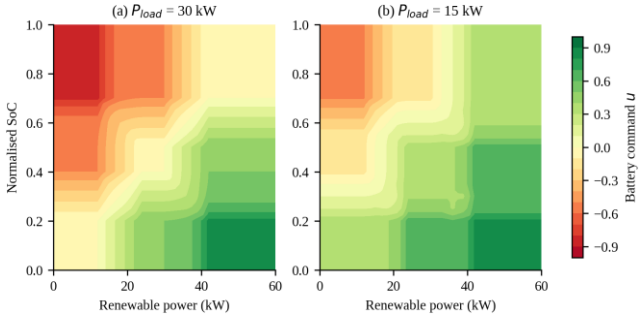


Fig. 5 Control surface of the FLC: battery command versus renewable power and normalized SoC at (a) 30 kW and (b) 15 kW load

Table 3 Complete 27-rule base of the Mamdani FLC (consequent battery command)

SoC	Load	Renewable: Low	Renewable: Medium	Renewable: High
Low	Low	Charge	Full charge	Full charge
	Medium	Hold	Charge	Full charge
	High	Discharge	Hold	Charge
Medium	Low	Hold	Charge	Full charge
	Medium	Discharge	Hold	Charge
	High	Full discharge	Discharge	Hold
High	Low	Discharge	Hold	Charge
	Medium	Full discharge	Discharge	Hold
	High	Full discharge	Full discharge	Discharge

3.4.3 PSO-tuned policy

PSO searches the six-dimensional parameter space of the threshold policy offline. Each particle i updates according to Eqs. 13 and 14:

$$v_i^{k+1} = w^k v_i^k + c_1 r_1 (p_i - x_i^k) + c_2 r_2 (g - x_i^k) \tag{13a}$$

$$x_i^{k+1} = x_i^k + v_i^{k+1} \tag{13b}$$

$$w^k = w_0 - (w_0 - w_1)\frac{k}{K-1} \tag{14}$$

where p_i and g are the personal and global best positions, $r_1, r_2 \sim U(0,1)$, and inertia decreases linearly from $w_0 = 0.9$ to $w_1 = 0.4$. The fitness is the mean NPC over N_y training years with a penalty for exceeding the limit on the loss of power supply probability (LPSP) (Eq. 15):

$$J(\theta) = \frac{1}{N_y}\sum_{y=1}^{N_y}\left[\mathrm{NPC}_y(\theta) + \lambda \max\left(0, \mathrm{LPSP}_y(\theta) - \mathrm{LPSP}_{\max}\right)\right] \tag{15}$$

with $\mathrm{LPSP}_{\max} = 2\%$ and $\lambda = 5\times10^7$ \$. Bounds are handled by clipping and velocity is clamped to ±20% of each range. One particle is initialized at the expert thresholds, which guarantees that the tuned policy is never worse than the rule-based baseline in-sample; this is a methodological choice that is disclosed here. Algorithm 1 summarizes the procedure and Table 4 lists the settings.

Algorithm 1 PSO tuning of the dispatch policy

1: **Input:** training years $\mathcal{Y}$, bounds $[\theta_{\mathrm{lb}}, \theta_{\mathrm{ub}}]$, swarm size n, iterations K
2: Initialize positions x_i uniformly in bounds; set $x_1 \leftarrow \theta_{\mathrm{RB}}$; initialize velocities
3: Evaluate $J(x_i)$ by simulating all years in $\mathcal{Y}$; set $p_i \leftarrow x_i$, $g \leftarrow \mathrm{argmin} J$
4: **for** $k = 0, \ldots, K-1$ **do**
5: update w^k; update velocities and positions; clamp and clip
6: evaluate $J(x_i)$; update p_i and g
7: **end for**
8: **return** $\theta^* = g$ (best of five independent runs)

Table 4 PSO settings and tuned policy parameters (mean ± standard deviation over five independent runs)

Item	Value
Swarm size / iterations / training years	24 / 50 / 2 (fitness evaluations per run: 1224)
Inertia / acceleration coefficients	0.9→0.4 / $c_1 = c_2 = 2.0$
Velocity clamp / reliability limit	±20 % of range / LPSP ≤ 2 %
Parameter bounds (lower–upper)	SoC_lo 0.2–0.5; ΔSoC_on 0–0.3; g_off 0–0.5; κ_ch 0.3–1; κ_DG 0–1; SoC_hi 0.75–0.9
Rule-based θ_RB	0.25; 0; 0; 1; 0; 0.9
Tuned θ* (best run)	0.20; 0.00; 0.08; 1.00; 0.41; 0.90
Tuned θ (mean ± standard deviation of 5 runs)	0.20±0.00; 0.00±0.00; 0.02±0.03; 0.76±0.13; 0.40±0.00; 0.90±0.00

3.4.4 Tuned fuzzy controller

The FLC has six design parameters: a shift of the SoC membership functions, the reference power used to normalize the load and renewable inputs, a gain on the charge command, the base discharge fraction, the discharge floor above the hard SoC minimum, and a generator charging allowance (a fraction of the battery power rating, analogous to the diesel charging parameter of the threshold policy). They are tuned by the same PSO as in Section 3.4.3 (24 particles, 50 iterations, five independent runs, the same two training years and the same fitness); one particle is seeded at the untuned design. Fitness evaluations inside PSO use a fast lookup of the inference surface (LCOE error below 0.1%); all reported results use exact inference. A variant without the generator allowance is also tuned.

3.5 Battery aging model

The SoC trace is reduced to cycles with the ASTM rainflow counting algorithm. Each cycle of depth DoD_i (the depth of discharge, DoD, expressed as a fraction of nominal capacity) consumes life according to a power-law curve, temperature accelerates aging by an Arrhenius factor, and calendar aging adds a constant term, as given by Eqs. 16 and 17 (Dufek et al. 2022):

$$N_f(\mathrm{DoD}) = N_{80}\left(\frac{\mathrm{DoD}}{0.8}\right)^{-k} \tag{16a}$$

$$\mathrm{AF} = \exp\left[\frac{E_a}{R}\left(\frac{1}{T_{\mathrm{ref}}} - \frac{1}{T_b}\right)\right] \tag{16b}$$

$$L_{\mathrm{bat}} = \min\left\{L_{\max},\ \left[\mathrm{AF}\left(\sum_i \frac{n_i}{N_f(\mathrm{DoD}_i)} + \frac{1}{L_{\mathrm{cal}}}\right)\right]^{-1}\right\} \tag{17}$$

Here n_i is the cycle count (0.5 or 1), T_b the battery temperature taken as mean ambient plus 4 K, and the parameters are in Table 5. The aging parameters are generic lithium-ion values chosen for illustration; they should be replaced by cell-specific data in applications.

3.6 Economic and performance metrics

Costs are evaluated over a N = 20-year life at a real discount rate i = 10%. The net present cost combines capital, fixed and variable operation and maintenance (O&M) costs, diesel fuel, component replacements and salvage; the levelized cost of energy (LCOE) divides the annualized cost (the NPC multiplied by the capital recovery factor, CRF) by the energy actually served (Eqs. 18 and 19):

$$\mathrm{NPC} = C_{\mathrm{cap}} + (C_{\mathrm{om}} + C_{\mathrm{fuel}} + C_{\mathrm{DG}})\frac{1-(1+i)^{-N}}{i} + C_{\mathrm{rep}} - C_{\mathrm{sal}} \tag{18}$$

$$\mathrm{CRF} = \frac{i}{1-(1+i)^{-N}} \tag{19a}$$

$$\mathrm{LCOE} = \frac{\mathrm{NPC}\cdot\mathrm{CRF}}{\sum_{t=1}^{T}[P_{\mathrm{load}}(t) - P_{\mathrm{un}}(t)]\Delta t} \tag{19b}$$

Battery replacements occur every L_{bat} years and diesel overhauls every 15 000 running hours, each discounted to present value, with linear salvage of remaining life at year N. Reliability, renewable use, cycling and cost saving are quantified by the LPSP, the renewable fraction (RF), the equivalent full cycles (EFC) and the LCOE saving, as defined in Eqs. 20 and 21,

$$\mathrm{LPSP} = \frac{\sum_t P_{\mathrm{un}}(t)\Delta t}{\sum_t P_{\mathrm{load}}(t)\Delta t} \tag{20a}$$

$$\mathrm{RF} = 1 - \frac{\sum_t P_{\mathrm{DG}}(t)\Delta t}{\sum_t P_{\mathrm{load}}(t)\Delta t} \tag{20b}$$

$$\mathrm{EFC} = \frac{\sum_t P_{\mathrm{dis}}(t)\Delta t/\eta_d}{E_{\mathrm{bat}}} \tag{21a}$$

$$S_x = \frac{\mathrm{LCOE_{RB}} - \mathrm{LCOE}_x}{\mathrm{LCOE_{RB}}} \times 100\% \tag{21b}$$

where EFC is the equivalent number of full cycles per year and S_x the LCOE saving of strategy x relative to the rule-based baseline.

Table 5 Economic and battery-aging parameters (assumed values; to be replaced by project quotes)

Item	Value	Item	Value
Project life / real discount rate	20 yr / 10 %	PV / wind capex	800 / 2800 $/kW
Battery capex / replacement cost	320 / 250 $/kWh	Inverter / diesel capex	250 $/kW / 350 $/kW
Balance-of-system uplift	10 %	Fixed O&M (PV, wind, battery)	12 $/kW-yr, 80 $/kW-yr, 4 $/kWh-yr
Diesel price	1.25 $/L	Diesel O&M / start cost	1.0 $/h / 2.0 $ per start
N_{80} / exponent k	3500 / 1.6	Activation energy E_a	30 kJ/mol
Calendar life L_cal / cap L_max	15 / 20 yr	Battery temp. offset	+4 K above ambient

3.7 Experimental protocol and statistics

PSO is run five times with different swarm seeds on the same two training years (weather seeds 11 and 12), and the run with the lowest training fitness defines the PSO-tuned policy. Two years (17,520 simulated hours) were chosen as a compromise between tuning cost and representativeness: the policy has only six parameters, and its generalization is tested on years never used in tuning (Section 4.11). All strategies are then evaluated on 100 independent test years (seeds 1000–1099) that were never seen during tuning, using common random numbers so that comparisons are paired. A diesel-only plant of identical load and generator serves as a reference. Differences are tested with the Friedman test across the three strategies and pairwise Wilcoxon signed-rank tests with Holm correction; mean paired differences carry 10 000-sample bootstrap 95% confidence intervals (CIs). Because the data are simulated, these statistics describe consistency across synthetic weather realizations and do not measure uncertainty about a real site. To test sensitivity to the training set, the threshold policy was also tuned with 1, 2, 5 and 10 training years (five replicates per size, with disjoint training years and different swarm seeds) and evaluated on the same 100 test years. Further experiments use 30 test years for stress scenarios and economic sensitivities (policies held fixed) and 20 years for the 12-point sizing sweep (PSO re-tuned for each sizing with 16 particles and 30 iterations, a smaller budget than in the baseline run).

3.8 Software implementation and reproducibility

The framework consists of four Python modules (plant, controllers and economics; experiment driver; figure generation; manuscript builder). The fitness function, the hourly dispatch (Algorithm 2) and all random seeds are deterministic, so every number in this paper can be regenerated: the seed schedule is 11–12 for training, 1000–1099 for the baseline test, 2000–2029 and 3000–3029 for the stress and sensitivity experiments, and 4000–4019 for the sizing sweep. A single 8760-hour simulation takes about 20 ms for threshold policies and 0.2 s for the fuzzy controller on a single CPU core (timings are hardware dependent). All parameter values are listed in Tables 1–5, and the source code, the raw per-year outputs and a README with the run order are provided in the public repository cited in the Data and code availability statement.

Algorithm 2 Hourly dispatch with shared plant protections

1: **for** $t = 1, \ldots, 8760$ **do**
2: $\Delta P \leftarrow P_{\mathrm{PV}} + P_{\mathrm{WT}} - P_{\mathrm{load}}$; apply self-discharge; update generator state by $\mathrm{SoC_{off}}$ and availability
3: **if** $\Delta P \geq 0$ **then** charge battery with $\min(\Delta P,$ rate and SoC limits); dump the rest
4: **else** (deficit)
5: **if** policy requests generator and generator available **then** run generator at $\max(\text{min load},\ \text{deficit} + \text{allowance})$; surplus charges battery
6: **else** discharge battery above its floor; **if** residual $\geq$ min load and generator available **then** start generator;
7: **else if** residual < min load and battery above hard floor **then** battery covers residual
8: any remaining deficit is unserved load
9: update SoC, fuel, starts, running hours, dumped and unserved energy
10: **end for**; post-process SoC trace by rainflow counting to obtain L_{bat} and then NPC, LCOE

3.9 Use of generative AI

The authors used an AI assistant (ChatGPT, OpenAI) to draft parts of the manuscript text and to correct grammar. They reviewed and edited all content and take full responsibility for it; the assistant is not listed as an author.

4 Results and discussion

4.1 Verification of the implementation

Table 6 shows that the implementation reproduces the analytical values of the component models and closes the hourly power balance to numerical precision, and that the realized statistics of the generated data (annual irradiation, generator availability, mean and peak load) lie within about 4% of their targets. These tests verify the code and the generator against their own specifications, not the realism of the assumed site, whose weather statistics and cost parameters are assumptions: they do not validate the data against measurements, and the last two rows of Table 6 are plausibility checks without an external benchmark. Validation against measured hourly data was not possible in this study and is the main limitation (Section 4.12).

Table 6 Code verification checks

Check	Expected	Obtained
Power balance residual (3 strategies, 8760 h)	0	1.4e-14 kW (max abs.)
PV output at standard test conditions (G = 1000 W/m², T_c = 25 °C)	86 kW	86.0 kW
Wind output at rated speed / below cut-in	30 kW / 0 kW	30.0 kW / 0.0 kW
Diesel fuel at rated and at minimum load	16.5 / 7.9 L/h	16.5 / 7.9 L/h
Battery AC-to-AC round-trip efficiency	$0.96^2 \cdot 0.95^2 = 0.83$	0.83
Rainflow: 10 cycles of depth 0.4	10 cycles	10.0 cycles
Annual GHI of generated years (20-yr mean)	≈ 1717 kWh/m²	1728 kWh/m²
PV specific yield / wind capacity factor	plausible for 6.5° N	1354 kWh/kWp / 9.6 %
Generator availability (MTBF 400 h, MTTR 12 h)	97.1 %	97.4 %
Load mean / peak	20 kW / ≈ 34 kW	19.3 / 33.6 kW

4.2 Baseline performance on 100 out-of-sample years

PSO converged consistently: all five runs reached nearly identical fitness within roughly 10–15 iterations (Fig. 6a). All five placed the SoC start level at its lower bound with no start-level offset, and all five set the diesel charging allowance near 0.40 (Fig. 6b and Table 4), i.e. the swarm favors running the generator less often but at higher loading with cycle charging. The remaining parameters (charging-rate fraction) were weakly identified, with similar fitness across a range of values.

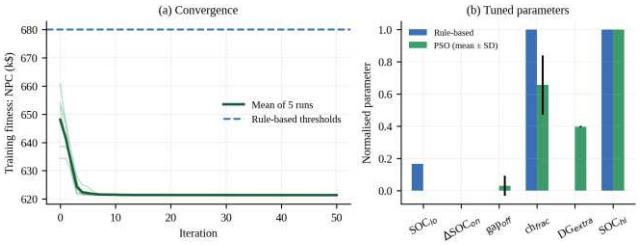


Fig. 6 PSO tuning: (a) convergence of the training fitness for five independent runs; (b) tuned policy parameters relative to the rule-based thresholds

Table 7 gives the baseline results. The PSO-tuned policy has the lowest NPC and LCOE (0.429 $/kWh versus 0.467 for the rule-based and 0.476 for the fuzzy controller), cutting diesel use by 28% and generator running hours by 64%. It does so by running the generator at higher loading with cycle charging, which is more fuel-efficient per kWh, at the cost of 23% more battery cycling and a battery life of 6.8 years versus 7.1. All three hybrid strategies are far cheaper than the diesel-only reference (0.711 $/kWh with 2.8% unserved load under the same outage model).

Table 7 Baseline results over 100 test years (mean ± 95% CI). Diesel-only reference: 30 years

Metric	Rule-based	Fuzzy (Mamdani)	PSO-tuned
NPC (k$)	669.3 ± 2.3	680.7 ± 2.2	614.4 ± 1.6
LCOE ($/kWh)	0.4674 ± 0.0016	0.4758 ± 0.0015	0.4290 ± 0.0011
LPSP (%)	0.392 ± 0.038	0.483 ± 0.042	0.366 ± 0.036
Renewable fraction (%)	83.5 ± 0.2	81.7 ± 0.2	81.7 ± 0.2
Renewable energy utilized (%)	94.9 ± 0.2	92.4 ± 0.2	94.8 ± 0.2
Diesel fuel (L/yr)	14157 ± 175	15635 ± 169	10193 ± 127
Generator running hours (h/yr)	1739 ± 21	1911 ± 20	619 ± 8
Generator starts (per yr)	492 ± 6	386 ± 4	597 ± 7
Battery life (yr)	7.10 ± 0.01	7.48 ± 0.01	6.81 ± 0.01
Equivalent full cycles (per yr)	162 ± 0	151 ± 0	200 ± 1
Dumped energy (MWh/yr)	8.1 ± 0.4	12.2 ± 0.4	8.3 ± 0.4

The untuned FLC extends battery life (7.5 years, 7% fewer equivalent cycles), but because it is conservative about discharging at low SoC the generator runs more (1911 h versus 1739 h), raising fuel cost by more than the replacement savings. Reliability is high for all strategies; the small LPSP arises from generator outages that coincide with low SoC. The fuzzy gain in battery life is therefore real but outweighed by diesel cost, and tuning the FLC changes that conclusion (Section 4.10).

4.3 Operational behavior

Fig. 7 compares seven days of dispatch in the rainy season. During the sunny first days all strategies behave alike. When resource drops (days 5–7) the rule-based controller runs the generator at partial load to follow the deficit, the fuzzy controller keeps SoC close to the floor with intermittent generator support, and the PSO-tuned policy runs the generator in short high-power bursts that also recharge the battery. Fig. 8a shows that the PSO policy spends far less time near the SoC floor (13% of hours below 25% SoC versus 21% for the rule-based and 34% for the fuzzy controller in the example year), and Fig. 8b shows that diesel use peaks in July–September but is lowest for the PSO policy in every month.

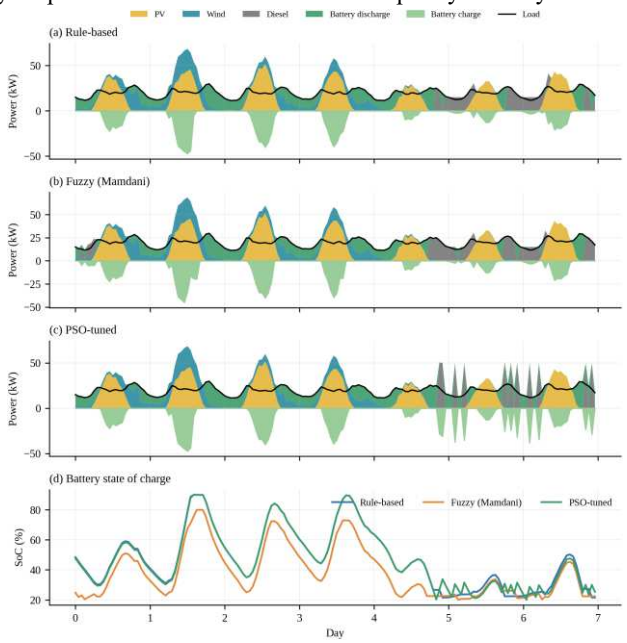


Fig. 7 Seven-day dispatch in the rainy season for (a) rule-based, (b) fuzzy and (c) PSO-tuned strategies, and (d) battery SoC

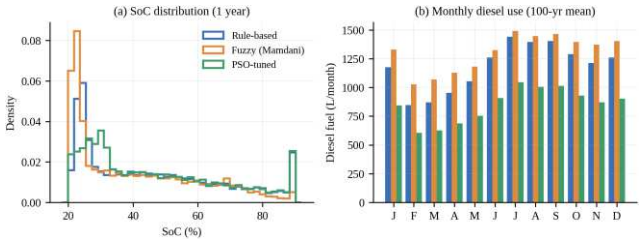


Fig. 8 (a) Annual SoC distribution and (b) mean monthly diesel consumption over 100 test years

4.4 Economics and statistical significance

Fig. 9 shows the distributions across the 100 test years and Fig. 10 decomposes NPC. Capital cost is common to all strategies; the savings of the PSO policy arise from diesel fuel and generator O&M, partly offset by slightly earlier battery replacement.

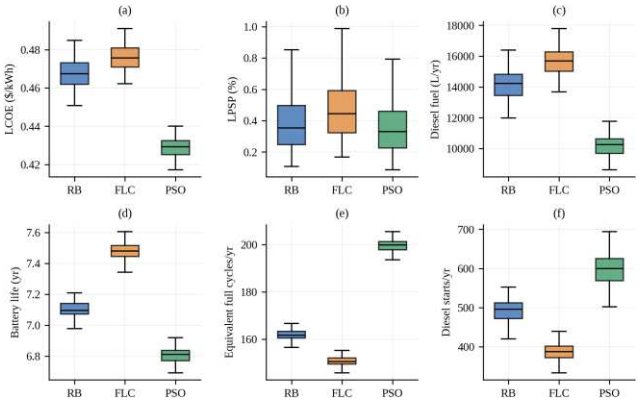


Fig. 9 Distribution over 100 test years of (a) LCOE, (b) LPSP, (c) diesel fuel, (d) battery life, (e) equivalent full cycles and (f) generator starts

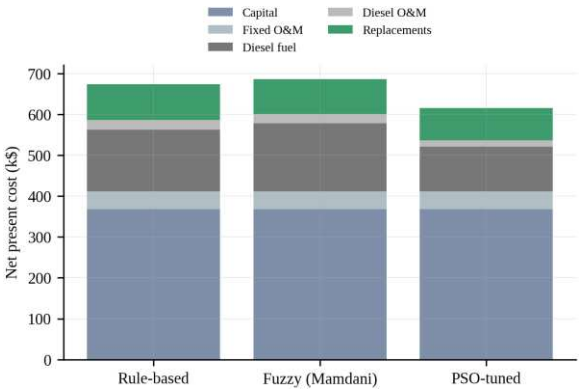


Fig. 10 Mean NPC decomposition by strategy

The Friedman test rejects equal performance ($\chi^2(2) = 200$, p = 3.7e-44). Every pairwise difference is significant after Holm correction (Table 8), and the ordering PSO < rule-based < fuzzy holds in all 100 test years. The mean LCOE saving of PSO over the rule-based controller is 8.2% (bootstrap 95% CI 8.1–8.3%), equivalent to an NPC reduction of about $55k, and the fuzzy controller is 1.8% costlier. With simulated data these p-values and intervals describe only consistency across synthetic years and can be made arbitrarily small by adding years, so the effect size is the informative quantity; they do not capture structural uncertainty in the data and cost assumptions, addressed next. The identical raw p-values in Table 8 arise because all 100 paired differences have the same sign.

Table 8 Paired LCOE differences (m$/kWh, positive = first strategy costlier), bootstrap 95% CI, years in which the first is costlier, and Wilcoxon p-values

Comparison	Mean difference [95% CI]	Years first costlier	Raw p	Holm p
Rule-based − PSO-tuned	+38.5 [+38.0, +39.0]	100/100	3.9e-18	1.2e-17
Rule-based − Fuzzy	-8.4 [-8.6, -8.1]	0/100	3.9e-18	1.2e-17
Fuzzy − PSO-tuned	+46.8 [+46.4, +47.3]	100/100	3.9e-18	1.2e-17

4.5 Economic sensitivity

Fig. 11 recomputes LCOE with the policies held fixed. The ordering is unchanged across diesel prices of 0.6–2.5 $/L, real discount rates of 6–20% and battery costs of 150–520 $/kWh. The PSO advantage grows with diesel price (from 5.6% at 0.6 $/L to 11.9% at 2.5 $/L), because its benefit is fuel-driven, whereas it shrinks as the discount rate rises (from 9.7% at 6% to 5.7% at 20%, since capital cost, which is identical across strategies, weighs more) and as battery cost rises (from 10.4% at 150 $/kWh to 6.5% at 520 $/kWh). The fuzzy controller's cost penalty relative to the rule-based baseline grows with diesel price (0.6% to 3.4%) and shrinks as battery cost rises (2.7% to 1.1%), consistent with its lower cycling.

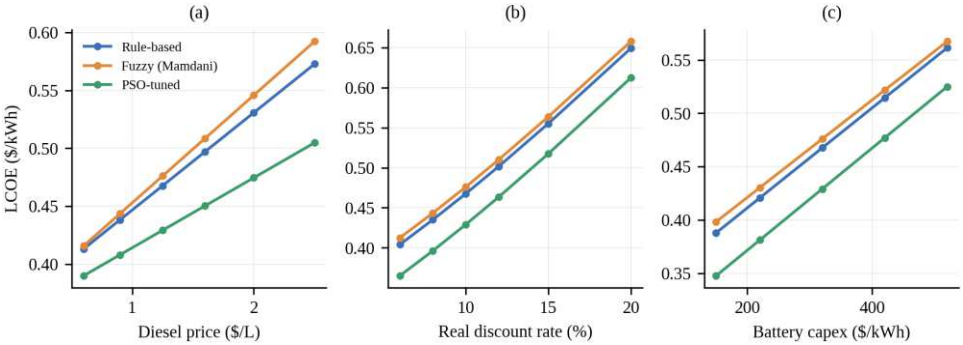


Fig. 11 LCOE sensitivity to (a) diesel price, (b) real discount rate and (c) battery capex, with policies fixed

4.6 Stress scenarios

Table 9 and Fig. 12 show that, among the three strategies of the baseline comparison, the PSO-tuned policy keeps the lowest LCOE and the lowest LPSP under +20% load, a 15% irradiance deficit and tripled generator outage rates, even though outages are not part of its state. This may reflect outages in its training years, but the mechanism was not isolated. The untuned fuzzy controller shows the highest LPSP in every scenario. The PSO-tuned FLC of Section 4.10 has LCOE within 0.7% of the PSO-tuned policy and lower LPSP in every scenario (0.31–0.79% versus 0.40–1.05%). Under the +20% load case, the LPSP of all strategies rises by roughly 65–85%, indicating that no EMS substitutes for adequate sizing. LCOE is slightly lower under tripled outages because unserved energy is not monetized and the generator burns less fuel when unavailable, so cost must be read together with LPSP.

Table 9 LCOE ($/kWh) and LPSP (%) under stress scenarios (30 years each, policies fixed)

Scenario	LCOE RB	LCOE FLC	LCOE PSO	LPSP RB	LPSP FLC	LPSP PSO
Baseline	0.468	0.477	0.429	0.43	0.54	0.40
Load +20%	0.479	0.486	0.430	0.78	0.89	0.72
Solar drought (-15% GHI)	0.519	0.520	0.465	0.64	0.74	0.61
DG outages x3 (MTBF 133 h)	0.465	0.474	0.428	1.14	1.42	1.05

Abbreviations: RB, rule-based controller; FLC, untuned fuzzy logic controller; PSO, PSO-tuned policy; LCOE, levelized cost of energy; LPSP, loss of power supply probability; DG, diesel generator; MTBF, mean time between failures; GHI, global horizontal irradiance.

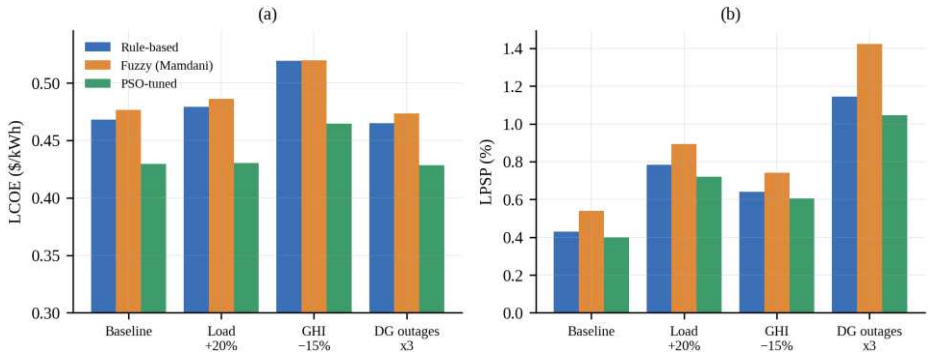


Fig. 12 Stress-test results: (a) LCOE and (b) LPSP for baseline, +20% load, −15% irradiance and tripled generator outage rate

4.7 Where does EMS sophistication pay? Sizing sweep

Fig. 13 and Table 10 show how strongly the benefit depends on sizing. When the plant is generous (150 kWp PV with 400–600 kWh storage, renewable fraction above 96%), the generator is rarely needed and the EMS matters little: the PSO saving falls to 0.8–1.8%. When PV is scarce, the generator dominates operation and PSO savings reach 16.1%. The FLC is within ±2% of the rule-based controller in nearly all cells, with small gains only for the smallest PV sizes combined with large batteries. Within the assumed climate, this map suggests that tuning effort is best spent on tightly sized plants.

Table 10 Sizing sweep (20 years per cell; PSO re-tuned per sizing). Savings are relative to the rule-based LCOE

PV (kWp) / battery (kWh)	RF, RB (%)	LCOE RB	LCOE FLC	LCOE PSO	FLC saving (%)	PSO saving (%)
60 / 200	59.0	0.554	0.556	0.465	-0.3	+16.1
60 / 400	59.9	0.598	0.588	0.512	+1.6	+14.3
60 / 600	60.1	0.658	0.646	0.569	+1.8	+13.5
80 / 200	68.8	0.505	0.509	0.436	-0.9	+13.6
80 / 400	72.6	0.528	0.525	0.467	+0.5	+11.7
80 / 600	73.2	0.582	0.575	0.521	+1.2	+10.6
100 / 200	74.0	0.485	0.487	0.428	-0.3	+11.8
100 / 400	83.6	0.467	0.476	0.429	-1.8	+8.1
100 / 600	85.1	0.514	0.515	0.479	-0.2	+6.8
150 / 200	78.5	0.487	0.485	0.442	+0.6	+9.3
150 / 400	96.4	0.413	0.427	0.405	-3.6	+1.8
150 / 600	98.4	0.453	0.459	0.449	-1.2	+0.8

Abbreviations: RF, renewable fraction; RB, rule-based controller; FLC, untuned fuzzy logic controller; PSO, PSO-tuned policy; LCOE, levelized cost of energy ($/kWh).

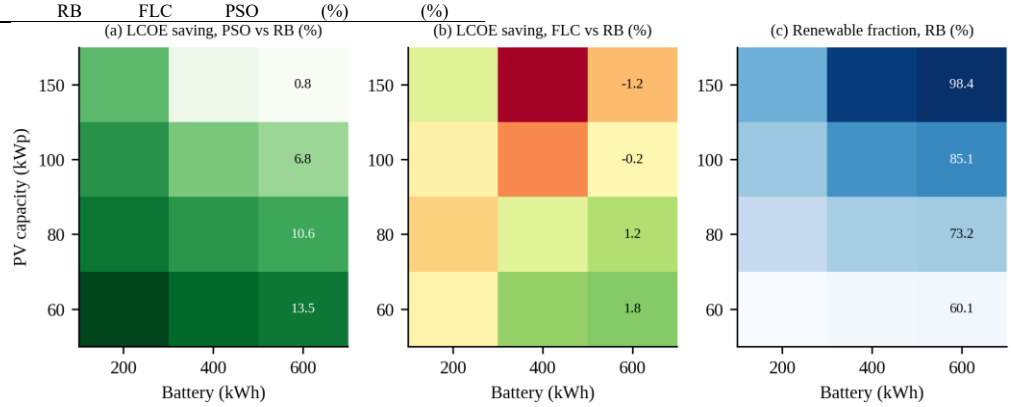


Fig. 13 Sizing sweep: (a) LCOE saving of PSO versus rule-based, (b) of FLC versus rule-based, (c) renewable fraction of the rule-based strategy

4.8 Computational cost

Table 11 shows that fuzzy inference is roughly 114 times slower per decision than a threshold policy, yet at 17 µs it is irrelevant for an hourly dispatch loop and is feasible on embedded controllers; the figures are indicative of relative cost only and were not measured on embedded hardware. PSO moves its cost offline: its tuned policy is as cheap online as the rule-based controller, while tuning took about 34 s per run here.

Table 11 Computational cost on a single x86 CPU core (Python implementation)

Strategy	Online decision time (µs/step)	Offline effort	Note
Rule-based	0.15	none	threshold comparisons
Fuzzy (Mamdani)	17.0	rule design	27 rules, 201-point centroid
PSO-tuned policy	0.15	34 s per run (1224 simulations)	same online form as rule-based

4.9 Multi-criteria view

The strategies of the baseline comparison trade off cost, reliability, battery life, diesel use and online cost: the PSO-tuned policy leads on cost, reliability (tied) and diesel savings and equals the rule-based controller online; the untuned fuzzy controller leads on battery life; the rule-based controller is simplest. The tuned FLC lies close to the PSO-tuned policy on cost and ahead of it on reliability (Section 4.10).

4.10 Equal-effort comparison with a tuned fuzzy controller

To separate controller structure from tuning effort, the FLC was tuned with the same PSO budget as the threshold policy (Section 3.4.4). Table 12 shows that the tuned FLC lowered mean LCOE by 7.7% relative to the rule-based baseline (paired difference -0.0359 $/kWh, bootstrap 95% CI -0.0365 to -0.0353; lower in all 100 years) and by 9.3% relative to the untuned FLC, and came within 0.6% of the PSO-tuned threshold policy, which remained cheaper in all 100 years. The tuned FLC had lower unserved energy (LPSP 0.27% versus 0.37%) and fewer equivalent full cycles (187 versus 200 per year), with a slightly longer battery life (7.0 versus 6.8 years).

Table 12 Equal-effort comparison on the 100 test years (means). The two PSO-tuned strategies use the same PSO budget, training years and fitness; EFC, equivalent full cycles

Strategy	LCOE ($/kWh)	NPC (k$)	LPSP (%)	Diesel (L/yr)	Battery life (yr)	EFC (1/yr)
Rule-based (expert-set)	0.4674	669.3	0.392	14157	7.10	162
FLC, untuned	0.4758	680.7	0.483	15635	7.48	151
FLC, PSO-tuned	0.4315	618.7	0.266	10731	6.97	187
FLC, PSO-tuned, no generator allowance	0.4580	656.1	0.346	13005	7.35	151
PSO-tuned threshold policy	0.4290	614.4	0.366	10193	6.81	200

Without the generator charging allowance, tuning improved the FLC by only 2.0% relative to the baseline (LCOE 0.458 $/kWh), so most of the gain of both tuned strategies comes from letting the

generator charge the battery at higher loading rather than from the choice between fuzzy and threshold structure. Two tuned FLC parameters ended at their bounds (SoC shift and charge gain), so a wider search space might improve it further. The cost penalty of the untuned FLC in Section 4.2 therefore reflects missing tuning, not a weakness of fuzzy inference, and the remaining 0.6% gap to the PSO-tuned policy is small compared with the 9.3% gained by tuning.

4.11 Robustness to the training conditions

To test sensitivity to the training set, the threshold policy was re-tuned with 1, 2, 5 and 10 training years (five replicates per size, each with different, disjoint training years and a different swarm seed), and every tuned policy was evaluated on the same 100 test years (Table 13). The saving relative to the rule-based baseline was 8.17 ± 0.02% with one year and 8.21 ± 0.03% with ten years, and all 20 tuned policies had lower cost than the rule-based controller in every one of the 100 test years. The mean difference between test and training NPC was between -1.0 and 2.0 k$ for the four sizes, small compared with a test NPC of about 614 k$, so there is no noticeable overfitting.

Table 13 Sensitivity of the PSO-tuned threshold policy to the number of training years (five replicates per size, mean ± SD; each replicate uses different training years and a different swarm seed)

Training years	Saving vs rule-based (%)	Range (%)	Years better than rule-based (min of 100)	Test NPC minus training NPC (k$)
1	8.17 ± 0.02	8.16–8.19	100	-1.0 ± 10.0
2	8.18 ± 0.05	8.13–8.23	100	2.0 ± 6.6
5	8.19 ± 0.03	8.17–8.23	100	1.0 ± 2.3
10	8.21 ± 0.03	8.18–8.23	100	-0.6 ± 2.5

The tuned thresholds themselves were less stable than the cost: with one training year the standard deviation across replicates of the diesel charging allowance was 0.22 and that of the charging-rate fraction 0.24, falling to 0.07 and 0.17 with ten years, while the SoC limits varied by less than 0.01. Cost is therefore nearly flat along these parameters, so two training years suffice for this policy structure and generator, but individual parameters should not be read as precisely identified from short training sets. Because only the sampled years varied, this does not test sensitivity to errors in the assumed climate (Section 4.12).

4.12 Limitations

The study has several limitations that qualify its conclusions. (i) Synthetic data: weather and load are generated from assumed climatology, so absolute LCOE values (≈ 0.43–0.48 $/kWh) depend on assumed capex, diesel price, discount rate and generic aging parameters and are not site forecasts. Rankings were stable across the stress and sensitivity ranges tested, but the generator's temporal correlations, extremes and load–weather dependence were not validated against measurements, and transferability to other climates or loads is not demonstrated. (ii) Tuning scope: the tuned FLC used six design parameters chosen by the authors, two of which ended at their bounds, and one PSO configuration; other tuned structures (for example neuro-fuzzy controllers) and other optimizers were not examined, so the remaining small gap between the tuned FLC and the PSO-tuned policy should not be generalized. (iii) Training conditions: the training-set study varied only the sampled years of the same synthetic generator, so it shows insensitivity to year sampling, not to errors in the assumed climate; the sizing sweep also used a smaller PSO budget. (iv) Model scope: hourly resolution, linear fuel curve, fixed-efficiency inverter, no voltage or frequency dynamics, one topology and one cost objective, and no forecasting-based or reinforcement-learning EMS.

5 Conclusions

This study compared an expert-set rule-based controller, an untuned expert-designed Mamdani fuzzy controller and a PSO-tuned threshold policy on a standalone PV–wind–battery–diesel microgrid, using identical plant models, protection logic, aging-aware economics and out-of-sample Monte Carlo testing on synthetic data. Within these limits, the main findings are:

- Tuning the thresholds of the rule-based policy with PSO reduced LCOE by 8.2% and diesel consumption by 28% relative to the expert-set baseline, with lower cost in all 100 simulated test years, but increased battery cycling by 23% and shortened battery life modestly.
- The untuned fuzzy controller reduced cycling by 7% but was 1.8% costlier than the expert-set baseline. The same controller tuned with the same PSO budget reached 7.7% lower LCOE than the baseline, 0.6% above the PSO-tuned threshold policy, with less unserved energy. Most of the difference between strategies therefore reflects tuning effort, in particular allowing the generator to charge the battery, not the control paradigm.
- The benefit of tuning depends strongly on sizing (0.8% for generous plants to 16.1% for undersized ones), persisted under the load growth, solar deficit, tripled outage rate and economic sensitivities tested, and was insensitive to the number of training years (8.2% with 1 year and 8.2% with 10 years) within the synthetic generator.
- Tuning moves computation offline (≈ 34 s here), whereas fuzzy inference costs about 114 times more per step than a threshold policy but remains negligible for hourly dispatch.

These conclusions are conditional on the assumed synthetic climate, cost and aging parameters, which were not validated against measured data. Future work should validate the models with measured hourly data for specific sites, examine other tuned structures and optimizers (for example neuro-fuzzy controllers), add forecasting-based and reinforcement-learning EMS, co-optimize sizing and control, and validate on a hardware-in-the-loop testbed with cell-specific aging data.

References

Adenuga OT, Krishnamurthy S (2023) Economic power dispatch of a grid-tied photovoltaic-based energy management system: Co-optimization approach. Mathematics 11:3266. https://doi.org/10.3390/math11153266

Ahmad S, Shafiullah M, Ahmed CB, Alowaifeer M (2023) A review of microgrid energy management and control strategies. IEEE Access 11:21729–21757. https://doi.org/10.1109/ACCESS.2023.3248511

Al Sumarmad KA, Sulaiman N, Wahab NIA, Hizam H (2022) Energy management and voltage control in microgrids using artificial neural networks, PID, and fuzzy logic controllers. Energies 15:303. https://doi.org/10.3390/en15010303

Aziz A et al (2021) Hybrid wind/PV/battery energy management-based intelligent non-integer control for smart DC-microgrid of smart university. IEEE Access 9:98948–98961. https://doi.org/10.1109/ACCESS.2021.3095973

Battula AR, Vuddanti S, Salkuti SR (2021) Review of energy management system approaches in microgrids. Energies 14:5459. https://doi.org/10.3390/en14175459

Bihari SP et al (2021) A comprehensive review of microgrid control mechanism and impact assessment for hybrid renewable energy integration. IEEE Access 9:88942–88958. https://doi.org/10.1109/ACCESS.2021.3090266

Bouaouda A, Sayouti Y (2022) Hybrid meta-heuristic algorithms for optimal sizing of hybrid renewable energy system: A review of the state-of-the-art. Arch Comput Methods Eng 29. https://doi.org/10.1007/s11831-022-09730-x

Dufek EJ, Tanim TR, Chen BR, Kim S (2022) Battery calendar aging and machine learning. Joule 6:1363–1367. https://doi.org/10.1016/j.joule.2022.06.007

Ejiyi CJ et al (2025) Comprehensive review of artificial intelligence applications in renewable energy systems: Current implementations and emerging trends. J Big Data 12:169. https://doi.org/10.1186/s40537-025-01178-7

Elmouatamid A et al (2021) Review of control and energy management approaches in micro-grid systems. Energies 14:168. https://doi.org/10.3390/en14010168

Eyimaya SE, Altin N (2024) Review of energy management systems in microgrids. Appl Sci 14:1249. https://doi.org/10.3390/app14031249

Giedraityte A, Rimkevicius S, Marciukaitis M, Radziukynas V, Bakas R (2025) Hybrid renewable energy systems—A review of optimization approaches and future challenges. Appl Sci 15:1744. https://doi.org/10.3390/app15041744

Gulzar MM, Iqbal A, Sibtain D (2023) An innovative converterless solar PV control strategy for a grid connected hybrid PV/wind/fuel-cell system coupled with battery energy storage. IEEE Access 11:23245–23259. https://doi.org/10.1109/ACCESS.2023.3252891

Ibrahim O et al (2023) Development of fuzzy logic-based demand-side energy management system for hybrid energy sources. Energy Convers Manag X 18:100354. https://doi.org/10.1016/j.ecmx.2023.100354

Jurasz J, Guezgouz M, Campana PE, Kies A (2022) On the impact of load profile data on the optimization results of off-grid energy systems. Renew Sustain Energy Rev 159:112199. https://doi.org/10.1016/j.rser.2022.112199

Kafando JG, Yamegueu D, Houdji ET (2024) Review on sizing and management of stand-alone PV/wind systems with storage. Heliyon 10:e38080. https://doi.org/10.1016/j.heliyon.2024.e38080
Kechida A et al (2024) Smart control and management for a renewable energy based stand-alone hybrid system. Sci Rep 14:32039. https://doi.org/10.1038/s41598-024-83826-1
Kermani M et al (2021) Intelligent energy management based on SCADA system in a real microgrid for smart building applications. Renew Energy 171:1115–1127. https://doi.org/10.1016/j.renene.2021.03.008
Khan AA, Minai AF, Pachauri RK, Malik H (2022) Optimal sizing, control, and management strategies for hybrid renewable energy systems: A comprehensive review. Energies 15:6249. https://doi.org/10.3390/en15176249
Kumar P, Kumar K, Adhikary N, Tesfaye EL (2025) Analysis of control and computational strategies for green energy integration for sociotechnical ecological power infrastructure in Indian and African markets. Sci Rep 15. https://doi.org/10.1038/s41598-025-96773-2
Magsumbol J-AV et al (2025) Intelligent battery management in a hybrid photovoltaic using fuzzy logic system. Technologies 13:214. https://doi.org/10.3390/technologies13060214
Menzri F et al (2024) Applications of novel combined controllers for optimizing grid-connected hybrid renewable energy systems. Sustainability 16:6825. https://doi.org/10.3390/su16166825
Ramesh S et al (2025) Comparative analysis of Q-learning, SARSA, and deep Q-network for microgrid energy management. Sci Rep 15:694. https://doi.org/10.1038/s41598-024-83625-8
Riaz MU et al (2024) Advanced energy management in a sustainable integrated hybrid power network using a computational intelligence control strategy. Energies 17:5040. https://doi.org/10.3390/en17205040
Roy P, He J, Zhao T, Singh YV (2022) Recent advances of wind-solar hybrid renewable energy systems for power generation: A review. IEEE Open J Ind Electron Soc 3:81–104. https://doi.org/10.1109/OJIES.2022.3144093
Salehi N, Martinez-Garcia H, Velasco-Quesada G, Guerrero JM (2022) A comprehensive review of control strategies and optimization methods for individual and community microgrids. IEEE Access 10:15935–15955. https://doi.org/10.1109/ACCESS.2022.3142810
Sayed ET et al (2023) Renewable energy and energy storage systems. Energies 16:1415. https://doi.org/10.3390/en16031415
Shafiullah M et al (2022) Review of recent developments in microgrid energy management strategies. Sustainability 14:14794. https://doi.org/10.3390/su142214794
Singh A, Sharma A, Rajput S, Bose A, Hu X (2022) An investigation on hybrid particle swarm optimization algorithms for parameter optimization of PV cells. Electronics 11:909. https://doi.org/10.3390/electronics11060909
Soliman MS et al (2021) Supervisory energy management of a hybrid battery/PV/tidal/wind sources integrated in DC-microgrid energy storage system. Energy Rep 7:7728–7740. https://doi.org/10.1016/j.egyr.2021.11.056
Talaat FM, Kabeel AE, Shaban WM (2025) Powering a sustainable future: AI-driven integration of renewables for optimized grid management. Sustain Futures 10:100821. https://doi.org/10.1016/j.sftr.2025.100821
Ukoba K, Olatunji KO, Adeoye E, Jen T-C, Madyira DM (2024) Optimizing renewable energy systems through artificial intelligence: Review and future prospects. Energy Environ 35:3833–3879. https://doi.org/10.1177/0958305X241256293
Zaini FA, Sulaima MF, Razak IAWA, Zulkafli NI, Mokhlis H (2023) A review on the applications of PSO-based algorithm in demand side management: Challenges and opportunities. IEEE Access 11:53373–53400. https://doi.org/10.1109/ACCESS.2023.3278261

Statements and Declarations

Data and code availability. No measured field data were used. All weather, load and outage series come from the stochastic generator of Section 3.2 with fixed random seeds, so the complete dataset can be regenerated with the supplied code. The source code (plant models, controllers, PSO, aging and economics, experiment driver, tuned-FLC and training-year experiments, and figure generation), the fixed seed schedule, the raw per-year outputs (results.pkl, exp_flc.pkl, exp_train.pkl) and a README giving the run order are deposited in a public Zenodo repository. The repository DOI is withheld from this anonymized manuscript for double-anonymous review; it is provided to the editor in the response letter and on the separate title page, and access can be arranged for reviewers through the editorial office.

Generative AI. The authors used ChatGPT (OpenAI) for drafting and grammar correction as described in Section 3.9, and take full responsibility for the content.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- README.pdf