

A two-stage MPPT controller with energy-regret coordination and transition memory for PV systems under complex partial shading conditions

Mingcheng Ye

Guangdong Ocean University

Limei Shi

Guangdong Ocean University

Peng Luo

luopeng@gdou.edu.cn

Guangdong Ocean University

Junyu Chen

Guangdong Ocean University

Lei Guo

Guangdong Ocean University

Yigeng Sun

Guangdong Power Grid Co., Ltd

Article

Keywords: partial shading, hybrid algorithm, energy-regret-guided coordination, transition-memory warm restart, global maximum power point tracking, photovoltaic maximum power point tracking

Posted Date: October 8th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10946262/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

A two-stage MPPT controller with energy-regret coordination and transition memory for PV systems under complex partial shading conditions

Mingcheng Ye¹, Limei Shi², Peng Luo^{2*}, Junyu Chen², Lei Guo²,
Yigeng Sun³

¹College of Mechanical Engineering, Guangdong Ocean University,
Zhanjiang, 524088, China.

²College of Electronic and Information Engineering, Guangdong Ocean
University, Zhanjiang, 524088, China.

³Zhanjiang Power Supply Bureau, Guangdong Power Grid Co., Ltd,
Zhanjiang, 524200, China.

*Corresponding author(s). E-mail(s): luopeng@gdou.edu.cn;

Abstract

Partial shading creates multiple peaks in the photovoltaic power-voltage characteristic and complicates maximum power point tracking (MPPT). Existing two-stage methods often ignore online probing losses and restart randomly after recurrent transitions. This study proposes ER-ISSA-IVY, a two-stage hybrid MPPT framework that combines an improved sparrow search algorithm (ISSA) for global exploration with the Ivy algorithm (IVY) for local refinement. An energy-regret-guided (ER) rule switches stages by comparing the uncertainty-adjusted marginal return of further exploration with the measured cost of off-MPP probing, and a training-free transition memory reuses verified peaks for recurrent power-change signatures while a globally initialized subpopulation preserves coverage of unseen conditions. Experiments show that ER-ISSA-IVY achieves **100%** static efficiency, **94.35%** dynamic efficiency under step irradiance, and **99.99%** day-integrated efficiency on a satellite-derived profile, with the fewest first-hit evaluations among three recent hybrids, **0.77%** overshoot and **0.07%** ripple at the switching level, and complementary ablation-confirmed gains from both mechanisms under hidden change boundaries.

Keywords: partial shading, hybrid algorithm, energy-regret-guided coordination, transition-memory warm restart, global maximum power point tracking, photovoltaic maximum power point tracking

1 Introduction

Photovoltaic (PV) generation is an important component of the transition to renewable energy, and its energy yield depends on operation near the maximum power point (MPP) of the array. At the network level, the optimal placement and planning of PV generation and storage under uncertainty have received extensive attention [1, 2]. At the conversion interface, a DC–DC converter executes MPPT commands while matching the PV voltage to the downstream bus. Continuous-input-current ultrahigh-step-up converters can facilitate MPPT in high-gain PV interfaces [3], and three-port converters can additionally route power independently among renewable sources, storage units, and DC buses [4]. Despite these advances in system planning and power conversion, maximum energy extraction still depends on accurate selection of the PV operating point. Under uniform irradiance, the power–voltage (P–U) characteristic is unimodal, so conventional perturb-and-observe (P&O) and incremental-conductance (INC) methods provide practical MPP tracking; recent P&O refinements further reduce oscillation and improve robustness [5]. Field arrays, however, are frequently exposed to partial shading caused by buildings, clouds, and dust, as well as electrical mismatch among modules. The resulting current mismatch activates bypass diodes and creates multiple local peaks in addition to the global peak [6]. Conventional hill-climbing methods can then converge to a local peak and forfeit a substantial fraction of the available energy, particularly when shading changes rapidly [7, 8]. Reliable global MPP tracking is therefore essential under partial-shading conditions (PSCs).

Swarm-intelligence methods address this multimodal problem by searching beyond a local neighbourhood. Representative methods include particle swarm optimization (PSO) and its stability-oriented MPPT variants [9], the grey wolf optimizer (GWO) [10], the sparrow search algorithm (SSA) [11], and the whale optimization algorithm (WOA). Simulated annealing, the Archimedes optimization algorithm, and other population-based updates have also been used to improve global localization under PSCs [12–14]. Beyond direct duty-cycle search, metaheuristics have also been applied offline to tune PI regulators for MPPT and dynamic-voltage control of grid-connected PV plants [15]. More recent developments include a modified butterfly optimization method evaluated under partial shading, uniform shading, and rapid load variation [16], and a single-sensor global MPPT method for partially shaded PV systems [17]. Although these search methods increase the probability of locating the global MPP (GMPP), global search alone does not resolve the balance between exploration and exploitation. Exploration-oriented methods may converge slowly or continue to oscillate after locating the GMPP, whereas strongly exploitative methods may converge prematurely to a local peak [14]. Effective GMPPT therefore requires both broad initial coverage and stable local refinement.

Hybrid and adaptive strategies seek this balance by combining global localization with local tracking. Representative recent designs include fusion firefly search [18], Salp-swarm optimization [19], hierarchical pigeon-inspired optimization [20], and fuzzy tracking strategies [21]. Spline-based, simplex-based, and predictive methods provide alternative mechanisms for global localization and local refinement [7, 22, 23], while dynamic-particle and repeatable-irradiance trackers improve recovery under changing PSCs [8, 24]. For example, the recent PO-WOA method combines variable-step P&O

with the Whale Optimization Algorithm [25]. Such global–local combinations represent a common two-stage architecture, and the same pairing of complementary search operators has proven effective in wider power-system optimization under renewable uncertainty [26]. Two limitations nevertheless remain in many implementations. The global stage is often terminated according to an iteration count or a geometric measure of stagnation, without determining whether further probing is justified by the energy that it may recover. A random restart also treats each detected environmental change as a new problem, even when the same transition has occurred previously.

Learning-based GMPPT studies indicate that recurrent operating conditions can provide useful prior information. Reinforcement-learning trackers reuse information acquired from previous states, but depend on a learned value function or policy [27]. In parallel, dynamic methods based on adaptive control, incremental conductance, and predictive control improve adaptation to changing conditions but do not explicitly retain transition-specific information [23, 28, 29]. Moreover, neither class of method directly relates the decision to continue global exploration to the energy harvested during candidate evaluation. Replacing only the global optimizer, for example with Tuna Swarm Optimization [30] or variable-step GWO/WOA [25], does not address both limitations simultaneously.

This distinction is important because online MPPT is an energy-harvesting process rather than an abstract optimization task. Each candidate is applied to the DC–DC converter and delivers power during its evaluation. Probing away from the MPP after an environmental change therefore incurs a measurable energy loss that is not represented by a static global-success rate. Consequently, two methods with the same static success can harvest different amounts of energy during repeated re-tracking. In addition, a converged swarm can lose diversity and respond slowly after the GMPP moves. Dynamic MPPT therefore benefits from an energy-based stage-transition criterion and a mechanism for reusing verified recurrent transitions, provided that global coverage is retained for previously unseen changes.

To address these limitations, this study proposes the energy-regret-guided improved sparrow search algorithm–Ivy hybrid MPPT method (ER-ISSA-IVY). The method combines energy-regret-guided stage coordination with a transition-aware restart. The complete terminal ISSA population is passed to a contracting Ivy refinement stage, rather than reducing the global-search output to a single best point. The transition between the two stages is determined by the estimated marginal energy return of continued exploration relative to its measured probing loss. After an environmental change, a training-free transition memory reuses a verified peak through a transition-conditioned partial warm restart, while a reserved globally initialized subpopulation preserves search diversity. The main contributions are:

1. Hybrid ISSA–IVY architecture. A two-stage ER-ISSA-IVY hybrid MPPT algorithm is developed in which an improved sparrow search performs global peak localization and a contracting Ivy search provides low-ripple local refinement. The complete terminal ISSA population is retained at the stage transition instead of restarting the local stage from a single best point.

2. Energy-regret-guided coordination. An energy-regret-guided rule determines the transition between the global and local stages. Unlike fixed-iteration or stagnation-based coordination, the rule compares the uncertainty-adjusted marginal energy return of further exploration with its measured off-MPP probing cost.
3. Transition-memory warm restart. A training-free mechanism associates each power-change signature with a verified post-change peak. A matched recurrent transition activates a partial warm restart, while a reserved globally initialized subpopulation maintains coverage of unseen shading conditions.
4. Mechanism and hybrid validation. Static, dynamic, switching-level, historical-profile, hidden-boundary, and ablation experiments evaluate the roles of hybridization, energy-regret coordination, and transition memory under common function-evaluation budgets.

The remainder of this paper is organized as follows. Section 2 describes the PV array model under partial shading. Section 3 presents the ER-ISSA-IVY search architecture, energy-regret-guided stage coordination, and transition-memory restart. Section 4 reports the static, dynamic, switching-level, historical-profile, recent-hybrid, and ablation results. Section 5 discusses the evidence and limitations, and Section 6 summarizes the main findings.

2 PV array modelling under partial shading

2.1 PV module model and environmental correction

PV modules are commonly represented by an implicit single-diode model whose unknown parameters are typically extracted by iterative or metaheuristic identification [31], and whose evaluation requires numerical root finding. Because MPPT repeatedly evaluates the power map, this study uses a closed-form engineering model derived from the rated electrical points to reduce evaluation cost.

The model uses the rated values U_{oc} , I_{sc} , and (U_m, I_m) at $S_{ref} = 1000 \text{ W m}^{-2}$ and $T_{ref} = 25^\circ\text{C}$.

At these standard test conditions (STC), the module current is

$$I(U) = I_{sc} \left\{ 1 - C_1 \left[\exp\left(\frac{U}{C_2 U_{oc}}\right) - 1 \right] \right\} \quad (1)$$

where $I(U)$ is the module current at voltage U , and C_1 and C_2 are dimensionless shape coefficients determined from the rated maximum-power point:

$$\begin{aligned} C_2 &= \frac{U_m/U_{oc} - 1}{\ln(1 - I_m/I_{sc})}, \\ C_1 &= \left(1 - \frac{I_m}{I_{sc}}\right) \exp\left(-\frac{U_m}{C_2 U_{oc}}\right). \end{aligned} \quad (2)$$

Equations (1) and (2) construct the characteristic from the rated short-circuit, open-circuit, and maximum-power points without iterative parameter identification.

Table 1 PV array and boost-converter parameters.

Parameter	Value	Parameter	Value
Module U_{oc}	87.2 V	Modules in series	5
Module I_{sc}	33.4 A	Bypass-diode drop	0.7 V
Module U_m	70.0 V	Boost inductor L	1.5 mH
Module I_m	30.4 A	Input capacitor C_{in}	220 μ F
I_{sc} coefficient a	0.0025 / $^{\circ}$ C	Reference temperature T_{ref}	25 $^{\circ}$ C
Irradiance coefficient b	0.5	Control period T_c	20 / 5 ms [†]
U_{oc} coefficient c	0.00288 / $^{\circ}$ C	Integration step	2 μ s

[†]20 ms for the static and dynamic $P(U)$ experiments; 5 ms for the switching-level boost-converter simulation.

Table 1 summarizes the module ratings, environmental-correction coefficients, series-string configuration, and converter parameters used in the simulations. The coefficients a and c set the temperature corrections of I_{sc} and U_{oc} , respectively, while b controls the irradiance correction of U_{oc} . The control periods distinguish the direct $P(U)$ evaluations (20 ms) from the switching-level converter test (5 ms).

Irradiance and temperature modify the two intercepts. For irradiance S and module temperature T , the corrected ratings are

$$I'_{sc} = I_{sc}s(1 + a\Delta T) \quad (3)$$

$$U'_{oc} = U_{oc} \ln(e + b\Delta s)(1 - c\Delta T) \quad (4)$$

where the primed quantities denote corrected ratings, $s = S/S_{ref}$, $\Delta s = s - 1$, and $\Delta T = T - T_{ref}$. Substitution into Eq. (1) gives the environmental characteristic, whose inverse is

$$U(I; S, T) = C_2 U'_{oc} \ln \left(1 + \frac{1 - I/I'_{sc}}{C_1} \right) \quad (5)$$

Equation (5) provides the module voltage at a prescribed current without numerical root finding.




Fig. 1 Module characteristics at STC (1000 W m⁻², 25 °C): (a) I - U . (b) P - U .

Figure 1 compares the resulting I - U and P - U curves with the four STC ratings. Panel (a) checks the short-circuit current, open-circuit voltage, and rated maximum-power operating point, while panel (b) locates that operating point on the power curve. This comparison checks consistency with the rated points used to construct the model.

2.2 Series array under partial shading

Series-connected modules carry the same current. When I exceeds the corrected short-circuit current of a shaded module, its bypass diode conducts and clamps the module

voltage near $-V_D$. For N_s modules, the string voltage is

$$U_{arr}(I) = \sum_{i=1}^{N_s} u_i(I),$$

$$u_i(I) = \begin{cases} U_i(I; S_i, T_i), & I \leq I'_{sc,i}, \\ -V_D, & I > I'_{sc,i}. \end{cases} \quad (6)$$

where $U_i(I; S_i, T_i)$ follows Eq. (5), and V_D is the bypass-diode drop. The string power and maximizing current are

$$P_{arr}(I) = I U_{arr}(I), \quad I_G = \arg \max_{0 \leq I \leq \max_i I'_{sc,i}} P_{arr}(I) \quad (7)$$

The GMPP is defined by $U_G = U_{arr}(I_G)$ and $P_G = P_{arr}(I_G)$. Different bypass states form distinct characteristic segments, allowing nonuniform irradiance to produce multiple local maxima.

The simulated array comprises five series-connected modules with individual bypass diodes.




Fig. 2 Five-module series string with per-module bypass diodes: (a) uniform irradiance (S1). (b) partial shading (S2). (c) severe shading (S3).

Figure 2 shows the module-level irradiance patterns and bypass states for uniform irradiance (S1), partial shading (S2), and severe shading (S3). The bypass paths allow the string current to continue when a shaded module can no longer support it, reducing that module's contribution to the string voltage.




Fig. 3 P–U characteristics of the five-module PV array: (a) uniform irradiance (S1). (b) partial shading (S2). (c) severe shading (S3).

Figure 3 shows the corresponding P–U characteristics. Uniform irradiance produces a single peak, whereas unequal module irradiances introduce distinct local peaks through changes in bypass state. Increasing mismatch can shift the GMPP to a low-voltage, high-current region, explaining why a local hill-climbing tracker may settle at a suboptimal peak.

2.3 Interface with the MPPT controller

In the direct $P(U)$ experiments, each candidate voltage U is evaluated using the characteristic generated from Eqs. (1)–(7). In the switching-level experiment, a duty cycle D is applied, and the converter dynamics determine (U_{pv}, I_{pv}) . All methods use $P = U_{pv}I_{pv}$ as the fitness and share the same module parameters, shading patterns, environmental corrections, and bypass logic. The pair (U_G, P_G) is used only as an offline reference for performance assessment.

3 The energy-regret-guided ISSA–IVY hybrid algorithm

3.1 Improved sparrow search algorithm (ISSA)

The standard SSA [11] divides the population into producers, scroungers, and scouts. The ISSA stage incorporates three modifications. First, *Skew-Tent chaotic initialization* generates the initial population with a skew-Tent map to improve coverage of the bounded search space [32]:

$$z_{k+1} = \begin{cases} z_k/\alpha, & z_k < \alpha \\ (1 - z_k)/(1 - \alpha), & z_k \geq \alpha. \end{cases} \quad (8)$$

where $z_k \in [0, 1]$ is the normalized state at map iteration k , z_{k+1} is the subsequent state, and $\alpha \in (0, 1)$ is the skew parameter that separates the two branches. Second, the *nonlinear inertia weight* $w = 0.2 + 0.7((M - t)/M)^2$, where t is the current optimization iteration and M is the maximum number of iterations, decreases over the search to favour broad producer motion initially and finer updates later. Third, *Lévy-flight producers and global-best-guided mutation* expand the exploratory range

and direct poorly performing individuals towards the current best, thereby reducing the risk of premature stagnation [32].

3.2 The Ivy algorithm (IVY)

IVY [33] models the growth and spreading behaviour of ivy plants. Position updates depend on individual fitness and a growth vector GV ; the neighbour-guided update used for fitter individuals is

$$X_i^{new} = X_i + |N(0, 1)| \odot (X_{i+1} - X_i) + N(0, 1) \odot GV_i \quad (9)$$

$$GV_i \leftarrow GV_i \odot (r_2 N(0, 1)). \quad (10)$$

where X_i and X_i^{new} are the current and updated position vectors of individual i , respectively; X_{i+1} is the position of a neighbouring individual; GV_i is the growth vector; $N(0, 1)$ denotes an independently sampled standard normal random vector; $|\cdot|$ and $\odot$ denote element-wise absolute value and multiplication, respectively; and $r_2 \in [0, 1]$ is an independent random scaling factor. Weaker individuals are instead directed towards the current population best. The merging, sorting, and competitive-elimination steps retain the better candidates and provide strong local exploitation. Local exploitation alone does not establish global capture on multimodal PV characteristics; the standalone IVY is therefore included in the static comparison to evaluate this limitation.




Fig. 4 Energy-regret-guided switching and transition-memory control flow of ER-ISSA-IVY.

3.3 Energy-regret-guided two-stage coordination

Figure 4 summarizes the control flow of ER-ISSA-IVY. The global ISSA stage compares the estimated return of further exploration with its measured probing-energy cost; after the minimum search budget and hysteresis conditions are met, it transfers the complete population to the contracting IVY stage. The transition-memory branch stores verified peaks and uses a matched entry for a partial warm restart, while retaining a globally initialized subpopulation for coverage. The diagram thus connects the energy-based stage switch with the restart mechanism used after environmental changes.

Stage 1 (global exploration). Let P_b^t denote the best power measured by the end of generation t , and let P_i^t denote the power delivered while candidate i is applied to the plant or simulation model. The non-negative marginal improvement and the probing-energy cost are

$$\Delta P_t = [P_b^t - P_b^{t-1}]_+. \quad (11)$$

$$C_t = T_c \sum_{i=1}^N [P_b^{t-1} - P_i^t]_+. \quad (12)$$

where ΔP_t is the improvement in the best measured power at generation t ; C_t is the probing-energy deficit accumulated while evaluating that generation relative to the preceding best power; T_c is the control period; N is the population size; and $[x]_+ = \max(x, 0)$ is the positive-part operator. Because the true GMPP is not observable online, both quantities are referenced to the best-so-far power P_b . They therefore form an *online proxy*, rather than an oracle measure, of the energy that further search may recover and the probing loss incurred to obtain it. Over a four-generation window, an uncertainty-adjusted upper estimate of the potential return is

$$G_t^U = HT_c \left(\overline{\Delta P}_t + 1.96 s_{\Delta P,t} / \sqrt{4} \right) \quad (13)$$

where G_t^U is the upper estimate of recoverable energy at generation t ; H is the expected remaining holding horizon in control samples, estimated as the mean observed dwell length minus the number of samples already spent in the current segment and lower-bounded by $H_{\min}$; and T_c converts this horizon to time. The quantities $\overline{\Delta P}_t$ and $s_{\Delta P,t}$ are the sample mean and standard deviation of ΔP over the latest four generations. The multiplier 1.96 is the two-sided Gaussian 95% quantile, but it is used only as a dispersion scale, not as an exact confidence bound, because the four increments are serially correlated and truncated at zero. After the prescribed minimum global-search budget ($t \geq 8$), the switch is triggered only when

$$G_t^U \leq \lambda \overline{C}_t \quad (14)$$

for h consecutive generations, where $\overline{C}_t$ is the mean probing-energy cost over the same four-generation window, λ is a dimensionless cost-weight coefficient, and h is the hysteresis count. Exploration therefore stops only when the estimated upper return remains no greater than λ times the measured probing cost for h successive generations. **Stage 2 (local refinement).** The complete terminal ISSA population from Stage 1 is transferred to a contracting IVY stage. This transfer preserves the candidate distribution obtained during global search instead of initializing the local stage from only the best point. The factor $\text{contract} = 1 - (t - t_{sw}) / (M - t_{sw})$ decreases linearly after the stage transition and restricts candidate motion to a shrinking neighbourhood of the best-so-far position X_{best} :

$$X_i^{new} = X_{best} + \text{contract} \cdot [N(0, 1)(X_{best} - X_i) + N(0, 1) \Delta GV_i]. \quad (15)$$

where X_i^{new} is the updated position of individual i ; X_i and X_{best} are the current and best-so-far positions, respectively; t_{sw} is the iteration at which the algorithm switches to IVY; M is the maximum iteration count; and ΔGV_i is the stochastic growth-vector perturbation of individual i . In addition, the worst individual performs a global Lévy jump with residual-exploration probability p_{exp} . Greedy retention preserves the stored best solution, although the residual Lévy evaluations can still introduce transient probing loss and ripple. These effects are included in P_{app} and η_{dyn} . The residual jump retains limited global coverage after the switch, and the resulting global-capture ability is evaluated empirically rather than assumed.

3.4 Online MPPT implementation and objective

In the switching-level implementation, the optimization variable is the boost-converter duty cycle $D \in [0.1, 0.9]$. The direct $P(U)$ experiments instead optimize the candidate operating voltage. During each control period T_c , one candidate is applied, and the sampled power $P = V_{pv} I_{pv}$ serves as both the fitness and the energy-accounting signal, where V_{pv} and I_{pv} are the measured PV terminal voltage and current.

Environmental changes are detected without access to phase labels. Three consecutive same-direction power deviations greater than 1% relative to a nine-sample rolling median trigger a new search. The resulting transition signature comprises the normalized power-change magnitude, the normalized pre-change operating point, and the short-time power slope. After the new search converges, the signature and the verified post-change peak are stored. A subsequent transition matches an archived entry only if the power-change directions agree and the scale-normalized signature distance does not exceed the fixed match radius. For a match, half of the restart population is initialized around the archived peak and its $\pm 1\%$ neighbourhood, while the other half is initialized globally. An unmatched transition uses entirely global initialization and therefore does not restrict the search space for an unseen shading pattern. Because the low-dimensional signature is not a unique representation of the spatial shading pattern, distinct conditions may still produce similar signatures. The direction constraint and multi-component distance reduce this collision risk, while the globally initialized subpopulation limits its effect.

The control objective is to maximize the energy delivered over the complete dynamic profile:

$$\max \eta_{dyn} = \frac{\sum_k P_{app}(k) T_c}{\sum_k P_{ideal}(k) T_c} \quad (16)$$

where η_{dyn} is the dynamic energy efficiency; k indexes the control periods over the evaluation horizon; and $P_{app}(k)$ and $P_{ideal}(k)$ are the delivered power and ideal GMPP power in period k , respectively. The numerator and denominator are therefore the delivered and ideal energies over the same horizon. For evaluation, the *offline oracle regret* is also reported:

$$R_K = T_c \sum_{k=1}^K [P_G(k) - P_{app}(k)]_+, \quad (17)$$

where R_K is the regret accumulated over the first K control periods, $P_G(k) = P_{ideal}(k)$ is the true GMPP power in period k , and $[\cdot]_+$ retains only positive power deficits. The simulated profiles provide P_G , so R_K can be calculated offline but is unavailable to the controller. By contrast, the switching quantities ΔP_t and C_t in Eqs. (11)–(13) are best-so-far-referenced surrogates formed from P_b . In particular, C_t underestimates the oracle probing deficit whenever $P_b < P_G$. The switching rule therefore has no guaranteed conservative relation to oracle regret. Its energy performance is instead assessed empirically using R_K in Eq. (17) and η_{dyn} in Eq. (16).

4 MPPT experiments and results

4.1 Experimental setup

The evaluation comprises three comparison sets. First, ER-ISSA-IVY is compared with conventional P&O, variable-step P&O (VS-P&O), PSO [9], GWO [10], SSA [11], ISSA, IVY [33], and Tuna Swarm Optimization (TSO) [30]. Second, a common-budget study compares ER-ISSA-IVY with three recent hybrid MPPT methods. Third, the controller-level study compares three ablated configurations, complete ER-ISSA-IVY, and PO-WOA [25]. All population-based methods use $N = 6$. P&O and VS-P&O perform six sequential perturbations per nominal generation so that every method receives the same candidate-evaluation budget. The static scenarios allow up to 50 iterations and are repeated for 30 independent runs. The dynamic scenario allows up to 15 search iterations per segment over 20 runs, followed by a longer steady-state holding phase.

The open-circuit voltage of the five-module string is approximately 436 V. The P&O perturbation is therefore set to 0.8% of the duty-cycle span, corresponding to approximately 3.7 V and remaining below the inter-peak voltage spacing. For ER-ISSA-IVY, global exploration lasts for at least eight generations, and the energy-return estimate uses a four-generation window. The switching parameters are selected on an independent calibration sequence: $\lambda = 0.50$, $h = 2$ generations, and $H_{\min} = 30$ control periods. The residual-exploration probability is $p_{exp} = 0.15$. The normalized transition-memory match radius is 1.60, with signature scales of 0.03, 0.08, and 0.015 for the power-change, operating-point, and slope components, respectively. Half of the restart population is reserved for global initialization.

Table 2 Static multi-peak performance.

Scenario	Algorithm	$\eta/\%$	Success/%	First-hit gen.	Osc. (%)
S1 (Uniform)	P&O	99.99	100	5.0	~ 0
	VS-P&O	100.00	100	19.0	~ 0
	PSO	100.00	100	1.2	~ 0
	GWO	100.00	100	2.2	0.031
	SSA	100.00	100	2.6	~ 0
	ISSA	100.00	100	0.9	~ 0
	IVY	98.34	86.7	14.6	11.50
	TSO	100.00	100	4.4	0.10
	ER-ISSA-IVY	100.00	100	0.7	~ 0
S2 (Partial)	P&O	100.00	100	5.0	~ 0
	VS-P&O	100.00	100	20.0	~ 0
	PSO	100.00	100	2.3	~ 0
	GWO	100.00	100	2.7	0.044
	SSA	100.00	100	5.4	0.033
	ISSA	100.00	100	2.2	~ 0
	IVY	99.72	90.0	16.6	11.12
	TSO	100.00	100	7.0	0.21
	ER-ISSA-IVY	100.00	100	2.1	~ 0
S3 (Severe)	P&O	64.15	0	50.0	~ 0
	VS-P&O	64.15	0	50.0	~ 0
	PSO	97.62	93.3	6.1	~ 0
	GWO	100.00	100	3.1	~ 0
	SSA	97.63	93.3	8.0	~ 0
	ISSA	100.00	100	4.8	13.15
	IVY	96.32	76.7	23.8	8.83
	TSO	100.00	100	5.3	0.03
	ER-ISSA-IVY	100.00	100	2.9	~ 0

4.2 Static multi-peak results

Table 2 reports the static results. The *first-hit generation* is the generation at which the best-so-far power first enters the prescribed GMPP band; it is not the termination generation. ER-ISSA-IVY reaches this band after mean values of 0.7, 2.1, and 2.9 generations in S1–S3, respectively, which are the lowest values in the table. Its global stage nevertheless continues for at least eight generations before the energy-regret-guided criterion can initiate local refinement. Under S1 and S2, all methods except IVY achieve 100% success, and their mean efficiencies range from 99.99% to 100.00%. The severe-shading case is more discriminating: GWO, ISSA, TSO, and ER-ISSA-IVY achieve 100% efficiency and success; PSO and SSA achieve 93.3% success; IVY achieves 76.7% success; and P&O and VS-P&O remain at the local peak with 64.15% efficiency and zero success. ER-ISSA-IVY also has the lowest mean first-hit generation in S3.




Fig. 5 Convergence curves under three static scenarios: (a) uniform irradiance. (b) partial shading. (c) severe shading.

Figure 5 compares the static convergence curves in S1–S3. The approach to the GMPP efficiency level shows how quickly each method locates the best peak; the separation under severe shading highlights the difficulty of global capture. These curves complement the mean first-hit generations and success rates in Table 2.




Fig. 6 Power-tracking trajectories under three static scenarios: (a) uniform irradiance. (b) partial shading. (c) severe shading.

Figure 6 shows the corresponding power-tracking trajectories. The trajectories show the rise towards the GMPP and the final power level reached by each method. In S3, the lower power plateau of P&O and VS-P&O illustrates local-peak entrapment, consistent with their 64.15% efficiency in Table 2.




Fig. 7 Box plots of tracking efficiency over 30 independent runs: (a) uniform irradiance. (b) partial shading. (c) severe shading.

Figure 7 summarizes the efficiency distributions over 30 independent runs. The spread and low-efficiency outcomes reveal run-to-run variability that a mean value alone cannot describe. ER-ISSA-IVY remains concentrated at the GMPP efficiency level in all three scenarios, consistent with its 100% static success, whereas the severe-shading case exposes unsuccessful capture by several baselines.

4.3 Dynamic step-irradiance results

A five-segment step-irradiance profile represents a cloud moving across the string. The sequence comprises uniform irradiance (1000 W m^{-2} on all five modules), a cloud edge (modules 4 and 5 at 600 and 400 W m^{-2}), a cloud core (modules 3–5 at 300 , 300 , and 200 W m^{-2}), a cloud tail (modules 3–5 at 400 , 400 , and 300 W m^{-2}), and a return to uniform irradiance. In segments 3 and 4, the GMPP lies in the low-voltage region while a strong local peak remains at a higher voltage. These deep-shading segments are therefore particularly difficult for local hill-climbing. Each segment contains a search

interval followed by a longer steady-state holding interval, so the controller must locate the new GMPP after each step and then maintain it.

Table 3 Dynamic step-irradiance performance.

Algorithm	$\eta_{dyn}/\%$ (mean $\pm$ s.d.)	$L_{dyn}/\%$	t_{resp}/ms	$\eta_{ss}/\%$
P&O	92.78 ± 0.05	7.22	1980.8	89.93
VS-P&O	85.40 ± 1.83	14.60	3665.4	85.17
PSO	92.48 ± 1.45	7.52	442.8	99.40
GWO	89.76 ± 0.93	10.24	525.8	99.57
SSA	83.71 ± 1.67	16.29	959.4	98.53
ISSA	93.14 ± 0.75	6.86	478.6	99.20
IVY	89.09 ± 4.54	10.91	1539.0	96.47
TSO	71.98 ± 0.75	28.02	977.8	99.63
ER-ISSA-IVY	94.35 ± 0.88	5.65	342.4	99.90

Table 3 reports the dynamic results, where $L_{dyn} = 100 - \eta_{dyn}$ is the energy-loss percentage. ER-ISSA-IVY achieves the highest dynamic energy efficiency (94.35%), the shortest mean response time (342.4 ms), and the highest steady-state efficiency (99.90%). The next-highest dynamic efficiency is obtained by ISSA (93.14%). P&O illustrates why the metrics must be interpreted jointly: its dynamic efficiency is 92.78%, but its steady-state efficiency is only 89.93% and its mean response time over the profile is 1980.8 ms. Because the boundary of each irradiance segment is supplied to the algorithms in this experiment, the results primarily evaluate the search core rather than the transition-detection and memory mechanisms.

4.4 Boost-converter circuit validation

The switching-level test evaluates the MPPT controller with the boost-converter duty cycle D as its control input. For reference, the corresponding averaged converter dynamics are

$$\begin{aligned}
C_{in} \frac{dV_{pv}}{dt} &= I_{pv}(V_{pv}, S) - i_L, \\
L \frac{di_L}{dt} &= V_{pv} - i_L R_L - (1 - D)V_{out}, \\
C \frac{dV_{out}}{dt} &= (1 - D)i_L - \frac{V_{out}}{R_{load}}.
\end{aligned} \tag{18}$$

where V_{pv} and $I_{pv}(V_{pv}, S)$ are the PV terminal voltage and current, respectively; i_L is the inductor current; C_{in} is the input capacitance; L and R_L are the inductance and series resistance of the boost inductor; D is the PWM duty cycle; V_{out} is the output voltage; and C and R_{load} are the output capacitance and load resistance. The converter supplies a resistive load of $R_{load} = 28 \Omega$ through an output capacitor of $C = 470 \mu\text{F}$. Because the output is not connected to a stiff DC bus, V_{out} floats between

approximately 340 and 550 V while remaining above V_{pv} during boost operation. The remaining parameters are $L = 1.5$ mH, $R_L = 0.1$ Ω , $C_{in} = 220$ μ F, $f_{sw} = 10$ kHz, $T_c = 5$ ms, and an integration step of 2 μ s. The constraint $i_L \geq 0$ enforces unidirectional diode current in continuous and discontinuous conduction modes.




Fig. 8 Circuit-level schematic of the boost-converter-based PV MPPT system.

Figure 8 shows the boost-converter-based MPPT system and its feedback path. The controller uses the measured PV voltage and current to evaluate power and commands the PWM duty cycle D , which changes the operating point presented to the PV array. The converter is simulated in the time domain at the pulse-width modulation (PWM) switching level to evaluate tracking under electrical transients and switching ripple.

The irradiance sequence is identical to that of the dynamic experiment. Candidate duty cycles are applied sequentially to one converter, with one candidate evaluated per control period T_c ; no parallel evaluation is used. The electrical transient is allowed to settle within each control period, and the switching ripple is calculated over the final switching cycle. Two time metrics are distinguished. The first-hit time t_{99} is the earliest time at which the best-so-far power reaches 99% of the GMPP, whereas the search-completion time t_{search} is the time at which the search ends and the final duty cycle is retained.




Fig. 9 Online MPPT waveforms on the boost converter: (a) PV output power. (b) PV terminal voltage. (c) duty cycle. Separate Zoom A and Zoom B panels show the re-tracking transient and deep-shading entrapment, respectively.

Figure 9 shows the PV output power, terminal voltage, and duty-cycle trajectories during online MPPT. The power trace indicates the energy delivered during re-tracking, while the voltage and duty-cycle traces show the associated operating-point changes. Zoom A resolves the re-tracking transient, and Zoom B highlights deep-shading entrapment, linking the converter response to the difficulty of escaping a local peak.

Table 4 Circuit-level metrics on the boost converter.

Algorithm	t_{99} (s)	t_{search} (s)	Over- shoot (%)	Ripple (%)	η_{dyn} (%)
P&O	0.342	0.365	3.13	0.18	82.11
GWO	0.019	0.390	4.06	0.15	85.80
TSO	0.040	0.390	17.39	0.43	71.73
ER-ISSA-IVY	0.102	0.408	0.77	0.07	86.87

Table 4 quantifies the first-hit time, search-completion time, overshoot, ripple, and dynamic efficiency for these converter tests. ER-ISSA-IVY achieves the lowest overshoot (0.77%) and ripple (0.07%). The corresponding values are 3.13% and 0.18% for P&O, 4.06% and 0.15% for GWO, and 17.39% and 0.43% for TSO. ER-ISSA-IVY also achieves the highest dynamic efficiency (86.87%), followed by GWO (85.80%), P&O (82.11%), and TSO (71.73%). The search-completion times lie between 0.365 and 0.408 s. GWO (0.019 s) and TSO (0.040 s) reach the 99% band sooner than ER-ISSA-IVY (0.102 s), whereas P&O reaches it later (0.342 s). The faster first hits of GWO and TSO do not translate into higher energy efficiency over the full profile. Thus, ER-ISSA-IVY provides the lowest ripple and overshoot and the highest energy capture in this switching-level test, with an intermediate first-hit time.

4.5 Simulation driven by a historical satellite-derived irradiance profile

To complement the synthetic profiles, the algorithms are evaluated using a historical daily irradiance profile from the NASA POWER database for Ain El Melh, Algeria, on 21 June 2023. The profile contains 14 daylight hours at hourly resolution. The NASA POWER surface-irradiance product is satellite-derived and model-assisted rather than measured by a ground sensor; this experiment is therefore a profile-driven simulation, not field or hardware validation. Three module-level shadow events are superimposed on the five-module string: a mid-morning single-module shadow, an afternoon two-module deep multi-peak event, and a late-afternoon shadow. The remaining hours produce single-peak characteristics. Each segment is evaluated over 20 independent runs, and the delivered energy is integrated over the day.

Table 5 Simulation driven by a historical satellite-derived irradiance profile.

Algorithm	$\eta_{dyn}/\%$	$\eta_{ss}/\%$	$\eta_{win}/\%$	E_{loss}/Wh
P&O	99.80	99.86	99.71	158.9
GWO	99.84	100.00	99.20	127.4
TSO	99.75	99.73	99.55	196.3
ISSA	99.98	99.94	100.00	5.5
IVY	97.98	97.86	97.89	1577.5
ER-ISSA-IVY	99.99	100.00	100.00	0.6

Table 5 shows that ER-ISSA-IVY achieves the highest day-integrated dynamic efficiency (99.99%), followed by ISSA (99.98%), GWO (99.84%), P&O (99.80%), TSO (99.75%), and IVY (97.98%). ER-ISSA-IVY and GWO both achieve 100.00% steady-state efficiency, whereas the remaining values range from 97.86% to 99.94%. The small separation in day-integrated efficiency is expected because most hours are clear-sky, single-peak intervals. Table 5 therefore also reports the efficiency during the three shadow-change windows and the day-integrated energy deficit relative to the true GMPP. ER-ISSA-IVY and ISSA both round to 100.00% change-window efficiency, but their daily energy deficits are 0.6 and 5.5 Wh, respectively. The corresponding deficits are 127.4 Wh for GWO, 158.9 Wh for P&O, 196.3 Wh for TSO, and 1577.5 Wh for IVY. Relative to P&O, ER-ISSA-IVY recovers approximately 158 Wh over the simulated day. Reporting both the aggregate and change-window metrics prevents the long single-peak periods from obscuring losses concentrated around irradiance transitions.

4.6 Head-to-head comparison with recent hybrid MPPT algorithms

To compare ER-ISSA-IVY with recent hybrid MPPT methods as well as classical baselines, three representative hybrids are reimplemented on the *identical* engineering PV model and shading scenarios: (i) PO-WOA, which combines variable-step P&O with the Whale Optimization Algorithm [25]; (ii) PLGSA, which combines PSO and gravitational search with Lévy flight [34]; and (iii) GWO-PSO, a low-level hybrid of GWO leader-guided exploration and PSO velocity updates [35]. Every method uses a population of $N = 6$ and the *same total number of function evaluations*. The iteration count of each method is adjusted to satisfy this common budget.

Table 6 Fair head-to-head comparison with three recent hybrid MPPT algorithms under the static scenarios.

Scenario	Method	$\eta/\%$	Success/%	First-hit (NFE)	Osc. (%)
S1 (Uniform)	PO-WOA	100.00	100.0	12.8	0.011
	PLGSA	100.00	100.0	8.4	~ 0
	GWO-PSO	100.00	100.0	21.0	0.085
	ER-ISSA-IVY	100.00	100.0	4.2	~ 0
S2 (Partial)	PO-WOA	100.00	100.0	16.1	0.023
	PLGSA	100.00	100.0	18.0	0.006
	GWO-PSO	99.99	100.0	41.8	0.277
	ER-ISSA-IVY	100.00	100.0	12.6	~ 0
S3 (Severe)	PO-WOA	100.00	100.0	28.0	0.016
	PLGSA	100.00	100.0	19.6	~ 0
	GWO-PSO	98.81	96.7	37.1	0.214
	ER-ISSA-IVY	99.97	96.7	17.4	~ 0

Static multi-peak results. Table 6 compares efficiency, success, first-hit function evaluations (NFE), and residual oscillation under the common static budget. PO-WOA and PLGSA achieve 100% mean efficiency and 100% success in S1–S3. ER-ISSA-IVY achieves 100% efficiency and success in S1 and S2, followed by 99.97% efficiency and 96.7% success in S3. GWO-PSO achieves 100% efficiency and success in S1, 99.99% efficiency and 100% success in S2, and 98.81% efficiency with 96.7% success in S3. ER-ISSA-IVY requires the fewest mean first-hit function evaluations in every scenario: 4.2, 12.6, and 17.4 in S1–S3, respectively. In S3, the corresponding value is 19.6 for PLGSA, 28.0 for PO-WOA, and 37.1 for GWO-PSO. ER-ISSA-IVY and PLGSA also retain near-zero residual oscillation, whereas GWO-PSO has the largest oscillation in all three scenarios. Thus, under the common budget, ER-ISSA-IVY provides the fastest mean initial capture, while PO-WOA and PLGSA provide the highest severe-shading success.






Fig. 10 Static convergence of ER-ISSA-IVY against three recent hybrid MPPT algorithms: (a) uniform irradiance. (b) partial shading. (c) severe shading.

Figure 10 compares the static convergence paths against the number of function evaluations in S1–S3. Using this common horizontal axis allows direct comparison of the different search operators. The early approach of ER-ISSA-IVY to the GMPP level is consistent with its lowest mean first-hit NFE in Table 6; the table separately reports the severe-shading success rates across repeated runs.



(a)



(b)

Fig. 11 Dynamic comparison with recent hybrids under the moving-cloud profile: (a) phase-normalized power trajectories during re-tracking. (b) per-run distribution of dynamic energy efficiency.

Figure 11 shows phase-normalized re-tracking trajectories in panel (a) and per-run dynamic-efficiency distributions in panel (b). The trajectories locate losses around irradiance changes, while the distributions show how consistently each method recovers energy across runs. These panels distinguish initial GMPP capture from energy recovery over the complete dynamic profile.

Table 7 Fair head-to-head comparison under the dynamic moving-cloud profile.

Method	$\eta_{dyn}/\%$	t_{resp}/ms	$\eta_{ss}/\%$
PO-WOA	91.93 ± 0.63	433.4	99.79
PLGSA	93.70 ± 1.68	566.2	98.70
GWO-PSO	82.74 ± 1.80	623.4	99.41
ER-ISSA-IVY	94.35 ± 0.88	342.4	99.90

Dynamic energy results. Table 7 compares energy capture, response time, and steady-state efficiency over the moving-cloud profile. For the moving-cloud profile, ER-ISSA-IVY achieves 94.35% dynamic energy efficiency. The corresponding values are 93.70% for PLGSA, 91.93% for PO-WOA, and 82.74% for GWO-PSO. ER-ISSA-IVY also has the shortest mean response time (342.4 ms), followed by PO-WOA (433.4 ms), PLGSA (566.2 ms), and GWO-PSO (623.4 ms). The steady-state efficiencies are 99.90% for ER-ISSA-IVY, 99.79% for PO-WOA, 99.41% for GWO-PSO, and 98.70% for PLGSA. ER-ISSA-IVY therefore ranks first in all three dynamic metrics in this common-budget comparison, while the principal weakness of GWO-PSO is the energy lost during re-tracking.

4.7 Energy-regret and transition-memory ablation

The controller-level experiment removes the known phase boundaries used in the preceding search-core benchmark. Each method must therefore detect environmental changes from the noisy power stream. All configurations use $N = 6$, at most 15 generations, a 120-sample holding horizon, $T_c = 20$ ms, 0.4% relative sensor noise, and the same 30 controller/noise runs. Four configurations isolate the energy-regret and transition-memory mechanisms on the shared ISSA/IVY search core. Stall-Random uses a stagnation-based stage switch and a global random restart. Regret-Random changes only the stage-switching rule, Stall-Memory changes only the restart rule, and ER-ISSA-IVY combines the energy-regret switch with the memory-assisted restart. The match radius, cost weight, hysteresis, and dwell floor are selected once on a separate calibration cloud sequence and are then fixed for all reported profiles. The four configurations use the same change detector, global-search budget, ISSA/IVY operators, and evaluation horizon. PO-WOA uses the same detector, population, function-evaluation budget, and horizon. It is selected as the recent-hybrid benchmark because it combines 100% static efficiency and success in S1-S3 with the shortest response time and the highest steady-state efficiency among the three recent baseline hybrids.

Three stress profiles are considered. The non-recurrent moving-cloud profile follows the five-pattern sequence in Section 4.3. The recurrent-cloud profile repeats transitions over nine segments to test whether stored transition information improves a subsequent restart. In the temperature-cycle profile, the irradiances of modules 1–5 are 1000, 1000, 800, 500, and 300 W m^{-2} , respectively. Their temperatures cycle through 65, 45, 40,

35, and 50 °C; 25 °C for all five modules; and 27, 43, 52, 55, and 32 °C. The reported metrics are dynamic efficiency, steady-state efficiency, response time, oracle energy regret, mean absolute error (MAE), root mean square error (RMSE), error standard deviation (SD), and threshold-based success.

Table 8 Hidden-boundary ablation and recent-hybrid comparison.

Profile	Method	$\eta_{dyn}/\%$	$\eta_{ss}/\%$	t_{resp}/ms	Regret/Wh
Moving cloud	Stall–Random	93.324 ± 1.107	99.629	309.9	2.885
	Regret–Random	93.717 ± 0.871	99.907	239.3	2.715
	Stall–Memory	93.324 ± 1.107	99.629	309.9	2.885
	ER-ISSA-IVY	94.017 ± 0.871	99.977	235.3	2.615
	PO-WOA	93.153 ± 0.645	99.860	289.2	2.726
Recurrent cloud	Stall–Random	93.586 ± 0.692	99.785	268.7	4.724
	Regret–Random	93.564 ± 0.733	99.774	275.3	4.740
	Stall–Memory	94.197 ± 0.681	99.849	223.0	4.477
	ER-ISSA-IVY	94.921 ± 0.731	99.906	203.9	4.274
	PO-WOA	94.121 ± 1.195	99.441	313.8	4.330
Temperature cycle	Stall–Random	94.657 ± 0.708	98.900	1177.4	2.436
	Regret–Random	94.470 ± 0.734	98.758	1221.0	2.521
	Stall–Memory	95.187 ± 0.539	98.934	1244.4	2.194
	ER-ISSA-IVY	97.609 ± 0.564	98.981	1152.6	1.090
	PO-WOA	94.688 ± 1.397	98.638	1174.4	2.422

Table 8 presents the controller-level results. ER-ISSA-IVY achieves the highest dynamic and steady-state efficiencies, the shortest mean response time, and the lowest oracle regret in every profile. Its dynamic efficiencies are 94.017%, 94.921%, and 97.609% under moving clouds, recurrent clouds, and temperature cycling, respectively. The single-mechanism ablations clarify how these gains arise. Under moving clouds, replacing the stagnation switch with the energy-regret switch raises dynamic efficiency from 93.324% for Stall–Random to 93.717% for Regret–Random. Under recurrent clouds, adding transition memory to the stagnation-based configuration raises the value from 93.586% to 94.197%. Under temperature cycling, the corresponding memory-only increase is from 94.657% to 95.187%. In all three profiles, the complete ER-ISSA-IVY configuration exceeds both single-mechanism ablations, providing numerical evidence that the two mechanisms make complementary contributions on the shared search core. Relative to PO-WOA, the dynamic-efficiency margins are 0.864, 0.800, and 2.921 percentage points, with the largest margin under temperature cycling.

Under the common population and function-evaluation budget, PO-WOA records dynamic efficiencies of 93.153%, 94.121%, and 94.688%, which are below the corresponding ER-ISSA-IVY values.

Table 9 Dynamic tracking error metrics.

Profile	Method	MAE/%	RMSE/%	SD/%	Success/%
Moving	Stall–Random	6.126	18.455	17.378	99.3
	ER-ISSA-IVY	5.556	17.901	17.012	100.0
	PO-WOA	5.975	18.010	17.122	99.3
Recurrent	Stall–Random	5.640	17.743	16.822	99.6
	ER-ISSA-IVY	5.231	17.289	16.476	98.5
	PO-WOA	5.279	17.469	16.699	98.5
Temperature	Stall–Random	5.351	14.689	13.671	74.3
	ER-ISSA-IVY	2.403	7.197	6.771	72.9
	PO-WOA	5.324	14.445	13.397	73.3

Table 9 provides a complementary continuous-error assessment for Stall–Random, ER-ISSA-IVY, and PO-WOA. ER-ISSA-IVY has the lowest MAE, RMSE, and SD in all three profiles and the highest moving-cloud success (100.0%). Stall–Random, however, has the highest threshold-based success under recurrent clouds (99.6% versus 98.5% for ER-ISSA-IVY) and temperature cycling (74.3% versus 72.9%). Thus, ER-ISSA-IVY provides the best energy, response, and continuous-error results in the reported tests, while the binary success metric does not rank it first in every profile.

5 Discussion

The experiments evaluate global capture, transient energy recovery, converter-level behaviour, and operation without known phase boundaries. In the baseline static comparison, ER-ISSA-IVY achieves 100% efficiency and success in all three scenarios and has the lowest mean first-hit generation. In the separate common-budget comparison with recent hybrids, its severe-shading efficiency and success are 99.97% and 96.7%, respectively, but it still requires the fewest mean first-hit evaluations in every static scenario. Under step irradiance, it provides the highest dynamic and steady-state efficiencies and the shortest mean response time. It also achieves the highest day-integrated efficiency in the historical-profile simulation. These results indicate that the principal advantage of the proposed coordination is rapid energy recovery across the complete tracking interval, rather than universal dominance of every static success metric.

The switching-level converter test supports the same energy-based interpretation. ER-ISSA-IVY achieves the lowest overshoot and ripple and the highest dynamic efficiency, although GWO and TSO reach the 99% band earlier. Its first-hit time is therefore intermediate rather than the shortest. In the recent-hybrid comparison, ER-ISSA-IVY ranks first in dynamic efficiency, steady-state efficiency, and response time; PO-WOA is the strongest baseline for the latter two metrics. Taken together, these results show that first-hit time alone is insufficient to characterize energy capture

because the power delivered during the remaining candidate evaluations also contributes to the dynamic efficiency.

The hidden-boundary ablation directly tests the two proposed coordination mechanisms on the same ISSA/IVY search core. In the moving-cloud profile, Regret–Random improves dynamic efficiency over Stall–Random by 0.393 percentage points, isolating the contribution of the energy-regret-guided stage switch. In the recurrent-cloud and temperature-cycle profiles, Stall–Memory improves on Stall–Random by 0.611 and 0.530 percentage points, respectively, isolating the contribution of the memory-assisted restart. The complete ER-ISSA-IVY configuration outperforms both single-mechanism variants in every profile and achieves the lowest oracle regret. Among the methods reported in Table 9, it also achieves the lowest MAE, RMSE, and SD. Stall–Random retains slightly higher threshold-based success under recurrent clouds and temperature cycling, showing that the binary success metric and the continuous energy-error metrics emphasize different aspects of tracking. The ablation therefore supports complementary, rather than interchangeable, roles for energy-regret coordination and transition memory.

Table 10 Structural and computational comparison of the evaluated MPPT methods.

Method	N	Evals/ search	Persistent state	Tuned parameters	Reference
ER-ISSA-IVY	6	$\leq NG$	population + transition archive, $O(N+A)$	5	This work
P&O	1	budget-matched	last sample (V, P) , $O(1)$	1	[5]
GWO	6	$\leq NG$	population, $O(N)$	1	[10]
ISSA	6	$\leq NG$	population, $O(N)$	3	[32]
PO-WOA	6	$\leq NG$	population + P&O state, $O(N)$	3	[25]
IVY	6	$\leq NG$	population + growth vectors, $O(N)$	2	[33]

Table 10 summarizes the implementation structure of the principal methods without assigning subjective method ratings. The tuning-parameter count includes method-specific control coefficients but excludes the common population size, evaluation budget, and fixed normalization scales. The symbols G and A denote the permitted number of generations per search and the number of stored transition entries, respectively. All methods use the same PV voltage and current measurements and require no offline training; their main differences are the search state and method-specific control parameters. In each control period, one candidate is applied and the controller performs only scalar position updates and the energy-accounting comparisons of Eqs. (11)–(14); no matrix operations, gradients, or offline training are involved, so the per-period computation of every evaluated method, including ER-ISSA-IVY, is compatible with the 5–20 ms control periods used here and with typical embedded MPPT controllers.

P&O has the smallest state and parameter count but loses global capture in the deceptive multi-peak case (Table 2). The population methods require $O(N)$ persistent state. ER-ISSA-IVY additionally stores a transition archive, giving $O(N+A)$ state, and uses several coordination parameters. This additional state is accompanied by the highest energy metrics in the dynamic (Table 3), boost-converter (Table 4), and hidden-boundary (Tables 8–9) tests, without additional sensors or offline training.

6 Conclusion

This study presents ER-ISSA-IVY for MPPT under dynamic partial shading. The two-stage method transfers the complete ISSA population to a contracting Ivy refinement stage. An energy-regret rule balances expected exploration gains against measured probing cost to govern this transition. A training-free transition memory reuses verified peaks through partial warm restarts while retaining global coverage for unseen conditions. Across static, dynamic, profile-driven, switching-level, and hidden-boundary tests, ER-ISSA-IVY achieved the strongest overall results among the evaluated baselines. Representative efficiencies were 100% under static conditions, 94.35% under step irradiance, and 99.99% over the satellite-derived daily profile. Under common budgets, the method required the fewest mean first-hit evaluations in S1–S3. Switching-level tests also showed the lowest overshoot and ripple, although the first-hit time was intermediate. Ablation results show that energy-regret coordination mainly improves moving-cloud tracking, whereas transition memory benefits recurrent-cloud and temperature-cycle conditions. The complete controller outperformed both single-mechanism variants across all tested profiles, supporting their complementary roles within the evaluated conditions. The present results are based on detailed PV and switching-level converter simulations under representative static, dynamic, and satellite-derived profiles and therefore do not constitute field validation. Future work will evaluate the complete controller through embedded hardware experiments and long-term field operation, with particular attention to sensor uncertainty and transition-memory scalability.

Declarations

Funding. This work was supported in part by the National Natural Science Foundation of China under Grant 62272109 and in part by the Guangxi Zhuang Autonomous Region Key Technologies R&D Program under Grant GK AB23026069.

Competing interests. The authors declare that they have no competing interests.

Ethics approval and consent to participate. Not applicable.

Consent for publication. Not applicable.

Data availability. The datasets analyzed during the current study are available from the corresponding author on reasonable request.

Materials availability. Not applicable.

Code availability. The simulation code is available from the corresponding author on reasonable request.

Author contributions. M.Y.: conceptualization, methodology, software, and writing—original draft. L.S.: supervision, project administration, funding acquisition, and validation. P.L.: writing—review and editing, and investigation. J.C.: data curation and visualization. L.G.: formal analysis and investigation. Y.S.: resources and validation. All authors reviewed the manuscript.

References

- [1] Ullah, Z., Elkadeem, M. R., Wang, S., Sharshir, S. W. & Azam, M. Planning optimization and stochastic analysis of RE-DGs for techno-economic benefit maximization in distribution networks. *Internet of Things* **11**, 100210 (2020).
- [2] Ullah, Z. *et al.* Advanced optimization of renewables and energy storage in power networks using metaheuristic technique with voltage collapse proximity and dynamic thermal rating technology. *Journal of Energy Storage* **108**, 115005 (2025).
- [3] Li, F., He, J., Huang, D., Luo, P. & Jiang, H. Synchronous dual-switch ultrahigh step-up DC–DC converter based on coupled inductor and voltage multiplier for photovoltaic systems. *IEEE Transactions on Industrial Electronics* **71**, 4807–4817 (2024).
- [4] Luo, P., Liang, T.-J., Chen, K.-H., Chen, S.-M. & Chen, J.-F. Syntheses of three-port DC–DC converters. *IEEE Transactions on Power Electronics* **38**, 16196–16207 (2023).
- [5] Kahani, R., Jamil, M. & Iqbal, M. T. An improved perturb and observed maximum power point tracking algorithm for photovoltaic power systems. *Journal of Modern Power Systems and Clean Energy* **11**, 1165–1175 (2023).
- [6] Yang, B. *et al.* Comprehensive overview of maximum power point tracking algorithms of PV systems under partial shading condition. *Journal of Cleaner Production* **268**, 121983 (2020).
- [7] Ostadrahimi, A. & Mahmoud, Y. Novel spline-MPPT technique for photovoltaic systems under uniform irradiance and partial shading conditions. *IEEE Transactions on Sustainable Energy* **12**, 524–532 (2021).
- [8] Zhao, Z. *et al.* A dynamic particles MPPT method for photovoltaic systems under partial shading conditions. *Energy Conversion and Management* **220**, 113070 (2020).
- [9] Kennedy, J. & Eberhart, R. Particle swarm optimization. *Proceedings of ICNN'95 – International Conference on Neural Networks* **4**, 1942–1948 (1995).
- [10] Mirjalili, S., Mirjalili, S. M. & Lewis, A. Grey wolf optimizer. *Advances in Engineering Software* **69**, 46–61 (2014).
- [11] Xue, J. & Shen, B. A novel swarm intelligence optimization approach: sparrow search algorithm. *Systems Science & Control Engineering* **8**, 22–34 (2020).

- [12] Lyden, S. & Haque, M. E. A simulated annealing global maximum power point tracking approach for PV modules under partial shading conditions. *IEEE Transactions on Power Electronics* **31**, 4171–4181 (2016).
- [13] Sajid, I. *et al.* Archimedes optimization algorithm (AOA)-based global maximum power point tracking for a photovoltaic system under partial and complex shading conditions. *Energy* **283**, 129169 (2023).
- [14] Pal, R. S. & Mukherjee, V. A novel population based maximum point tracking algorithm to overcome partial shading issues in solar photovoltaic technology. *Energy Conversion and Management* **244**, 114470 (2021).
- [15] Hasanien, H. M. Performance improvement of photovoltaic power systems using an optimal control strategy based on whale optimization algorithm. *Electric Power Systems Research* **157**, 168–176 (2018).
- [16] Shams, I., Mekhilef, S. & Tey, K. S. Maximum power point tracking using modified butterfly optimization algorithm for partial shading, uniform shading, and fast varying load conditions. *IEEE Transactions on Power Electronics* **36**, 5569–5581 (2021).
- [17] Zhou, G., Bi, Q., Tian, Q., Leng, M. & Xu, G. Single sensor based global maximum power point tracking algorithm of PV system with partial shading condition. *IEEE Transactions on Industrial Electronics* **69**, 2669–2683 (2022).
- [18] Huang, Y.-P., Huang, M.-Y. & Ye, C.-E. A fusion firefly algorithm with simplified propagation for photovoltaic MPPT under partial shading conditions. *IEEE Transactions on Sustainable Energy* **11**, 2641–2652 (2020).
- [19] Mirza, A. F., Mansoor, M., Ling, Q., Yin, B. & Javed, M. Y. A salp-swarm optimization based MPPT technique for harvesting maximum energy from PV systems under partial shading conditions. *Energy Conversion and Management* **209**, 112625 (2020).
- [20] Zhao, Z. *et al.* Hierarchical pigeon-inspired optimization-based MPPT method for photovoltaic systems under complex partial shading conditions. *IEEE Transactions on Industrial Electronics* **69**, 10129–10143 (2022).
- [21] Rajesh, R. & Mabel, M. C. Design and real time implementation of a novel rule compressed fuzzy logic method for the determination operating point in a photovoltaic system. *Energy* **116**, 140–153 (2016).
- [22] Song, P. Y., Liu, Y. H., Wang, S. C. & Pai, H. Y. A novel global maximum power point tracking algorithm based on Nelder-Mead simplex technique for complex partial shading conditions. *Applied Energy* **321**, 119380 (2022).
- [23] Zhao, L. & Yin, L. Multi-step depth model predictive control for photovoltaic maximum power point tracking under partial shading conditions. *International Journal of Electrical Power & Energy Systems* **151**, 109196 (2023).

- [24] Wang, S.-C. A novel and efficient global maximum power tracking method for photovoltaic systems under complicated partial shading with repeatable irradiance conditions. *Applied Energy* **383**, 125406 (2025).
- [25] Zemmit, A. *et al.* Gwo and woa variable step mppt algorithms-based pv system output power optimization. *Scientific Reports* **15**, 7810 (2025).
- [26] Ullah, Z., Wang, S., Radosavljević, J. & Lai, J. A solution to the optimal power flow problem considering WT and PV generation. *IEEE Access* **7**, 46763–46772 (2019).
- [27] Gao, F., Hu, R. & Yin, L. Variable boundary reinforcement learning for maximum power point tracking of photovoltaic grid-connected systems. *Energy* **264**, 126278 (2023).
- [28] Manna, S., Akella, A. K. & Singh, D. K. Novel lyapunov-based rapid and ripple-free MPPT using a robust model reference adaptive controller for solar PV system. *Protection and Control of Modern Power Systems* **8**, 13 (2023).
- [29] Shang, L., Guo, H. & Zhu, W. An improved MPPT control strategy based on incremental conductance algorithm. *Protection and Control of Modern Power Systems* **5**, 14 (2020).
- [30] Fu, C. & Zhang, L. A novel method based on tuna swarm algorithm under complex partial shading conditions in PV system. *Solar Energy* **248**, 28–40 (2022).
- [31] Hasanien, H. M. Shuffled frog leaping algorithm for photovoltaic model identification. *IEEE Transactions on Sustainable Energy* **6**, 509–515 (2015).
- [32] Awadallah, M. A., Al-Betar, M. A., Doush, I. A., Makhadmeh, S. N. & Al-Naymat, G. Recent versions and applications of sparrow search algorithm. *Archives of Computational Methods in Engineering* **30**, 2831–2858 (2023).
- [33] Ghasemi, M. *et al.* Optimization based on the smart behavior of plants with its engineering applications: Ivy algorithm. *Knowledge-Based Systems* **295**, 111850 (2024).
- [34] Wei, L. & Li, Y. Research on maximum power point tracking of photovoltaic power generation based on improved hybrid optimization algorithm. *Scientific Reports* **15**, 3991 (2025).
- [35] Chtita, S. *et al.* A novel hybrid GWO-PSO-based maximum power point tracking for photovoltaic systems operating under partial shading conditions. *Scientific Reports* **12**, 10637 (2022).