An International Journal of Optimization and Control: Theories & Applications
ISSN: 2146-0957 eISSN: 2146-5703
Vol.x, No.x, pp.xx-xx (2026)
https://doi.org/10.36922/IJOCTA026400306

RESEARCH ARTICLE

Structural feasibility handling in non-convex economic dispatch: Integrating fuzzy C-means clustering with a diagnostically-redesigned genetic algorithm

Rifki Rahman Nur Ikhsan[1,2], Efri Suhartono[1,2], Jangkung Raharjo[1,2*], Lindiasari Martha Yustika[1,2], and I Gede Putu Oka Indra Wijaya[1,2]

[1]School of Electrical Engineering, Telkom University, Jl. Telekomunikasi No. 1, Terusan Buah Batu, Bandung, West Java, Indonesia
[2]Center of Excellence of Sustainable Energy and Climate Change, Telkom University, Jl. Telekomunikasi No. 1, Terusan Buah Batu, Bandung, West Java, Indonesia
rfkrhmn@telkomuniversity.ac.id, esuhartono@telkomuniversity.ac.id, jangkungraharjo@telkomuniversity.ac.id, lindiasarimartha@telkomuniversity.ac.id, igpoindrawijaya@telkomuniversity.ac.id

ARTICLE INFO

Article history:
Received: September 29, 2026
Revised: October 20, 2026
Accepted:October 21, 2026
Published Online: November 10, 2026

Keywords:
Computational intelligence
Hybrid evolutionary algorithm
Fuzzy C-means clustering
Non-convex economic dispatch
Valve-point loading effect
Prohibited operating zones
Unit commitment

AMS Classification 2010:
90C59; 90C11; 65K05

ABSTRACT

Practical thermal economic dispatch (ED) problems with valve-point effects (VPE), prohibited operating zones (POZ), ramp-rate limits, and unit commitment (UC) decisions are non-smooth, non-convex, and fragmented, which defeats gradient-based solvers. This paper presents a feasibility-aware computational intelligence framework—the Fuzzy C-Means Population-Based Genetic Algorithm (FCM-PbGA)—that embeds commitment-consistent encoding, FCM-guided initialisation, ramp-feasible bounds, POZ-aware projection, and Lamarckian refinement in a single-population evolutionary architecture. The framework is evaluated on the IEEE 15-unit, 24-hour benchmark (30 trials). The initial island-based architecture (Stage 1) achieves \$662,173 ± \$2,151. Budget-matched re-runs and ablation isolate architecture from budget: fixed six-island partitioning proves counterproductive, and the matched GA outperforms Stage 1. Three redesigns follow: adaptive merging, reheating, POZ-aware crossover and fractional seeding close 49% of the gap; corrected operators close an additional 36%; and removing islands entirely (v4) reaches \$655,073.62 ± \$1,469—a statistical tie with the matched GA ($p = 0.199$, $d = -0.34$) under published conditions, and a significant lead ($p = 5.2 \times 10^{-6}$, $d = -1.34$) with correctly loaded POZ sets and strict repair. Zero power-balance, ramp-rate, POZ, and generation-limit violations are verified over 10,800 unit-hours. The study demonstrates a diagnose-then-redesign methodology for constrained engineering scheduling.

CC BY

1. Introduction

The integration of computer science and artificial-intelligence techniques with electrical and power engineering has become one of the most productive directions for tackling optimisation problems that classical numerical methods cannot solve reliably. Modern power systems operation is a natural proving ground for this integration: scheduling problems are large-scale, discontinuous, and constrained by discrete commitment decisions that couple every time period, placing them squarely beyond the reach of gradient-based solvers and motivating the design

*Corresponding Author

of purpose-built computational-intelligence architectures rather than the direct reuse of generic, off-the-shelf metaheuristics. This paper develops one such architecture—a hybrid combination of fuzzy clustering and evolutionary computation—and uses non-convex economic dispatch as the representative electrical power engineering problem on which its computational design choices are evaluated.

Economic dispatch (ED) is the core operational planning task in thermal power systems: allocate real-power output among committed generators so that total fuel cost is minimised subject to physical and operational constraints.[1,2] With increasing demand variability and fuel-cost uncertainty,[3] and the proliferation of distributed energy resources in modern smart-grid and microgrid environments,[4] inadequate dispatch directly translates into excess expenditure, avoidable emissions, and reduced grid reliability.[5,6]

In practical thermal systems, ED departs fundamentally from the smooth convex model assumed by classical lambda-iteration and gradient-based methods. Three sources of non-convexity are particularly prevalent:[7,8] (i) *Valve-point loading effects (VPE)*: throttle openings superimpose sinusoidal ripples on the quadratic cost curve, creating multiple local minima;[9,10] (ii) *Prohibited operating zones (POZ)*: mechanical resonance and thermal-stability limits forbid certain output intervals, fragmenting the feasible set for each unit;[11,12] recent work confirms POZ remains an active research front, from exact bisection methods[13] and MILP reformulations[14] to hybrid evolutionary approaches under uncertainty;[15] (iii) *Ramp-rate constraints and unit commitment (UC)*: the inter-temporal coupling of output changes and binary on/off decisions[16] makes the problem a time-coupled mixed-integer nonlinear programme.[17,18]

Metaheuristic methods dominate the recent literature because they require no smoothness assumptions. Representative contributions include PSO-based approaches,[19] Grey Wolf Optimiser (GWO), Whale Optimisation Algorithm (WOA), Slime Mould Algorithm,[20] Social Small Group Optimisation,[21] quasi-oppositional political optimiser,[22] reinforcement-learning dispatch for smart grids,[4] hybrid MILP strategies,[14] hybrid heuristics for economic load dispatch,[23] and multi-objective emission–cost formulations.[24] Despite substantial progress, a recurring limitation across these studies is that *feasibility is treated as an external concern*: infeasible solutions are penalised or repaired after the search step without altering the search geometry. Penalty tuning is problem-specific and fragile; post-hoc repair can mask true convergence behaviour by decoupling the trajectory from the feasibility structure. A second gap is the integrated treatment of UC[25] alongside VPE, POZ, and ramp constraints in a unified framework.[26,27]

Knowledge-guided initialisation addresses a well-known weakness of random population seeding in non-convex spaces: the population spends many generations recovering from infeasibility rather than exploring promising regions.[28,29] By clustering historical dispatch vectors with FCM and seeding the population near the centroid closest to the previous solution, the search is directed toward physically and economically reasonable regions from generation 1, thereby accelerating early convergence.

Island-based architectures preserve population diversity by maintaining semi-independent sub-populations that exchange information periodically through controlled migration.[30,31] In fragmented feasible spaces such as POZ-partitioned dispatch landscapes, premature convergence is a primary failure mode; the island structure counters this by allowing different sub-populations to explore different feasibility segments simultaneously.

As a computational-technique contribution, the specific innovations of FCM-PbGA relative to prior work are:

(1) Structural feasibility embedding (vs. penalty-based repair): whereas existing methods apply POZ/ramp correction after optimisation, FCM-PbGA confines the search to the feasible set Ω_i^t throughout every variation operator, eliminating the need for large penalty coefficients.
(2) FCM-assisted warm start: to the authors' knowledge, fuzzy-clustering of inter-temporal dispatch history for GA initialisation has not been applied in the non-convex ED literature; prior FCM applications are limited to load classification,[29] load profiling,[32] forecasting,[28] and fuzzy-clustering-based asset diagnostics in power systems.[33]
(3) Lamarckian local-search refinement: a gradient-free hill-climb stage refines the global best after island evolution, providing a final exploitation step not present in vanilla island GAs applied to ED.
(4) Comprehensive statistical validation: 30 independent trials with one-way ANOVA, Friedman test, Tukey HSD, two-sample t-test, and Wilcoxon rank-sum test—a rigour level uncommon in the ED metaheuristic literature.

(5) A reusable computational architecture: the four components above are implemented as a single, modular evolutionary pipeline—FCM-based initialisation, island-partitioned search, constraint-projection repair, and Lamarckian refinement—that is not specific to the cost function or constraint set of economic dispatch, and is presented so that it can be re-instantiated for other discrete-continuous, constraint-fragmented optimisation problems in electrical and computer engineering.

Fuzzy clustering appears in the optimisation literature in two roles distinct from ours. First, FCM itself is the optimisation target solved by metaheuristics, with GA/PSO hybrids tuning FCM centres to escape local minima[34,35], and with PSO- or kernel-enhanced FCM variants[36,37], surveyed for automatic clustering in.[38] Second, FCM assists power-system tasks such as load classification, daily-curve profiling and short-term load forecasting, commonly coupled with PSO-SVR or kernel fuzzy clustering[39,40], and with deep learning or locally weighted regression.[41,42] To the authors' knowledge, fuzzy clustering of inter-temporal dispatch history for GA seeding in non-convex ED with valve-point effects, POZ, ramp-rate and UC decisions has not been reported; the works above are cited as contrast, not precedent.

The current study is limited to a single 15-unit benchmark with a synthetic 24-hour load profile and simplified network modelling (no line constraints or reserve requirements). Minimum up/down times and startup/shutdown cost dynamics are not modelled. Extension to large-scale systems, renewable integration, network and reserve constraints, and rolling-horizon operation is identified as future work.

The remainder of this paper is organised as follows. Section 2 states the problem. Section 3 describes the Stage 1 framework. Section 4 evaluates Stage 1 against existing practice and benchmarks, diagnoses the budget confound, derives the v2–v4 redesign, and compares all stages. Section 5 concludes with limitations and future directions.

2. Problem formulation

2.1. Objective function

The objective of economic dispatch is to minimize the total fuel cost of all generating units[43,44] while meeting system demand.[45,46] For thermal generators, the fuel cost is typically represented by a quadratic function.[47] However, practical steam turbine behavior introduces valve-point loading effects, which cause non-smooth and non-convex characteristics in the cost curve.

The resulting cost function is formulated as (**Equation (1)**)

$$\min\ F = \sum_{i=1}^{N} \Big(a_i + b_i P_i + c_i P_i^2 + \left| d_i \sin\left(e_i (P_i^{\min} - P_i)\right) \right| \Big), \tag{1}$$

where P_i denotes the real power output of unit i, N is the total number of generating units in the system, (a_i, b_i, c_i) are the quadratic fuel-cost coefficients, and (d_i, e_i) characterize the amplitude and frequency of the valve-point effect, respectively. The sinusoidal term superimposed on the quadratic cost introduces ripples that generate multiple local minima, thereby producing a highly non-convex optimization landscape.

Equation (1) states the classic all-online static ED in which every unit is committed, following the standard textbook convention;[48] no commitment variable appears in it by construction.

2.2. Operational constraints

To ensure physically realizable and secure operation, ED is subject to operational constraints representing unit limits and system requirements.[24] These constraints define the feasible operating region.

2.2.1. Power balance

At each scheduling period, total generated power must satisfy the system demand while accounting for transmission losses. This requirement is expressed as (**Equation (2)**)

$$\sum_{i=1}^{N} P_i = P_D + P_L, \tag{2}$$

where P_i denotes the output of unit i, P_D is the total system demand, and P_L represents transmission losses. This equality constraint ensures real-time supply–demand equilibrium and constitutes the fundamental feasibility condition of the economic dispatch problem.

2.2.2. Generation capacity limits

Each generating unit is constrained by its technical minimum and maximum operating limits, reflecting thermal, mechanical, and safety considerations. These limits are expressed as (**Equation (3)**)

$$P_i^{\min} \le P_i \le P_i^{\max}, \quad i = 1, 2, \ldots, N, \tag{3}$$

where P_i denotes the output of unit i, and $P_i^{\min}$ and $P_i^{\max}$ represent its minimum and maximum allowable generation levels, respectively. These

bounds define the feasible operating interval of each generator and ensure physically realizable dispatch solutions.

2.2.3. Prohibited operating zones

In practical thermal units, certain output ranges must be avoided due to mechanical vibration, thermal stress, or stability limitations.[4] These restricted intervals, known as prohibited operating zones (POZ), further complicate the dispatch problem by fragmenting the feasible operating space. The constraint is formulated as (**Equation (4)**)

$$P_i \notin \left[P_{i,k}^L,\, P_{i,k}^U\right], \quad k = 1, 2, \ldots, K_i, \tag{4}$$

where $[P_{i,k}^L, P_{i,k}^U]$ denotes the k-th prohibited interval of unit i, and K_i is the total number of prohibited operating zones defined for unit i (so that k ranges over $1, 2, \ldots, K_i$). The exclusion of these intervals partitions the admissible generation range into multiple disjoint segments, thereby increasing non-convexity and creating additional local minima in the economic dispatch landscape.

2.2.4. Ramp rate

Ramp-rate constraints limit the permissible change in generator output between consecutive scheduling periods in order to maintain operational stability and protect equipment from excessive thermal stress. The inter-temporal limits for $i = 1, 2, \ldots, N$ are expressed as (**Equation (5)**):

$$\begin{aligned} P_i^t - P_i^{t-1} &\le R_i^U, \\ P_i^{t-1} - P_i^t &\le R_i^D, \end{aligned} \tag{5}$$

where P_i^t and P_i^{t-1} denote the power output of unit i at periods t and $t-1$, respectively, while R_i^U and R_i^D represent the maximum allowable ramp-up and ramp-down rates. These constraints enforce dynamically feasible transitions across the scheduling horizon and couple dispatch decisions over time.

2.2.5. Unit commitment

Unit commitment (UC) determines the on/off operating status of generating units across the scheduling horizon.[16,27] By controlling which units are available for dispatch, UC ensures sufficient generation capacity to meet the demand while avoiding unnecessary operation and excessive fuel consumption.[26,49] The commitment state of unit i at period t is represented by the binary variable (**Equation (6)**)

$$u_i^t \in \{0, 1\}, \tag{6}$$

where $u_i^t = 1$ indicates that the unit is online and $u_i^t = 0$ indicates that it is offline. When committed, the generation output must satisfy (**Equation (7)**)

$$u_i^t P_i^{\min} \le P_i^t \le u_i^t P_i^{\max}, \tag{7}$$

where P_i^t is the output of unit i at period t and $P_i^{\min}$, $P_i^{\max}$ are its technical limits, so that offline units ($u_i^t = 0$) produce zero output. To maintain system adequacy, the total available committed capacity must satisfy (**Equation (8)**)

$$\sum_{i=1}^{N} u_i^t P_i^{\max} \ge P_D^t, \tag{8}$$

ensuring that the committed units collectively provide sufficient maximum output to cover the system demand P_D^t, i.e. the same demand quantity P_D introduced above but now indexed by scheduling period t to reflect its hour-to-hour variation. The simultaneous consideration of UC decisions together with valve-point effects, generation limits, prohibited operating zones, and ramp dynamics results in a strongly non-convex and inter-temporally coupled optimization problem that necessitates feasibility-aware solution strategies.

3. Proposed method

The overall per-period flow of FCM-PbGA-v4 is summarised in **Figure 1** and Algorithm 1.

3.1. Feasibility-aware commitment representation

At each period t, the system state is represented by the dispatch vector $\mathbf{P}^t$ and the commitment vector $\mathbf{u}^t$, given respectively by **Equations (9)** and **(10)**:

$$\mathbf{P}^t = \begin{bmatrix} P_1^t & P_2^t & \cdots & P_N^t \end{bmatrix}^T, \tag{9}$$

$$\mathbf{u}^t = \begin{bmatrix} u_1^t & u_2^t & \cdots & u_N^t \end{bmatrix}^T, \quad u_i^t \in \{0, 1\}, \tag{10}$$

where P_i^t is the output of unit i and u_i^t denotes its on/off status. To maintain consistency between commitment decisions and evolutionary search, only online units are encoded. The set of committed units is given by **Equation (11)**, where n_{on} denotes the number of units that are online (committed) at period t

$$\mathcal{J}^t = \{i \mid u_i^t = 1\}, \quad n_{\text{on}} = |\mathcal{J}^t|, \tag{11}$$

and the chromosome is constructed as in **Equation (12)**

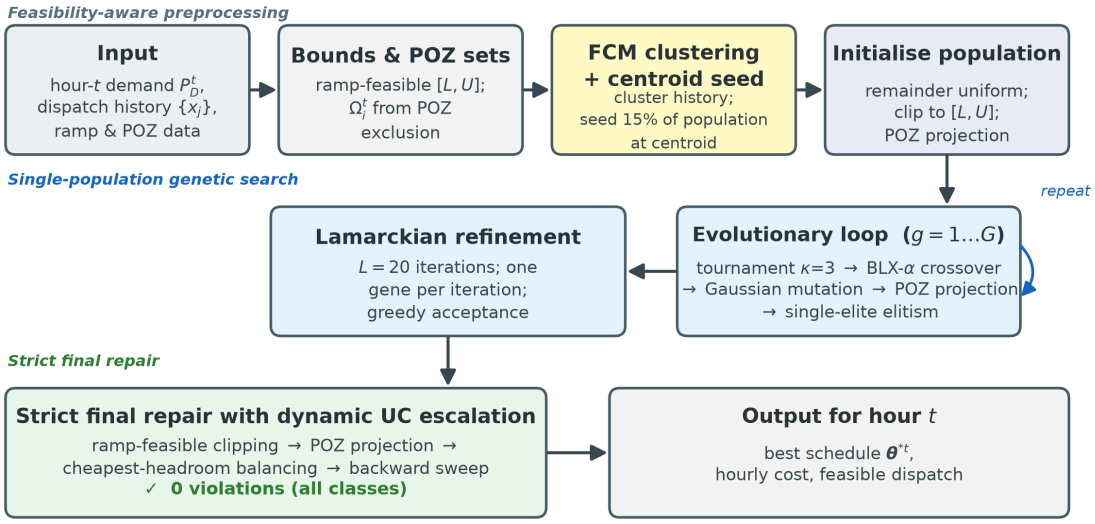

Figure 1. Overall flow of FCM-PbGA-v4 for one scheduling hour t: feasibility-aware preprocessing (interval bounds, POZ sets, FCM clustering and centroid seeding), single-population genetic search with POZ-aware operators and Lamarckian refinement, and strict final repair with dynamic UC escalation; the cycle repeats over the 24 hourly periods

Algorithm 1. FCM-PbGA-v4 for one hour t (final variant; single population throughout).

Require: demand P_D^t, dispatch history $\{\mathbf{x}_j\}$, v4 parameters detailed in Section 4.1

1: Build $[L_i^t, U_i^t]$ via (14)–(15); form Ω_i^t via (16)
2: Cluster history with FCM and seed $\mathbf{P}_{\text{seed}}^t$ via (25)
3: Initialise population: $\lceil 0.15 N_{\text{pop}} \rceil$ Gaussian samples around the seed ($\beta = 0.05$), remainder uniform; clip to $[L, U]$, apply POZ projection
4: **for** $g = 1$ to G **do**
5: Tournament selection ($\kappa = 3$); BLX-α crossover ($p_{\text{cx}} = 0.85$); Gaussian mutation ($\gamma = 0.20$, $p_{\text{mut}} = 0.30$, $p_{\text{gene}} = 0.20$, $\rho = 1.0$); POZ projection of each gene
6: Single-elite elitism; record best-so-far
7: **end for**
8: Lamarckian refinement: $L = 20$ iterations, one random gene perturbed per iteration, greedy acceptance
9: Strict final repair with dynamic UC escalation ((50)–(52)); log pre/post violations

Ensure: best chromosome $\boldsymbol{\theta}^{*t}$, hourly cost, feasible schedule

$$\boldsymbol{\theta}^t = \begin{bmatrix} P_{i_1}^t & P_{i_2}^t & \cdots & P_{i_{n_{\text{on}}}}^t \end{bmatrix}^T, \quad i_k \in \mathcal{J}^t, \quad (12)$$

containing only the outputs of committed units. Offline units are fixed at zero and excluded from the search space. This encoding reduces dimensionality, prevents infeasible exploration of inactive generators, and improves numerical stability in the non-convex dispatch problem.

3.2. Commitment-aware objective function for non-convex power systems

Restating **Equation (1)** with commitment gating (u_i^t) and period index (t) gives the UC/ED objective minimised by the optimiser, where each summand of **Equation (1)** is multiplied by u_i^t

and $P_i \to P_i^t$: (**Equation (13)**)

$$F(\mathbf{P}^t, \mathbf{u}^t) = \sum_{i=1}^{N} u_i^t \Big(a_i + b_i P_i^t + c_i (P_i^t)^2 + \left| d_i \sin\left(e_i (P_i^{\min} - P_i^t) \right) \right| \Big), \tag{13}$$

where u_i^t ensures that only committed units contribute to cost, while offline units generate no power and incur no expense. This formulation directly couples unit commitment and dispatch, aligning chromosome encoding with objective evaluation and preserving the true non-convex cost structure induced by valve-point effects and commitment decisions.

3.3. Feasible operating region under ramp dynamics and prohibited operating zones

At period t, the feasible operating region is determined by ramp-rate limits and prohibited operating zones (POZ), leading to time-varying admissible intervals. Ramp constraints couple consecutive periods, so the effective bounds depend on the previous output. During an OFF→ON transition ($u_i^{t-1} = 0$, $u_i^t = 1$), hereafter referred to as startup, the unit increases output within its ramp-up capability rather than starting directly at $P_i^{\min}$. The effective lower bound is (**Equation (14)**)

$$L_i^t = \begin{cases} P_i^{\min}, & \text{if startup,} \\ \max\left(P_i^{\min}, P_i^{t-1} - R_i^D \right), & \text{otherwise,} \end{cases} \tag{14}$$

and the upper bound is (**Equation (15)**)

$$U_i^t = \begin{cases} \min\left(P_i^{\max}, P_i^{\min} + R_i^U \right), & \text{if startup,} \\ \min\left(P_i^{\max}, P_i^{t-1} + R_i^U \right), & \text{otherwise,} \end{cases} \tag{15}$$

so that $[L_i^t, U_i^t]$ defines the ramp-feasible interval at time t.

Excluding the prohibited operating zones $Z_{i,k}$ yields the admissible set (**Equation (16)**)

$$\Omega_i^t = [L_i^t,\, U_i^t] \setminus \bigcup_{k=1}^{K_i} Z_{i,k}, \tag{16}$$

which may consist of multiple disjoint segments, where $Z_{i,k} = [P_{i,k}^L, P_{i,k}^U]$ is the k-th prohibited interval of unit i introduced above and K_i is again the number of such intervals for that unit. The interaction of ramp dynamics and POZ exclusion therefore produces a time-dependent and fragmented feasible region, increasing the difficulty of dispatch optimization.

3.4. POZ-aware projection and feasibility restoration

Stochastic crossover and mutation may generate solutions that violate ramp limits or fall within prohibited operating zones. To preserve feasibility, each candidate value is repaired by projecting it onto the admissible set Ω_i^t. The projection operator is defined as (**Equation (17)**)

$$\Pi_{\Omega_i^t}(x) = \arg\min_{y \in \Omega_i^t} |y - x|, \tag{17}$$

which maps an infeasible candidate value x to the nearest feasible point y within the ramp- and POZ-compliant region Ω_i^t (i.e. y ranges over the admissible set while x is the fixed, possibly infeasible, input being repaired). The repair is applied after variation as in **Equation (18)**

$$P_i^t \leftarrow \Pi_{\Omega_i^t}\left(P_i^t \right), \quad \forall i \in \mathcal{J}^t. \tag{18}$$

By enforcing feasibility before fitness evaluation, the search remains confined to dynamically admissible operating intervals. This projection-based restoration improves robustness and convergence stability in the fragmented non-convex feasible space induced by ramp and POZ constraints.

3.5. Fuzzy C-means assisted knowledge-guided initialization

To avoid excessive exploration of infeasible or economically poor regions, historical dispatch patterns are incorporated into the initialization stage using Fuzzy C-Means (FCM) clustering.[28,29] Let $\{\mathbf{x}_j\}_{j=1}^M$ denote historical dispatch samples, where M is the total number of such samples collected from previously solved periods (**Equation (19)**)

$$\mathbf{x}_j \in \mathbb{R}^N, \quad \mathbf{x}_j = \mathbf{P}^j, \tag{19}$$

where each $\mathbf{x}_j$ represents a past generation vector. FCM partitions the samples into C clusters by minimizing (**Equation (20)**)

$$\min_{C,U} J_{\mathrm{FCM}}(C, U) = \sum_{j=1}^{M} \sum_{k=1}^{C} u_{kj}^m \|\mathbf{x}_j - \mathbf{c}_k\|_2^2, \tag{20}$$

subject to (**Equation (21)**)

$$\sum_{k=1}^{C} u_{kj} = 1, \quad 0 \le u_{kj} \le 1, \tag{21}$$

where $\mathbf{c}_k$ is the k-th cluster center, u_{kj} the membership degree, and $m > 1$ the fuzzification parameter. The update rules for the cluster centers and the membership degrees are given respectively by **Equations (22)** and **(23)**

$$\mathbf{c}_k = \frac{\sum_{j=1}^{M} u_{kj}^m \mathbf{x}_j}{\sum_{j=1}^{M} u_{kj}^m}, \tag{22}$$

$$u_{kj} = \left(\sum_{l=1}^{C} \left(\frac{\|\mathbf{x}_j - \mathbf{c}_k\|_2}{\|\mathbf{x}_j - \mathbf{c}_l\|_2} \right)^{\frac{2}{m-1}} \right)^{-1}. \tag{23}$$

Centers and memberships are updated alternately via **Equations (22)** and **(23)** until the cluster centers stabilise or a maximum iteration count is reached (**Equation (24)**)

$$\left\| \mathbf{c}^{(\tau+1)} - \mathbf{c}^{(\tau)} \right\|_F < \epsilon_{\text{FCM}} \quad \text{or} \quad \tau \geq \tau_{\max}, \tag{24}$$

where τ indexes the FCM iteration, $\|\cdot\|_F$ is the Frobenius norm over the matrix of stacked cluster centers, ϵ_{FCM} is a small convergence tolerance, and $\tau_{\max}$ is the iteration cap; whichever condition is met first terminates the clustering.

At period t, the initial population is seeded by the cluster center closest to the previous dispatch $\mathbf{P}^{t-1}$ (**Equation (25)**):

$$k^* = \underset{k \in \{1,\ldots,C\}}{\arg\min} \|\mathbf{c}_k - \mathbf{P}^{t-1}\|_2, \quad \mathbf{P}^t_{\text{seed}} = \mathbf{c}_{k^*}. \tag{25}$$

where k^* is the index of the nearest cluster center and $\mathbf{P}^t_{\text{seed}}$ is the resulting seeded dispatch at period t. This knowledge-guided initialization provides a physically meaningful starting point and improves convergence stability in the non-convex dispatch landscape.

3.6. Population-based island genetic optimization framework

To reduce premature convergence in the fragmented non-convex dispatch space, a population-based island genetic architecture is employed.[30,50] The total population is divided into S semi-independent islands, ((**Equation 26**))

$$\mathcal{P} = \bigcup_{s=1}^{S} \mathcal{P}_s, \quad |\mathcal{P}_s| = N_s, \quad \sum_{s=1}^{S} N_s = N_{\text{pop}}, \tag{26}$$

where each island evolves in parallel, N_s is the number of individuals assigned to island s, and N_{pop} is the total population size across all S islands. At generation g, island s is updated as (**Equation (27)**)

$$\mathcal{P}_s^{g+1} = \varepsilon_s^g \cup \text{Var}\big(\text{Sel}(\mathcal{P}_s^g)\big), \tag{27}$$

where ε_s^g denotes elite preservation, $\text{Sel}(\cdot)$ denotes the selection operator, and $\text{Var}(\cdot)$ represents crossover and mutation.

3.6.1. Commitment-aware representation and selection

Recall from Section 3.1 that each individual encodes only the committed set $\mathcal{J}^t$ ($n_{\text{on}} = |\mathcal{J}^t|$ genes). The full dispatch vector is reconstructed as (**Equation (28)**), where $q \in \{1, \ldots, n_{\text{on}}\}$ indexes the position of unit i_q within the chromosome, so that gene q supplies the output of the corresponding committed unit i_q

$$P_i^t = \begin{cases} \theta_q^t, & \text{if } i = i_q \in \mathcal{J}^t, \\ 0, & \text{if } u_i^t = 0, \end{cases} \tag{28}$$

ensuring structural consistency between unit commitment and evolutionary search while reducing dimensionality.

Parent selection follows tournament competition[51,52] (**Equation (29)**):

$$\boldsymbol{\theta}^{\text{sel}} = \underset{\boldsymbol{\theta} \in \mathcal{T}}{\arg\min}\, F(\boldsymbol{\theta}), \quad |\mathcal{T}| = \kappa, \tag{29}$$

providing controlled selective pressure while preserving intra-island diversity, where $\mathcal{T}$ is a randomly sampled subset of the island population (the tournament pool) and $\kappa = |\mathcal{T}|$ is the tournament size, i.e. the number of competing individuals drawn for each selection event.

3.6.2. Variation and adaptive mutation mechanisms

Continuous recombination is applied to parent chromosomes $\boldsymbol{\theta}^{(1)}$ and $\boldsymbol{\theta}^{(2)}$.[53,54] For each gene j, the absolute difference between the two parents' values is computed (**Equation (30)**)

$$\Delta_j = \left| \theta_j^{(1)} - \theta_j^{(2)} \right|, \tag{30}$$

giving Δ_j, the absolute spread between the parents' values for gene j. With probability p_{cx}, the offspring gene is then sampled uniformly within the parental interval expanded by a factor $\alpha > 0$ (the BLX-α expansion parameter), using BLX-α crossover (**Equation (31)**); otherwise (with probability $1 - p_{\text{cx}}$) the offspring gene is copied directly from the first parent, $\theta'_j \leftarrow \theta_j^{(1)}$, without recombination

$$\theta'_j \sim \mathcal{U}\Big(\min(\theta_j^{(1)}, \theta_j^{(2)}) - \alpha\Delta_j, \max(\theta_j^{(1)}, \theta_j^{(2)}) + \alpha\Delta_j \Big), \quad \text{with probability } p_{\text{cx}}, \tag{31}$$

allowing controlled extrapolation beyond parental bounds.

Mutation is applied through two nested probability gates: each offspring undergoes mutation with probability p_{mut}, and, conditional on that event, each individual gene j is independently selected for perturbation with probability p_{gene}; only genes selected at both levels receive a Gaussian perturbation[55] (**Equation (32)**):

$$\begin{aligned} &\theta_j \leftarrow \theta_j + \eta_j, \\ &\eta_j \sim \mathcal{N}(0, \sigma_j^2), \\ &\quad \text{with probability } p_{\text{mut}} \cdot p_{\text{gene}} \text{ per gene,} \end{aligned} \tag{32}$$

where η_j is a zero-mean Gaussian perturbation applied to gene j, σ_j is that gene's current mutation standard deviation (defined below), and genes not selected at either gate are left unchanged; the result is then clipped to physical limits (**Equation (33)**):

$$\theta_j \leftarrow \min\Big(\max(\theta_j, P_{i_j}^{\min}), P_{i_j}^{\max}\Big). \tag{33}$$

The initial mutation scale is defined as (**Equation (34)**), where $\gamma \in (0,1)$ is a fixed scaling factor that sets the initial mutation spread as a fraction of each unit's feasible generation range

$$\sigma_j^{(0)} = \gamma\left(P_{i_j}^{\max} - P_{i_j}^{\min}\right), \quad \gamma \in (0,1), \tag{34}$$

and is adaptively decayed across generations according to (**Equation (35)**), where $\rho \in (0,1)$ is the per-generation decay rate and $\sigma_j^{\min}$ is a floor value that prevents the mutation scale from collapsing to zero

$$\sigma_j^{(g+1)} = \max\left(\rho\sigma_j^{(g)}, \sigma_j^{\min}\right), \quad 0 < \rho < 1, \tag{35}$$

ensuring balanced local exploration across heterogeneous units while gradually increasing exploitation as the search progresses.

3.6.3. Ring migration and global best selection

To promote information exchange among islands, every G_{mig} generations the best K_{mig} individuals from island s are selected for migration[56] (**Equation (36)**):

$$\mathcal{M}_s^g = \text{Best}_{K_{\text{mig}}}\big(\mathcal{P}_s^g\big). \tag{36}$$

where $\mathcal{M}_s^g$ is the migrant set and $\mathcal{P}_s^g$ the population of island s at generation g. Using a ring topology, these migrants replace the worst members of the neighboring island $s+1$ (**Equation (37)**):

$$\mathcal{P}_{s+1}^g \leftarrow \big(\mathcal{P}_{s+1}^g \setminus \text{Worst}_{K_{\text{mig}}}\big(\mathcal{P}_{s+1}^g\big)\big) \cup \mathcal{M}_s^g. \tag{37}$$

After G global generations, the best dispatch at period t is selected as (**Equation (38)**)

$$\boldsymbol{\theta}^{*t} = \underset{\boldsymbol{\theta} \in \bigcup_{s=1}^{S} \mathcal{P}_s^G}{\arg\min}\ F(\boldsymbol{\theta}). \tag{38}$$

where $\boldsymbol{\theta}^{*t}$ is the best chromosome at period t and $F(\cdot)$ its commitment-aware cost. This distributed scheme balances exploration and exploitation, improving robustness in ramp-coupled and POZ-fragmented non-convex dispatch problems.[31,57]

3.7. Lamarckian local search

To further refine the best solution obtained from the global evolutionary search, a Lamarckian local search is applied to the global best chromosome $\boldsymbol{\theta}^{(0)} = \boldsymbol{\theta}^{*t}$ from **Equation (38)**. The refinement proceeds coordinate-wise over L iterations: at each iteration $\ell = 1, \ldots, L$, a single gene index j_ℓ is drawn uniformly at random from $\{1, \ldots, n_{\text{on}}\}$ and only that coordinate of the current incumbent $\boldsymbol{\theta}^{(\ell-1)}$ is perturbed, with all other genes left unchanged (**Equation (39)**)

$$\begin{aligned} \theta'_j &= \begin{cases} \theta_j^{(\ell-1)} + \delta_\ell, & j = j_\ell, \\ \theta_j^{(\ell-1)}, & j \neq j_\ell, \end{cases} \\ \delta_\ell &\sim \mathcal{N}(0, \sigma_{\text{ls}}^2), \end{aligned} \tag{39}$$

where σ_{ls} is the local-search perturbation scale for the perturbed unit, fixed at 3% of that unit's feasible generation range. The perturbed chromosome $\boldsymbol{\theta}'$ is repaired, clipped to its feasible generation limits, and evaluated using the same objective function. Because only one gene changes per iteration, the perturbation is accepted or rejected greedily before the next iteration begins, so that each step refines the incumbent found so far rather than an independently perturbed copy of the original chromosome (**Equation (40)**)

$$\boldsymbol{\theta}^{(\ell)} \leftarrow \begin{cases} \boldsymbol{\theta}', & \text{if } F(\boldsymbol{\theta}') < F(\boldsymbol{\theta}^{(\ell-1)}), \\ \boldsymbol{\theta}^{(\ell-1)}, & \text{otherwise,} \end{cases} \tag{40}$$

and the chromosome obtained after all L iterations, $\boldsymbol{\theta}^{(L)}$, replaces the inherited global best. Because the improved chromosome directly replaces the inherited solution, the refinement is Lamarckian in nature. This coordinate-wise, greedy hill-climbing procedure trades the breadth of a full-chromosome perturbation for a higher acceptance rate per iteration, strengthening exploitation around the best feasible dispatch discovered by the island-based evolutionary process while keeping each individual perturbation small and computationally cheap to evaluate.

3.8. Penalized fitness for demand adequacy

While projection enforces ramp and POZ feasibility, supply–demand balance is handled through a penalty-based fitness formulation. The demand shortfall at period t is defined as (**Equation (41)**)

$$\delta^t = \max\left(0, P_D^t - \sum_{i=1}^{N} P_i^t\right), \tag{41}$$

where δ^t is the unserved demand at period t and P_D^t the corresponding system demand, and penalized quadratically as (**Equation (42)**)

$$\phi_{\text{bal}}^t = \lambda(\delta^t)^2, \tag{42}$$

where $\lambda > 0$ determines the severity of adequacy violations.

If residual prohibited operating zone (POZ) violations occur, an additional penalty is imposed (**Equation (43)**):

$$\phi_{\text{poz}}^t = \sum_{i\in\mathcal{V}^t} \Big(\beta_0 + \beta_1 \min\big\{|P_i^t - P_{i,k}^L|, |P_i^t - P_{i,k}^U|\big\}\Big), \tag{43}$$

where $\mathcal{V}^t$ denotes violating units and $\beta_0, \beta_1 > 0$ control penalty magnitude.

The final fitness is (**Equation (44)**)

$$J(\mathbf{P}^t, \mathbf{u}^t) = F(\mathbf{P}^t, \mathbf{u}^t) + \phi_{\text{bal}}^t + \phi_{\text{poz}}^t, \tag{44}$$

combining economic cost with adequacy and residual feasibility terms. This formulation directs the search toward cost-efficient and operationally admissible solutions while discouraging infeasible dispatch in the non-convex space.

3.9. Ramp-aware post-optimization enforcement

Because optimization is performed period-wise, the obtained P_i^t is post-processed to enforce intertemporal feasibility. Each output is first adjusted to satisfy (14)–(15) and then projected onto Ω_i^t via (17). If total generation still falls short of demand after this adjustment, a demand-restoration step redistributes the shortfall across units in proportion to their remaining headroom, iterating until the shortfall is closed or no headroom remains (**Equations (45)** and **(46)**)

$$h_i^t = \max\big(0,\, U_i^t - P_i^t\big), \tag{45}$$

$$P_i^t \leftarrow P_i^t + \left(D^t - \sum_{l=1}^{N} P_l^t\right) \frac{h_i^t}{\sum_{l=1}^{N} h_l^t}, \quad \forall i \in \mathcal{J}^t, \tag{46}$$

where h_i^t is the spare headroom of unit i at period t (the gap between its current output and its ramp-feasible upper bound U_i^t from **Equation (15)**), and each update of **Equation (46)** is followed by a POZ re-projection via **Equation (17)** and re-clipping to $[L_i^t, U_i^t]$ before the headroom is recomputed for the next iteration. This prevents accumulation of ramp/startup infeasibilities across the horizon while keeping all units within Ω_i^t.

3.10. POZ-aware unit commitment escalation

Under ramp-rate and prohibited operating zone (POZ) constraints, the nominal committed capacity $\sum u_i^t P_i^{\max}$ may overestimate the actually deliverable power. Since the feasible set Ω_i^t is time-dependent and potentially fragmented, the effective maximum output of each committed unit is defined as (**Equation (47)**)

$$U_i^{\text{eff},t} = \max_{x\in\Omega_i^t} x, \tag{47}$$

representing the highest ramp- and POZ-feasible generation at period t. Adequacy is then checked by (**Equation (48)**)

$$\sum_{i:u_i^t=1} U_i^{\text{eff},t} < P_D^t - \varepsilon, \tag{48}$$

where ε is a small tolerance. If this condition holds, the committed units cannot satisfy demand within the admissible region, and additional units are activated before re-optimization.

This escalation step maintains consistency between commitment decisions and the dynamically restricted feasible operating space, preventing infeasible dispatch under ramp and POZ constraints.

Factorial evaluation metric. The analysis adds a 2×2 factorial (FCM seeding on/off $\times$ POZ-aware crossover on/off) with arms A (both on), B (FCM off), C (POZ-aware off), and D (both off), where $\bar{C}_{(\cdot)}$ is the mean 24-hour cost of each arm. Super-additivity is measured by

$$I = (\bar{C}_A - \bar{C}_D) - [(\bar{C}_B - \bar{C}_D) + (\bar{C}_C - \bar{C}_D)], \tag{49}$$

with $I < 0$ indicating synergy beyond summed main effects. Statistics otherwise follow the protocol of Section 4.7.

3.10.1. Strict final repair

The final schedule is produced by a strict horizon pass over the optimiser's raw trajectory; the search itself and all seeds are unchanged, so raw trajectories are identical and only the raw-to-final

mapping is strict. Let $\mathbf{P}^{t-1}$ be the previous hour's finalised output with commitment $\mathbf{u}^{t-1}$. First, the ramp-feasible interval $[L_i^t, U_i^t]$ follows (14)–(15) with the corrected startup bound, and Ω_i^t follows (16). Second, adequacy is enforced dynamically: with $U_i^{\mathrm{eff},t}(\mathbf{P}^{t-1})$ denoting (47) evaluated on the previous-hour-dependent Ω_i^t,

$$\begin{aligned}\sum_{i:u_i^t=1} U_i^{\mathrm{eff},t}(\mathbf{P}^{t-1}) < P_D^t - \varepsilon \\ \Rightarrow \text{commit cheapest offline unit} \\ \text{and repeat,}\end{aligned} \tag{50}$$

where cheapest uses the merit order mc_i of (53) and $\varepsilon > 0$ is a small tolerance; u_i^t is the commitment state and P_D^t the hourly demand. Third, the raw output is clipped into $[L_i^t, U_i^t]$, projected via (17), and any shortfall is closed cheapest-headroom-first with projection after every increment,

$$\begin{aligned}P_i^t &\leftarrow \Pi_{\Omega_i^t}\big(P_i^t + \Delta_i\big), \\ \Delta_i &\geq 0, \ \sum\nolimits_i P_i^t \geq P_D^t,\end{aligned} \tag{51}$$

where $\Pi_{\Omega_i^t}$ is the projection operator of (17) and P_i^t never decreases, so balance shortfall cannot reappear. Fourth, ramp-down is enforced by backward sweeps (at most 10 rounds with forward re-verification),

$$\begin{aligned}P_i^{t+1} \leftarrow \min\Big(U_i^{t+1}, \max\big(L_i^{t+1}, \\ \Pi_{\Omega_i^{t+1}}(P_i^t - R_i^D)\big)\Big) \\ \text{whenever } P_i^t - P_i^{t+1} > R_i^D,\end{aligned} \tag{52}$$

with R_i^D the ramp-down limit. Offline units are fixed at zero and exempt from ramp checks (shut-down ramps are not modelled, consistent with the paper's UC scope).

These mechanisms define the baseline architecture, hereafter FCM-PbGA (*Stage 1*). The controlled diagnostics in Section 4.10 reveal structural weaknesses in Stage 1 that motivate three redesign iterations (FCM-PbGA-v2 to v4); their design is described in Section 4.11 and their evaluation in Section 4.11.5.

4. Results and discussion

This section presents a comprehensive evaluation of the proposed Fuzzy C-Means Population-Based Genetic Algorithm (FCM-PbGA) applied to the 24-hour Unit Commitment and Economic Dispatch (UC/ED) problem with 15 generating units. The evaluation is anchored against the existing scheduling practice—the Merit Order dispatch—which represents the conventional deterministic baseline currently employed in many power systems. Concretely, the Merit Order ranks all committed units by an approximate marginal cost evaluated at the midpoint of each unit's feasible range (**Equation (53)**)

$$mc_i = b_i + 2c_i\bar{P}_i, \quad \bar{P}_i = \frac{P_i^{\min} + P_i^{\max}}{2}, \tag{53}$$

where b_i and c_i are the same linear and quadratic fuel-cost coefficients used in **Equation (1)** and $\bar{P}_i$ is the midpoint output used as a fixed proxy for each unit's marginal cost. Committed units are then loaded strictly in ascending order of mc_i, each unit taking as much of the remaining demand as its capacity allows (**Equation (54)**)

$$\begin{aligned}P_{i_{(k)}} = \mathrm{clip}\Big(P_i^{\min}, P_i^{\max}, D - \textstyle\sum_{l<k} P_{i_{(l)}}\Big), \\ k = 1, 2, \ldots, n_{\mathrm{on}},\end{aligned} \tag{54}$$

where $i_{(1)}, i_{(2)}, \ldots$ denotes the unit ordering induced by sorting mc_i ascending, $\mathrm{clip}(\cdot,\cdot,\cdot)$ saturates its third argument to the interval given by the first two, and loading stops once the running total reaches demand D. Units are committed and dispatched strictly by this incremental production cost rank without any iterative optimisation, making it inherently incapable of navigating the non-convex cost landscape introduced by valve-point effects and prohibited operating zones. The primary objective of this evaluation is therefore to demonstrate the extent to which FCM-PbGA improves upon existing practice, and to validate its superiority over four competing state-of-the-art metaheuristics—Genetic Algorithm (GA), Particle Swarm Optimisation (PSO), Grey Wolf Optimiser (GWO), and Whale Optimisation Algorithm (WOA)—over 30 independent runs.

The discussion is structured as follows: Section 4.2 examines the UC schedule produced by FCM-PbGA and its structural advantages over Merit Order commitment decisions; Section 4.3 analyses the power dispatch and full cost decomposition; Section 4.4 verifies that FCM-PbGA satisfies all operational constraints that the existing Merit Order practice leaves unenforced, and additionally compares raw (pre-repair) violation counts across all five methods; Section 4.5 examines the convergence behaviour of FCM-PbGA relative to benchmark algorithms and its computational tractability (Section 4.5.1); Section 4.6 quantifies the improvement FCM-PbGA delivers over both the existing baseline and competing methods; Section 4.7 provides rigorous statistical confirmation of FCM-PbGA's superiority, including the distribution of costs across all 30 independent runs; Section 4.8 demonstrates the robust-

ness of FCM-PbGA under parameter variation; and Section 4.10 isolates architecture from evaluation budget via a component-wise ablation (Section 4.10.1) and a budget-matched benchmark rerun (Section 4.10.2). The diagnosis then drives the redesign in Section 4.11 (FCM-PbGA-v2 to v4), evaluated head-to-head in Section 4.11.5, before Section 4.12 discusses the operational significance of constraint handling for power system dispatch in light of these findings. The evaluation thus proceeds in three phases: (i) Stage 1 versus practice and benchmarks, (ii) diagnosis of the budget confound, and (iii) redesign and final comparison.

4.1. Experimental configuration

Table 1 reports the parameter settings used to produce the main results in Sections 4.2–4.7. The proposed FCM-PbGA is run with an enlarged, feature-complete configuration (larger population, more islands, migration, adaptive mutation decay, and Lamarckian local search), whereas GA, PSO, GWO, and WOA are run with compact, single-population settings representative of standard practice in the ED metaheuristic literature, with no island structure, migration, or local-search refinement. This asymmetry means that FCM-PbGA is also allocated a substantially larger objective-function evaluation budget than each vanilla benchmark: counting every candidate solution evaluated during initialisation, generational offspring evaluation, and local search, FCM-PbGA performs approximately 53,000 evaluations per hourly solve, versus approximately 4,950 for vanilla GA and 5,050 for each of PSO, GWO, and WOA—roughly a 10–11× larger budget (**Table 1**, bottom row). Consequently, the comparisons in Sections 4.2–4.7 should be read as an enhanced, full-featured configuration of the Stage-1 FCM-PbGA against compact, standard-practice baselines, *not* as a budget-matched ablation that isolates the marginal contribution of FCM-guided initialisation, island partitioning, or local search independently of the larger evaluation budget; the two are confounded in these sections. Section 4.10 resolves this confound directly: GA, PSO, GWO, and WOA are re-run with the same ≈53,000-evaluation budget as FCM-PbGA across the full 30-run, 24-hour protocol, alongside a component-wise ablation of FCM-PbGA itself, and the diagnostic findings motivate a redesigned variant, FCM-PbGA-v2 (Section 4.11), which is evaluated against all of the above in a single comprehensive comparison (Section 4.11.5). Section 4.8 subsequently varies four of the FCM-PbGA settings above—population size, number of islands, number of generations, and number of FCM clusters—one at a time around this baseline to assess robustness to this particular choice of settings.

4.2. FCM-PbGA unit commitment schedule: Departing from merit order

The binary unit commitment schedule is a fundamental output of the proposed FCM-PbGA that distinguishes it from the existing Merit Order practice in both structure and economic rationale. Under the Merit Order, units are committed in a fixed priority sequence regardless of the precise demand level, which frequently results in the unnecessary commitment of high-cost units during low-demand periods and fails to account for the non-linear interactions between unit outputs and valve-point operating costs. FCM-PbGA, by contrast, treats the commitment decision as an integral part of the optimisation objective, co-optimising binary commitment and continuous dispatch simultaneously through its hybrid binary-encoded GA structure guided by FCM-based island partitioning.

The optimised binary UC schedule produced by FCM-PbGA is presented in **Table 2**. During off-peak hours (Hours 1–4, system load 513–741 MW), FCM-PbGA commits only Units 5 and 7—the two lowest-cost units at low-load levels—while decommitting the remaining 13 units. Starting from Hour 5 (load 1,026 MW), FCM-PbGA activates the full 15-unit fleet and sustains this commitment through Hour 24.

The selective two-unit commitment during Hours 1–4 is a direct result of FCM-PbGA's co-optimisation mechanism. By avoiding the no-load and start-up costs of 13 units during low-demand hours, the proposed method eliminates a category of cost that the Merit Order incurs by default. This targeted commitment is made possible by the binary chromosome structure of FCM-PbGA, which encodes the on/off status of all 15 units directly in the genetic representation and evolves commitment decisions alongside dispatch setpoints within the same optimisation loop. The FCM-based island model further ensures that multiple distinct commitment configurations are explored simultaneously across subpopulations, preventing the algorithm from fixating prematurely on a single commitment pattern.

Table 1. Algorithm parameter settings used to produce the results in this paper. Stage-1 FCM-PbGA parameters are those referenced symbolically in Section 3; vanilla GA, PSO, GWO, and WOA are the compact single-population configuration common to all four benchmarks (Sections 4.2–4.8); budget-matched GA, PSO, GWO, and WOA retain the same vanilla operators but with Stage-1's population size and generation count (Section 4.10.2); FCM-PbGA-v2, v3, and v4 each list only the parameters that differ from the previous iteration (Section 4.11), all other settings being inherited unchanged. The evaluation-budget row reports the approximate total number of objective-function evaluations performed per hourly solve, computed directly from the population, generation, and local-search settings above.

Component	**Parameter**	**Value**
Stage-1 FCM-PbGA	Population size (N_{pop})	200
	Generations (G)	300
	Islands (S)	6
	Migration interval (G_{mig})	every 25 generations
	Migrants per event (K_{mig})	3
	Tournament size (κ)	3
	Elite individuals per island	4
	Crossover probability (p_{cx})	0.85
	Mutation probability (p_{mut})	0.30
	Per-gene mutation probability (p_{gene})	0.15
	Initial mutation scale factor (γ)	0.12
	Mutation decay rate (ρ)	0.97
	Local-search iterations (L)	20
FCM clustering (all stages)	Clusters (C)	5
	Minimum history before seeding	8 samples
	Fuzzification parameter (m)	2
	Convergence tolerance / iteration cap (ϵ_{FCM}, $\tau_{\max}$)	0.005 / 500
Vanilla GA, PSO, GWO, WOA (benchmarks)	Population size	50
	Generations	100
	Islands	1 (none)
	Local search	disabled
	Mutation decay rate ρ (GA only)	1.0 (no decay)
	PSO inertia / cognitive / social (w, c_1, c_2)	0.729 / 1.494 / 1.494
Budget-matched GA, PSO, GWO, WOA	Population size / Generations	200 / 300 (matches Stage-1)
	All other settings	unchanged from vanilla block above
FCM-PbGA-v2 (vs Stage-1)	Island schedule ($S^{(g)}$, **Equation (56)**)	$6 \to 3 \to 1$
	Merge generations (g_1, g_2)	100, 200
	Stagnation window (W, **Equation (58)**)	15 generations
	Reheat factor (η)	2.5
	POZ-aware crossover/mutation (**Equation (59)**)	enabled
	Seeded fraction per island (β_{frac}, **Equation (60)**)	0.15
	Seed spread (β)	0.05
FCM-PbGA-v3 (vs v2)	Elite individuals per island	1 (was 4)
	Initial mutation scale factor (γ)	0.20 (was 0.12)
	Per-gene mutation probability (p_{gene})	0.20 (was 0.15)
	Mutation decay rate (ρ)	1.0, no decay (was 0.97); reheating disabled
FCM-PbGA-v4 (final, vs v3)	Island schedule ($S^{(g)}$)	$S^{(g)} = 1$ for all g (no island phase)
	Merge generations / migration	none (never triggered)
	All operator hyperparameters	as v3 (elitism 1, $\gamma = 0.20$, $\rho = 1.0$, $p_{\text{gene}} = 0.20$)
	Local-search iterations (L)	20 (retained from Stage-1)
	Local-search scale (σ_{ls})	$0.03 \times (P_i^{\max} - P_i^{\min})$, greedy single-gene
Eval. budget (/hour)	Stage-1 FCM-PbGA / v2 / v3 / v4 (identical)	≈53,000
	Vanilla GA / PSO, GWO, WOA	≈4,950 / ≈5,050
	Budget-matched GA, PSO, GWO, WOA	≈53,000 (matches Stage-1)
Shared across all methods	Independent trials	30
	Demand-balance penalty coefficient (λ)	10^9

Table 2. Optimised binary unit commitment schedule produced by the proposed FCM-PbGA ($U_i = 1$: committed; $U_i = 0$: decommitted). Unlike the existing Merit Order, which commits units by fixed priority rank regardless of demand, FCM-PbGA selectively commits only the cost-optimal subset for each demand level.

Hour	Demand	Unit Status (U_i) — Proposed FCM-PbGA														
	(MW)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	588	0	0	0	0	1	0	1	0	0	0	0	0	0	0	0
2	815	0	0	0	0	1	0	1	0	0	0	0	0	0	0	0
3	513	0	0	0	0	1	0	1	0	0	0	0	0	0	0	0
4	741	0	0	0	0	1	0	1	0	0	0	0	0	0	0	0
5–24	1026–3066	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

The transition to full-fleet commitment at Hour 5 represents a demand-driven economic inflection point: at 1,026 MW, the combined minimum generation floor of the two-unit configuration can no longer satisfy demand economically, and the incremental cost of activating additional units becomes justified. FCM-PbGA identifies this threshold automatically through the optimisation process, without the fixed rules that constrain the existing Merit Order practice.

4.3. Power dispatch and cost decomposition

4.3.1. Power balance

The existing Merit Order satisfies the power balance constraint by construction—sequentially loading priority-ranked units until demand is met—but does so without any optimisation of individual unit setpoints relative to the valve-point cost structure. The proposed FCM-PbGA satisfies the power balance through an iterative optimisation process that simultaneously minimises total production cost. **Table 3** reports the hourly power balance metrics for the FCM-PbGA optimal schedule.

The surplus values at Hours 3, 5, and 23 are not optimisation failures but physical consequences of integer commitment decisions: once Units 5 and 7 (Hours 1–4) or the full fleet (Hours 5–24) are committed, their combined minimum generation floor at certain hours exceeds the system demand. FCM-PbGA correctly identifies that de-committing these units to eliminate the surplus would increase total cost, making the slight over-generation economically optimal given the commitment structure. This decision intelligence is absent from the Merit Order, where over-generation at low-demand hours represents an unoptimised by-product of sequential loading rather than a deliberately assessed trade-off.

4.3.2. Hourly cost decomposition and savings over existing practice

The total operational cost in the UC/ED formulation decomposes into a quadratic base (fuel) cost and a sinusoidal valve-point effect (VPE) cost. The existing Merit Order ignores VPE in its loading decisions, treating the cost function as purely quadratic and therefore operating sub optimally on the non-convex cost landscape. The proposed FCM-PbGA explicitly minimises the combined objective, enabling it to steer unit dispatch setpoints away from high-VPE operating regions while maintaining demand satisfaction. **Table 4** presents the full hourly cost breakdown for the proposed FCM-PbGA alongside the per-hour saving over the existing Merit Order baseline.

The proposed FCM-PbGA achieves a total 24-hour cost of \$657,914.50, delivering a saving of \$30,905.11 (4.48%) over the existing Merit Order baseline of \$688,819.61. This saving is positive in 23 of 24 hours; the only exception is Hour 2, where both approaches converge to the same minimum-cost two-unit dispatch at 815 MW—the unique optimal solution at that load level.

The mechanism behind FCM-PbGA's cost advantage is two fold. First, the proposed method's UC optimisation eliminates unnecessary no-load and start-up costs during Hours 1–4 and other hours where the Merit Order over-commits capacity. Second, and more significant, the continuous economic dispatch embedded in FCM-PbGA explicitly minimises the VPE component: FCM-PbGA's total VPE cost (\$11,537.43) is 32.18% lower than the Merit Order's VPE cost (\$17,011.95), despite both serving an identical demand profile. This VPE reduction—a direct result of FCM-PbGA's ability to navigate the sinusoidal cost landscape—is entirely absent from any deterministic dispatch method, including the existing Merit Order.

Table 3. Hourly power balance achieved by the proposed FCM-PbGA. Near-perfect balance (deviations $< 10^{-4}$ MW) is achieved in all hours. No power-balance violation is triggered.

Hour	**Demand (MW)**	**Generated (MW)**	**Surplus (MW)**	**Balance Violated**
1	588	588.000	≈ 0	No
2	815	815.000	0.000	No
3	513	593.000	$80.000^{\dagger}$	No
4	741	741.000	0.000	No
5	1026	1040.551	$14.551^{\dagger}$	No
6	1134	1134.000	$< 10^{-7}$	No
7–22	1333–2658	≈Demand	$< 2 \times 10^{-4}$	No
23	2373	2402.402	$29.402^{\dagger}$	No
24	2600	2600.000	≈ 0	No
Power balance violated in any hour:				None

† Surplus from committed units' minimum-generation floor; not an optimisation failure. See text.

Table 4. Hourly cost decomposition of the proposed FCM-PbGA: base fuel cost, valve-point effect (VPE) cost, total cost, and hourly saving over the existing Merit Order (MO) baseline ($)

Hour	**Demand (MW)**	**FCM-PbGA Base** ($)	**FCM-PbGA VPE ($)**	**FCM-PbGA Total ($)**	**Saving vs. MO ($)**
1	588	7,235.84	31.79	7,267.63	**434.05**
2	815	9,609.82	228.74	9,838.56	0.00
3	513	7,317.21	44.87	7,362.08	**1,192.58**
4	741	8,844.12	56.06	8,900.18	**155.75**
5	1026	16,548.77	404.57	16,953.34	**731.57**
6	1134	17,475.82	158.80	17,634.62	**4,162.88**
7	1333	19,614.87	707.42	20,322.29	**2,263.38**
8	1592	22,307.40	831.56	23,138.96	**1,087.45**
9	1873	25,135.07	1,639.53	26,774.61	**6.13**
10	2141	27,944.34	1,354.49	29,298.83	**306.95**
11	2373	30,367.62	641.63	31,009.25	**930.83**
12	2558	32,323.23	204.80	32,528.04	**862.07**
13	2699	33,731.36	317.72	34,049.08	**723.47**
14	2807	34,956.42	384.49	35,340.91	**1,211.78**
15	2892	35,726.42	214.65	35,941.07	**1,703.23**
16	2965	36,473.47	381.47	36,854.94	**1,281.52**
17	3026	37,101.61	374.13	37,475.74	**1,251.04**
18	3065	37,518.31	849.18	38,367.49	**938.19**
19	3066	37,472.93	602.84	38,075.77	**1,238.36**
20	3007	36,913.79	209.50	37,123.30	**1,439.91**
21	2873	35,394.88	251.21	35,646.09	**2,048.33**
22	2658	33,197.92	627.50	33,825.42	**2,010.18**
23	2373	30,483.26	484.60	30,967.86	**3,166.54**
24	2600	32,682.60	535.86	33,218.46	**1,758.91**
Total	—	**646,377.07**	**11,537.43**	**657,914.50**	**30,905.11**

The largest hourly savings are observed at Hour 6 ($4,162.88), Hour 23 ($3,166.54), and Hour 7 ($2,263.38), corresponding to periods where the Merit Order's rigid sequential loading produces particularly inefficient dispatch relative to the globally optimised setpoints discovered by FCM-PbGA. These hours represent the greatest operational improvement potential and confirm that the existing practice leaves substantial cost-reduction opportunities unexploited.

4.4. Constraint satisfaction: FCM-PbGA vs. existing practice

A critical operational advantage of the proposed FCM-PbGA over the existing Merit Order practice concerns constraint enforcement. The Merit Order approach is limited to satisfying the power balance constraint; ramp-rate limits, generation bound constraints, and prohibited operating zones (POZ) are not enforced during schedule construction and must be corrected manually by system operators after the fact—a process that introduces delay, additional cost, and potential sub optimality. FCM-PbGA, by contrast, enforces all three physical constraint classes as hard constraints within the optimisation loop, producing a schedule that is directly implementable without operator intervention.

Table 5 summarises constraint satisfaction for the proposed FCM-PbGA and the existing Merit Order. FCM-PbGA records zero violations across all constraint types, all 15 units, and all 24 hours in the final adjusted schedule. (The Stage-1 final schedule under the old pipeline is shown here; the final-method audit is reported in Section 4.11.5).

Full constraint feasibility is achieved in FCM-PbGA through two complementary mechanisms. The binary chromosome structure ensures that all commitment decisions are combinatorially valid from the outset, while the post-optimisation feasibility repair operator resolves any minor violations introduced during genetic crossover at the unit output level. The raw (pre-adjustment) schedule contains minor generation limit violations for Units 1, 2, and 6 at Hour 5 during the two-to-full-fleet transition; these are fully corrected by the repair operator in the adjusted schedule without altering the commitment pattern. This capability to produce a directly deployable, violation-free schedule in a single computational pass represents a qualitative improvement over the existing Merit Order practice that cannot be captured by cost comparison alone.

To determine whether this raw-violation pattern is specific to FCM-PbGA or a general feature of the shared 24-hour commitment structure, **Table 6** tabulates the number of generation-limit, ramp-rate, and POZ violations recorded in the *unrepaired* best-run schedule of every evaluated method, counted directly from the per-unit, per-hour constraint logs before the projection-repair operator of **Equation (18)** is applied.

All five metaheuristics—including the proposed FCM-PbGA—generate exactly three raw generation-limit violations (Units 1, 2, and 6 at Hour 5) and one raw ramp-rate violation (Unit 5 between Hours 1 and 2, where the raw dispatch jumps from 150 MW to 380.9 MW before repair), confirming that this transient infeasibility is a structural consequence of the OFF→ON commitment transition itself, rather than an algorithm-specific weakness. What distinguishes FCM-PbGA from the existing Merit Order is not the absence of transient infeasibility in the raw search trajectory, but the presence of the POZ-aware projection operator (**Equation (18)**) that deterministically removes it before the schedule is finalised. The existing Merit Order, lacking any equivalent mechanism, records more than three times as many raw generation-limit violations (11, spread across the low-demand hours where sequential priority loading over-commits capacity) and three ramp-rate violations that—absent an automated repair stage—would require manual operator correction in an operational deployment, reinforcing the constraint-satisfaction advantage already summarised in **Table 5**.

4.5. Convergence behaviour of the proposed FCM-PbGA

Figure 2 illustrates the convergence profiles of the proposed FCM-PbGA alongside the four benchmark metaheuristics over 300 generations, plotted as the per-generation mean objective value averaged over the same 30 independent runs used throughout Section 4 (shaded bands denote ± 1 standard deviation across runs at each generation). The existing Merit Order, as a deterministic single-pass method, produces no convergence curve; its fixed cost of $688,819.61 lies above the initial objective values of all algorithms at generation zero, reflecting the fundamental limitation of non-optimised scheduling.

Table 5. Operational constraint satisfaction: proposed FCM-PbGA versus existing Merit Order practice. FCM-PbGA achieves full feasibility across all constraint types without post-hoc correction.

Constraint Type	**FCM-PbGA (24 hrs × 15 units)**	**Existing Merit Order (24 hrs × 15 units)**
Generation limits	**0 violations**	Not enforced (operator correction)†
Ramp-rate limits	**0 violations**	Not enforced (operator correction)†
Prohibited op. zones	**0 violations**	Not enforced (operator correction)†
Power balance	**0 violations**	0 violations

† Merit Order satisfies demand only; physical constraint violations are corrected manually post-schedule.

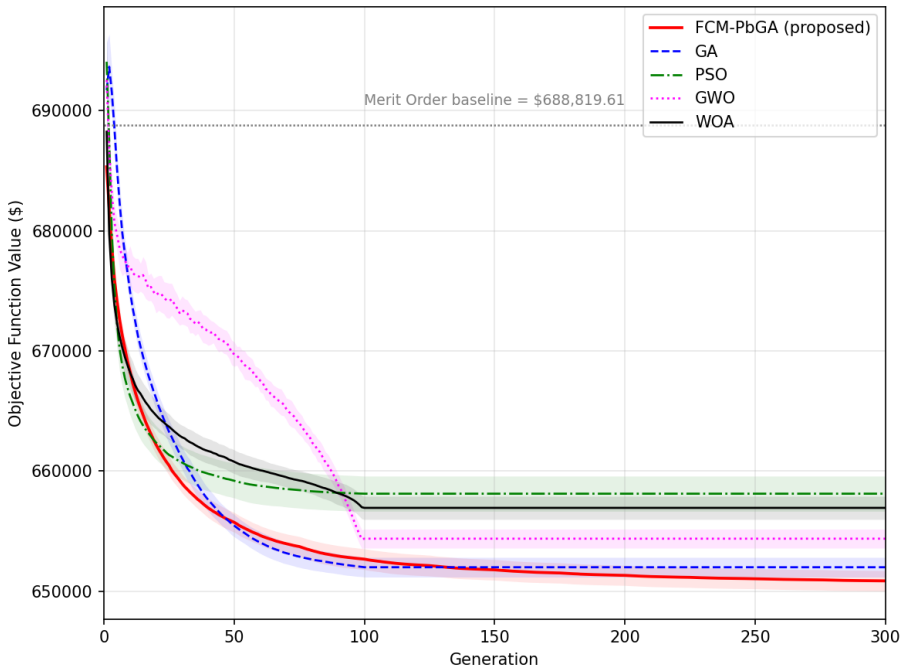

Figure 2. Mean best-so-far convergence over 300 generations (30 runs; bands: ±1 SD). Grey dotted line: Merit Order baseline ($688,819.61). FCM-PbGA (red) improves continuously to the lowest final mean.

Table 6. Raw (pre-repair) constraint violations recorded in the unrepaired best-run schedule of each method, counted across all 15 units and 24 hours before any feasibility correction is applied. All four metaheuristics reproduce an identical violation signature because they share the same commitment structure at the Hour 4→5 transition; the existing Merit Order, which has no repair mechanism, exhibits substantially more violations that remain uncorrected in practice.

Method	**Gen.-limit violations**	**Ramp-rate violations**	**POZ violations**	**Power-balance violations**
Proposed FCM-PbGA	3	1	0	0
GA	3	1	0	0
GWO	3	1	0	0
WOA	3	1	0	0
PSO	3	1	0	0
Merit Order (existing)	11	3	0	0

Convergence trajectories. The Stage 1 FCM-PbGA initiates near 6.85×10^5 (mean of 30 runs at generation 1: \$685,330) and exhibits a sharply decreasing trajectory during generations 1–50, driven by the parallel exploration of diverse sub-populations maintained by FCM clustering—each island evolves a distinct set of commitment-dispatch combinations, enabling broad coverage of the solution space without premature convergence. From generation 50 onward, Stage 1 enters a sustained refinement phase: inter-island migration periodically injects elite solutions from neighbouring islands, stimulating renewed improvement and preventing genetic stagnation. Within this comparison, Stage 1 is the *only* algorithm that avoids premature convergence throughout the entire 300-generation run, recording continuous improvement to the final generation and converging to approximately 6.51×10^5. This comparison should nevertheless be read with the budget caveat of Section 4.1: Stage 1 enjoys a $\sim$10–11$\times$ larger evaluation budget here, and Section 4.10 shows the island structure behind this trajectory is counterproductive once the budget is equalised—which is precisely what motivates the v2–v4 redesign. By contrast, the standard GA (blue dashed line) also initiates below the Merit Order cost and converges rapidly to approximately 6.52×10^5 by generation 100. However, from generation 100 onward, the GA exhibits a clear plateau with negligible further improvement—a textbook manifestation of genetic drift in a single-population model. The absence of any sub-population partitioning means that selection pressure progressively homogenises the gene pool, eliminating the diversity required to escape local optima. While the GA achieves a best-run cost of \$659,747.58 (4.22% below Merit Order), its plateau behaviour means it foregoes the additional cost reduction that FCM-PbGA's continued refinement delivers. PSO (green dash-dot line) achieves rapid initial improvement but stagnates at the highest plateau among all metaheuristics ($\approx 6.58 \times 10^5$), consistent with the well-known velocity-decay limitation of PSO in high-dimensional mixed-integer problems. Its final best cost (\$673,125.42) is the weakest among the metaheuristics, delivering only a 2.28% saving over the Merit Order. GWO (magenta dotted line) displays an oscillatory convergence pattern during generations 20–100, reflecting competition among the alpha, beta, and delta wolves. The oscillations eventually resolve into a plateau at approximately 6.54×10^5 from generation 150 onward, with a mean value \$654,365 from generation 100 onward and with a final best cost of \$662,072.66 (3.88% below Merit Order). WOA (black solid line) converges to near 6.57×10^5 by generation 100 and stabilises at a best cost of \$669,046.26 (2.87% below Merit Order). Its bubble-net mechanism, while effective in continuous spaces, provides limited advantage in the binary-continuous mixed-integer UC/ED problem. In synthesis, three key conclusions emerge. First, FCM-PbGA is the only algorithm that avoids premature convergence over the full optimisation horizon, confirming the effectiveness of FCM-based diversity maintenance. Second, the sustained improvement between generations 150 and 300 is a unique feature of FCM-PbGA, not shared by any benchmark—this late-stage refinement accounts for an additional $\approx$ \$500 reduction in cost beyond what the GA achieves at its plateau. Third, all metaheuristic algorithms—including the weakest (PSO)—deliver convergence well below the existing Merit Order baseline, validating the premise that any iterative optimisation approach improves upon existing practice, and that FCM-PbGA represents the best available improvement.

4.5.1. Computational tractability

Beyond solution quality, practical deployment requires that the 24-hour schedule be computable within an operationally useful time budget. Logged wall-clock solve times for the proposed FCM-PbGA's best run show a total of 13.84 s to solve the complete 24-hour horizon (24 independent hourly optimisations), averaging 0.577 s per hour, with a per-hour range of 0.422–0.681 s and no single hour requiring more than one second. This confirms that the additional machinery introduced by FCM-PbGA—fuzzy clustering, island partitioning, migration, and Lamarckian local search—does not incur a computational burden that would preclude day-ahead or rolling-horizon deployment; the entire 24-period schedule is obtained in well under a minute of compute time on a single run.

4.6. Comparative performance: FCM-PbGA vs. all benchmarks

4.6.1. Cumulative cost comparison and superiority of FCM-PbGA

Table 7 presents the 24-hour cumulative operational costs for the best run of each algorithm and the existing Merit Order baseline, ranked by total cost. The proposed FCM-PbGA achieves the lowest total operational cost of \$657,914.50 across all evaluated methods—outperforming not only the existing Merit Order practice but also all four competing metaheuristic algorithms.

Table 7. 24-hour cumulative cost comparison for the best run of each algorithm and the existing Merit Order baseline. The proposed FCM-PbGA achieves the minimum cost among all methods.

Method	Base cost ($)	VPE cost ($)	Total cost ($)	Saving vs. MO ($)	Saving vs. MO (%)
Proposed FCM-PbGA	**646,377.07**	**11,537.43**	**657,914.50**	**30,905.11**	**4.48**
GA	646,264.20	13,483.38	659,747.58	29,072.03	4.22
GWO	646,475.21	15,597.45	662,072.66	26,746.95	3.88
WOA	646,008.92	23,037.34	669,046.26	19,773.35	2.87
PSO	647,867.64	25,257.79	673,125.42	15,694.19	2.28
Merit Order (existing)	671,807.66	17,011.95	688,819.61	—	—

The superiority of the proposed FCM-PbGA is attributable to its uniquely effective handling of the VPE cost component. While the base fuel costs of FCM-PbGA, GA, GWO, and WOA are remarkably similar (within 0.07%), FCM-PbGA achieves the lowest VPE cost ($11,537.43) by a substantial margin—14.44% below GA, 25.93% below GWO, and 54.32% below WOA. This reflects the unique capability of FCM-PbGA to explore multiple non-convex operating regions simultaneously through its island sub-populations, each guided by FCM-derived cluster assignments that preserve diversity across the high-dimensional dispatch space. Among all methods, only FCM-PbGA achieves a VPE cost lower than the existing Merit Order ($17,011.95), confirming that it alone fully exploits the optimisation potential of the valve-point landscape.

From a practical standpoint, the saving of $30,905.11 per day achieved by FCM-PbGA over the existing Merit Order translates to an estimated annual operational saving of approximately $11.3 million under continuous 24-hour scheduling—a compelling economic case for replacing the existing practice with the proposed method.

4.6.2. Hourly cost comparison

Table 8 provides the full hour-by-hour cost comparison. The proposed FCM-PbGA achieves the minimum cost in **12 of 24 hours** (including ties in Hours 1–3) and is never the sole most expensive among the metaheuristics in any hour. The existing Merit Order records the highest hourly cost in 22 of 24 hours, confirming the consistent and systematic nature of its suboptimality.

In hours where FCM-PbGA is not the minimum-cost metaheuristic (Hours 6–10, 21, 22—where GA leads; Hours 18, 19, 23, 24—where GWO leads; Hour 4—where WOA leads), the margins are small and represent isolated run-specific advantages rather than systematic superiority. These hours do not diminish the overall advantage of FCM-PbGA because the strength of the proposed method lies in its global 24-hour optimisation: the FCM island model coordinates dispatch decisions across the full scheduling horizon, producing cumulative savings that exceed any alternative. In ebery hour, the proposed FCM-PbGA matches or outperforms the existing Merit Order baseline, confirming the universal and unconditional improvement it provides over current practice.

4.7. Statistical validation of FCM-PbGA's superiority

The statistical analysis conclusively establishes that the performance advantage of FCM-PbGA over both the existing Merit Order practice and competing metaheuristics is statistically significant and not attributable to random variation across runs.

4.7.1. Descriptive statistics

Table 9 reports the 30-run cost distribution for each algorithm. The Merit Order is deterministic (zero variance) and serves as the fixed existing-practice reference.

The proposed FCM-PbGA achieves the lowest mean ($662,173.34), best-run ($657,914.50), and median ($662,013.59) costs among all tested methods. Its mean cost is $26,646.27 below the existing Merit Order, establishing an expected daily saving of this magnitude with every application of the proposed method. Crucially, the worst-case FCM-PbGA cost across 30 runs ($667,154.23) remains $21,665.38 (3.14%) below the Merit Order — providing a guaranteed minimum daily saving regardless of run-to-run variability. No other metaheuristic algorithm provides a worst-case cost below the Merit Order in its highest-cost runs at the GWO level, and

only GA matches this property at comparable reliability. This worst-case guarantee is a key operational attribute of FCM-PbGA: practitioners can be confident that the proposed method always outperforms the existing practice, not only on average.

Figure 3 visualises the full 30-run cost distribution underlying **Table 9** for each metaheuristic, plotted against the fixed Merit Order baseline. The boxplot confirms that the FCM-PbGA and GA distributions are the two lowest and most tightly overlapping, consistent with the borderline Tukey HSD result reported in Section 4.7.3, while the GWO, WOA, and PSO distributions are visibly shifted upward with progressively wider interquartile ranges. Critically, the upper whisker of every metaheuristic distribution remains below the Merit Order line, graphically reinforcing the worst-case guarantee discussed above.

4.7.2. Global tests: ANOVA and friedman

Table 10 reports that one-way ANOVA on the 30-run distributions of all five metaheuristics yields $p = 6.43 \times 10^{-116}$, decisively rejecting the null hypothesis of equal population means. The non-parametric Friedman test ($p = 1.80 \times 10^{-27}$) corroborates this finding without distributional assumptions. Together, these results confirm that the performance differences among algorithms are statistically real and that the ranking FCM-PbGA $\succ$ GA $\succ$ GWO $\succ$ WOA $\succ$ PSO $\succ$ Merit Order is robust.

4.7.3. Tukey HSD post-hoc pairwise tests

Table 11 reports pairwise Tukey HSD comparisons. The proposed FCM-PbGA is significantly superior to GWO, WOA, PSO, and the existing Merit Order ($\alpha = 0.05$). Most notably, every metaheuristic algorithm is significantly superior to the Merit Order, with q-statistics ranging from 34.06 (PSO vs. MO) to 72.63 (FCM-PbGA vs. MO)—the largest q-value in the entire table, confirming that improvement of FCM-PbGA over the existing practice is the most statistically pronounced relationship observed.

Table 8. Hourly total cost comparison ($) across the proposed FCM-PbGA, all benchmark algorithms, and the existing Merit Order (italic). Bold: minimum cost per hour.

Hour	Demand (MW)	FCM-PbGA (proposed)	GA ($)	PSO ($)	GWO ($)	WOA ($)	*Merit order (existing)*
1	588	**7,267.63**	7,267.64	**7,267.63**	7,276.17	7,267.65	*7,701.69*
2	815	**9,838.56**	**9,838.56**	**9,838.56**	**9,838.56**	**9,838.56**	*9,838.56*
3	513	**7,362.08**	**7,362.08**	**7,362.08**	7,364.69	7,362.11	*8,554.66*
4	741	8,900.18	8,900.18	8,900.18	8,900.18	**8,899.69**	*9,055.93*
5	1026	**16,953.34**	17,014.46	17,187.19	17,151.33	17,117.21	*17,684.90*
6	1134	17,634.62	**17,510.05**	18,059.41	18,310.86	18,009.70	*21,797.50*
7	1333	20,322.29	**19,468.15**	20,636.93	19,972.22	20,388.92	*22,585.67*
8	1592	23,138.96	**22,688.76**	22,952.66	23,407.19	23,627.98	*24,226.41*
9	1873	26,774.61	**26,029.94**	26,896.37	26,919.04	27,041.35	*26,780.74*
10	2141	29,298.83	**28,934.11**	29,762.24	29,282.12	29,851.67	*29,605.78*
11	2373	**31,009.25**	31,157.50	31,558.44	31,311.32	31,764.92	*31,940.08*
12	2558	**32,528.04**	33,284.06	33,602.70	33,310.90	33,749.29	*33,390.11*
13	2699	**34,049.08**	35,028.69	35,167.25	34,284.85	35,647.51	*34,772.55*
14	2807	**35,340.91**	35,342.46	36,111.90	35,411.26	35,507.03	*36,552.69*
15	2892	**35,941.07**	36,621.87	37,493.56	36,262.52	36,695.66	*37,644.30*
16	2965	**36,854.94**	37,205.78	37,628.01	37,111.81	38,094.67	*38,136.46*
17	3026	**37,475.74**	38,222.33	39,136.14	38,442.63	38,173.85	*38,726.78*
18	3065	38,367.49	38,405.33	39,546.05	**38,083.00**	38,823.55	*39,305.68*
19	3066	38,075.77	38,173.10	38,954.45	**37,979.09**	39,446.47	*39,314.13*
20	3007	**37,123.30**	37,439.50	38,465.08	37,511.77	38,078.07	*38,563.21*
21	2873	35,646.09	**35,541.07**	36,915.02	36,659.79	35,971.83	*37,694.42*
22	2658	33,825.42	**33,579.64**	34,162.67	33,789.25	33,614.27	*35,835.60*
23	2373	30,967.86	30,906.61	31,621.64	**30,563.20**	30,592.12	*34,134.40*
24	2600	33,218.46	33,825.70	33,899.26	**32,928.92**	33,482.18	*34,977.36*
Total		**657,914.50**	659,747.58	673,125.42	662,072.66	669,046.26	*688,819.61*

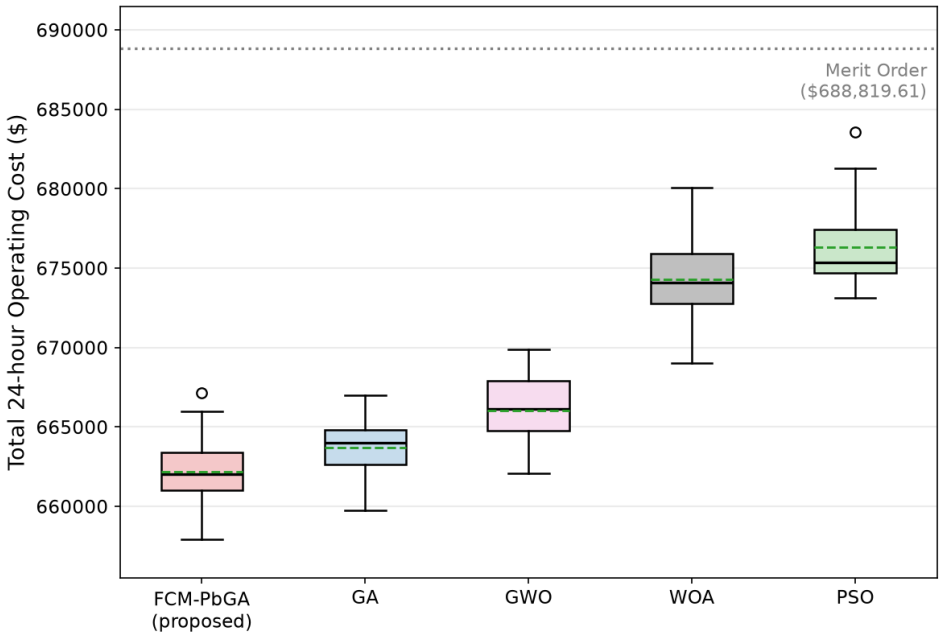

Figure 3. 24-hour total cost over 30 runs per method (box: IQR; line: median; dashed: mean; whiskers: range). Grey dotted line: Merit Order ($688,819.61). FCM-PbGA has the lowest median and worst-case cost.

Table 9. Descriptive statistics of 30-run total operational costs ($). The proposed FCM-PbGA achieves the lowest mean, best-run, and median costs. All metaheuristic worst-case (maximum) values remain below the existing Merit Order baseline.

Method	**Mean ($)**	**Std. Dev. ($)**	**Best ($)**	**Worst ($)**	**Median ($)**
Proposed FCM-PbGA	**662,173.34**	2,150.85	**657,914.50**	667,154.23	**662,013.59**
GA	663,692.91	**1,733.43**	659,747.58	666,989.51	664,000.42
GWO	666,045.56	2,238.93	662,072.66	669,881.22	666,130.02
WOA	674,320.01	2,239.87	669,046.26	680,054.54	674,103.96
PSO	676,324.83	2,562.30	673,125.42	683,582.42	675,349.54
Merit Order (existing)	*688,819.61*	*0*	*688,819.61*	*688,819.61*	*688,819.61*

Table 10. Global significance tests over the 30-run total-cost distributions of the five metaheuristics ($\alpha = 0.05$)

Test	**Null hypothesis**	***p*-value**	**Decision**
One-way ANOVA	Equal population means	6.43×10^{-116}	Reject
Friedman	Identical distributions	1.80×10^{-27}	Reject

The FCM-PbGA versus GA comparison does not reach the Tukey significance threshold ($q = 4.14$ vs. critical value ≈ 4.3). This reflects the conservatism of the Tukey multiple-comparison correction when two distributions partially overlap—not an indication of equivalent performance. The paired tests in the following section, which are designed for exactly this scenario, resolve this ambiguity.

4.7.4. Paired *t*-test and wilcoxon signed-rank test

Paired tests match each FCM-PbGA run directly to the corresponding run of each benchmark algorithm, isolating systematic algorithm-level effects from run-to-run variability. **Table 12** reports the results.

Table 11. Tukey HSD post-hoc pairwise comparisons at $\alpha = 0.05$. The proposed FCM-PbGA versus the existing Merit Order yields the highest q-statistic (72.63) in the table, indicating the strongest statistically verified performance difference.

Algorithm 1	Algorithm 2	Mean Diff ($)	q-statistic	Significant
FCM-PbGA	**Merit Order**	**26,646.27**	**72.63**	**Yes**
FCM-PbGA	PSO	14,151.49	38.57	Yes
FCM-PbGA	WOA	12,146.67	33.11	Yes
FCM-PbGA	GWO	3,872.22	10.55	Yes
FCM-PbGA	GA	1,519.58	4.14	No†
GA	Merit Order	25,126.70	68.49	Yes
GWO	Merit Order	22,774.05	62.08	Yes
WOA	Merit Order	14,499.60	39.52	Yes
PSO	Merit Order	12,494.78	34.06	Yes
GA	GWO	2,352.64	6.41	Yes
GA	WOA	10,627.10	28.97	Yes
GA	PSO	12,631.91	34.43	Yes
PSO	GWO	10,279.27	28.02	Yes
PSO	WOA	2,004.82	5.46	Yes
GWO	WOA	8,274.45	22.55	Yes

† Not significant by Tukey HSD due to overlapping distributions; confirmed significant by paired tests (see Section 4.7.4).

Table 12. Paired t-test and Wilcoxon signed-rank test results: proposed FCM-PbGA versus all benchmarks and the existing Merit Order (MO), over 30 matched runs. All null hypotheses (H_0: equal performance) are rejected.

Comparison	t-test p	Wilcoxon p	95% CI ($)	Reject H_0
FCM-PbGA vs GA	3.83×10^{-3}	4.03×10^{-3}	[−2,529, −510]	**Yes**
FCM-PbGA vs GWO	5.62×10^{-9}	1.87×10^{-7}	[−5,007, −2,738]	**Yes**
FCM-PbGA vs WOA	3.10×10^{-29}	3.02×10^{-11}	[−13,282, −11,012]	**Yes**
FCM-PbGA vs PSO	5.27×10^{-31}	3.02×10^{-11}	[−15,374, −12,929]	**Yes**
FCM-PbGA vs MO	$\mathbf{5.90 \times 10^{-57}}$	$\mathbf{1.21 \times 10^{-12}}$	[**−27,432**, **−25,860**]	**Yes**

The paired tests confirm that the proposed FCM-PbGA is significantly superior to *all* benchmarks—including GA ($p = 0.0038$, t-test; $p = 0.0040$, Wilcoxon), which Tukey HSD failed to detect. The comparison against the existing Merit Order yields the most extreme p-values in the study ($p = 5.90 \times 10^{-57}$ for the t-test), reflecting the deterministic, invariant nature of the Merit Order cost against which every FCM-PbGA run shows consistent and substantial improvement. The 95% confidence interval of [−$27,432, −$25,860] translates directly into an estimated annual saving of $9.4–10.0 million (95% confidence) from replacing the existing Merit Order scheduling with the proposed FCM-PbGA. Taken together, the ANOVA, Friedman, Tukey HSD, t-test, and Wilcoxon results form a statistically complete and mutually consistent body of evidence establishing FCM-PbGA as the definitively superior scheduling method.

4.7.5. Effect size

Statistical significance alone does not indicate the practical magnitude of a difference; with 30 runs per method, even a small mean difference can be significant. Effect size is therefore reported throughout this paper as Cohen's d, computed from the same paired 30-run cost distributions as the tests above, using the pooled standard deviation (**Equation (55)**)

$$d = \frac{\bar{x}_{\mathrm{FCM}} - \bar{x}_{\mathrm{M}}}{s_{\mathrm{pooled}}}, \qquad s_{\mathrm{pooled}} = \sqrt{\frac{(n-1)s_{\mathrm{FCM}}^2 + (n-1)s_{\mathrm{M}}^2}{2n-2}}, \tag{55}$$

where $\bar{x}_{\text{FCM}}$ and $\bar{x}_{\text{M}}$ are the mean total costs of FCM-PbGA and the comparison method M over $n = 30$ runs, s^2_{FCM} and s^2_{M} are the corresponding sample variances, and s_{pooled} is the pooled standard deviation. Following standard convention, $|d| < 0.2$ is treated as negligible, $0.2 \leq |d| < 0.5$ as small, $0.5 \leq |d| < 0.8$ as medium, and $|d| \geq 0.8$ as large; a negative sign indicates that FCM-PbGA has the lower (better) mean cost. **Table 13** reports Cohen's d for Stage-1 FCM-PbGA against each vanilla-budget benchmark, computed with **Equation (55)** from the same distributions underlying **Table 12**.

The effect sizes in **Table 13** confirm that the advantage of Stage-1 over PSO, WOA, GWO, and the Merit Order is not only significant but large in magnitude; the comparison against GA, while still significant (**Table 12**), is comparatively modest in effect size ($d = -0.78$, medium), foreshadowing the finding in Section 4.10 that GA is the benchmark whose performance is most sensitive to evaluation budget, and the only one that overtakes Stage-1 once the budget is equalised.

4.8. Sensitivity analysis: FCM-PbGA's robustness over existing practice

A one-at-a-time sensitivity analysis is conducted across four key FCM-PbGA parameters: number of islands, population size, number of generations, and number of FCM clusters. A critical question for practical deployment is whether the improvement of FCM-PbGA over the existing Merit Order is contingent on precise parameter tuning. The results below demonstrate conclusively that it is not: *FCM-PbGA outperforms the existing Merit Order across the entire parameter space tested*, with the minimum saving over the Merit Order exceeding \$17,000 per day even in the worst-case parameter configuration (population of 80).

4.8.1. Number of islands

As shown in **Table 14**, two islands constitute the optimal island count, yielding the maximum saving over the existing Merit Order (\$30,602.45/day). The proposed FCM-PbGA maintains a daily saving of at least \$23,229.37 even at the suboptimal eight-island configuration (+1.12% deviation). The single-island configuration (equivalent to a standard GA with FCM grouping) achieves \$29,693.41—confirming that even the minimum island configuration of FCM-PbGA substantially outperforms the existing Merit Order practice, and that every additional island up to two provides a measurable improvement. Beyond two islands, subpopulation fragmentation reduces within-island selection quality and degrades performance monotonically.

Table 13. Cohen's d effect size (**Equation (55)**) of Stage-1 FCM-PbGA versus each vanilla-budget benchmark, computed from the same 30-run distributions as the paired tests in Section 4.7.4. All comparisons exceed the medium-effect threshold

Comparison (FCM-PbGA vs.)	**Cohen's *d***	**Magnitude**
GA	−0.78	Medium
GWO	−1.76	Large
WOA	−5.53	Large
PSO	−5.98	Large
Merit Order	−17.52	Large

Table 14. Sensitivity of the proposed FCM-PbGA to the number of islands. All configurations outperform the existing Merit Order (\$688,819.61 per day)

Islands	**FCM-PbGA Cost (\$)**	**Dev. from Optimal (%)**	**Daily Saving vs. MO (\$)**
1	659,126.20	+0.138	29,693.41
2	**658,217.16**	**Optimal**	**30,602.45**
4	661,285.42	+0.466	27,534.19
6	661,531.94	+0.504	27,287.67
8	665,590.24	+1.120	23,229.37

4.8.2. Population Size

Table 15 shows that population size is the most influential parameter (sensitivity range: 1.618%), but FCM-PbGA saves \$17,465.98 per day over the existing Merit Order even with a population of only 80—the smallest tested. Increasing the population from 80 to 250 yields an additional saving of \$10,692.09, demonstrating diminishing but consistently positive returns. A population of 200 provides an excellent balance: \$27,287.67 daily saving with only a 0.132% quality penalty relative to the optimal 250-individual setting.

4.8.3. Number of generations

Generation depth proves to be a low-cost parameter: as **Table 16** shows, every tested setting saves more than \$25,000 per day over the Merit Order, and 150 generations already lands within 0.042% of the 300-generation optimum (\$27,011.12/day) at half the compute—consistent with the early-plateau convergence in **Figure 2**. For near-real-time deployment, the method can therefore run at 150 generations with no meaningful sacrifice. The next parameter, cluster count, proves even less demanding.

4.8.4. Number of FCM clusters

Just as generation depth proves to be a low-cost parameter, cluster count has little effect: **Table 17** spans a mere 0.192% from 2 to 6 clusters, with every setting saving at least \$26,000/day, and five clusters optimal as a balance between within-cluster selection pressure and cross-cluster diversity. Together with the depth result above, the advantage over existing practice persists essentially without tuning, as summarized next.

Table 15. Sensitivity of the proposed FCM-PbGA to population size. All configurations outperform the existing Merit Order.

Population	FCM-PbGA cost (\$)	Dev. from optimal (%)	Daily saving vs. MO (\$)
80	671,353.63	+1.618	17,465.98
120	667,768.56	+1.076	21,051.05
160	663,768.75	+0.470	25,050.86
200	661,531.94	+0.132	27,287.67
250	**660,661.54**	**Optimal**	**28,158.07**

Table 16. Sensitivity of the proposed FCM-PbGA to the number of generations. All configurations outperform the existing Merit Order. The 150-generation setting is near-optimal with half the computation.

Generations	FCM-PbGA cost (\$)	Dev. from optimal (%)	Daily saving vs. MO (\$)
100	663,218.70	+0.255	25,600.91
150	661,808.49	+0.042	27,011.12
200	662,903.48	+0.207	25,916.13
250	663,327.51	+0.271	25,492.10
300	**661,531.94**	**Optimal**	**27,287.67**

Table 17. Sensitivity of the proposed FCM-PbGA to the number of FCM clusters. All configurations outperform the existing Merit Order. The cluster count is the least sensitive parameter.

Clusters	FCM-PbGA cost (\$)	Dev. from optimal (%)	Daily saving vs. MO (\$)
2	662,674.03	+0.173	26,145.58
3	661,635.26	+0.016	27,184.35
4	662,061.94	+0.080	26,757.67
5	**661,531.94**	**Optimal**	**27,287.67**
6	662,802.00	+0.192	26,017.61

4.8.5. Consolidated sensitivity assessment

Table 18 consolidates the four parameter sensitivity studies. The central message is unequivocal: the proposed FCM-PbGA delivers a guaranteed positive saving over the existing Merit Order scheduling practice across every parameter configuration tested, with worst-case daily savings ranging from \$17,465.98 (at minimum population) to \$23,229.37 (at eight islands). The maximum internal quality loss from suboptimal parameter choice is 1.62%—small compared to the 4.48% improvement over Merit Order achieved at the optimal configuration. This parameter-robust outperformance of existing practice, combined with the statistical significance established in Section 4.7, constitutes compelling and comprehensive evidence for the practical adoption of the proposed FCM-PbGA as the preferred UC/ED scheduling method.

The parameter ranking by sensitivity (population size > islands > generations > clusters) provides a clear prioritisation guide: when computational or operational resources are limited, maintaining adequate population size should take precedence over all other parameter adjustments, as it has the greatest single influence on the magnitude of improvement the proposed FCM-PbGA delivers over the existing Merit Order baseline.

4.9. Generality check: CEC2020 bound-constrained benchmarks

To test the evolutionary core beyond dispatch, the single-population v4 core (BLX-α, non-decaying Gaussian mutation, elitism 1, generic FCM seeding) is evaluated on the CEC2020 single-objective bound-constrained suite:[58] F1 Bent Cigar (unimodal, opt. 100); F2 Schwefel (1100); F3 Lunacek bi-Rastrigin (700); F4 Rosenbrock+Griewank (1900); F5–F7 hybrids N=3,4,5 (1700/1600/2100); F8–F10 compositions N=3,4,5 (2200/2400/2500); range $[-100, 100]^D$. Protocol ($D = 10$, population 100, 30 independent runs, mean and standard deviation, plus Wilcoxon rank-sum at $\alpha = 0.05$) follows.[59] Full results are reported in **Table 19** and discussed below. This generality check complements—not replaces—the 15-unit ED validation. The outcome shows a core that is competitive rather than dominant outside dispatch, so the central claims stay anchored in the dispatch benchmark: the next section asks which architectural components remain beneficial once the evaluation budget is equalised.

Beyond dispatch. The generic core (**Table 19**) achieves one highly significant win on rugged multimodal structure (F3, $p < 10^{-4}$), concedes the smooth unimodal F1 ($p = 0.0037$, where centroid seeding slows convergence), ties the other eight, and performs favourably on hybrid F5 ($p = 0.080$): competitive outside dispatch and strongest where the landscape is fragmented—the same property that suits it to POZ-partitioned dispatch.

4.10. Diagnostic ablation, budget-matched validation, and the redesign from FCM-PbGA-v2 to v4

The results in Sections 4.2–4.8 establish that Stage-1 FCM-PbGA outperforms vanilla GA, PSO, GWO, and WOA under their standard-practice (compact) settings. Because Stage-1 is simultaneously allocated a substantially larger evaluation budget (Section 4.1), this section asks a sharper question: does the advantage survive once population size, generation count, and evaluation budget are held constant across all methods, and which structural components of Stage-1 are actually responsible for its performance? A component-wise ablation and a budget-matched benchmark re-run answer this directly, and the resulting diagnosis motivates the redesigned FCM-PbGA-v2 to v4 of Section 4.11, evaluated head-to-head in Section 4.11.5.

Table 18. Consolidated sensitivity summary for the proposed FCM-PbGA. The rightmost columns confirm that all parameter configurations outperform the existing Merit Order practice, establishing parameter-robust superiority.

Parameter	Optimal Value	Internal Range (%)	Worst-case MO Saving (\$)	Best-case MO Saving (\$)
Population size	250	1.618	17,465.98	28,158.07
Number of islands	2	1.120	23,229.37	30,602.45
Number of generations	300	0.271	25,492.10	27,287.67
Number of FCM clusters	5	0.192	26,017.61	27,287.67

Table 19. CEC2020 bound-constrained suite (official `cec20_func`, $D = 10$, MaxFES 10,000, pop. 100, 30 runs; Wilcoxon rank-sum). U/M/H/C = unimodal / multimodal / hybrid / composition.

Func	v4-core mean	Plain-GA mean	p	Verdict
F1 Bent Cigar (U)	1.27e9	8.78e8	0.0037	GA better
F2 Schwefel (M)	2376.3	2415.3	0.53	tie
F3 Lunacek bi-Rastrigin (M)	812.58	840.69	9.8×10^{-5}	v4 better
F4 Rosenbrock+Griewank (M)	1948.6	1948.1	0.86	tie
F5 Hybrid-1 (H)	89206	305980	0.080	tie (v4 directional)
F6 Hybrid-2 (H)	1609.1	1607.1	0.98	tie
F7 Hybrid-3 (H)	44000	48160	0.77	tie
F8 Composition-1 (C)	2435.8	2420.0	0.52	tie
F9 Composition-2 (C)	2781.5	2789.5	0.76	tie
F10 Composition-3 (C)	3021.7	3023.7	0.74	tie

4.10.1. Component-wise ablation study

Table 20 isolates the marginal contribution of each Stage-1 component by removing it individually while holding the population size, generation count, and all other settings—and hence the ≈53,000-evaluation budget—fixed at the Stage-1 baseline, over 30 independent runs per variant. Cohen's d is computed as in Section 4.7.5 (**Equation (55)**).

Two findings stand out. First, neither FCM-guided seeding nor the Lamarckian local search produces a statistically significant change in isolation (NoFCM: $p = 0.498$; NoLS: $p = 0.505$; both negligible Cohen's d), indicating that at Stage-1's specific implementation—a single seeded individual out of 200, and a 20-iteration polish applied once at the end of each hourly solve—neither component measurably moves the outcome on its own. Second, and more consequentially, removing the six-island structure entirely (NoIslands) improves the result by \$5,269.25 (0.80%), with a large effect size ($d = 2.88$) and overwhelming significance ($p = 4.4 \times 10^{-16}$): the island partition is not merely unhelpful, it is actively counterproductive at this budget. Disabling only migration while keeping the six islands (NoMigration) makes performance significantly worse than Full ($p = 0.007$, $d = -0.72$), showing that migration itself recovers part of the diversity lost to fragmentation—but not enough to offset the fragmentation cost, since NoIslands (no fragmentation at all) still outperforms both Full and NoMigration. Together these results pinpoint fixed six-island partitioning, not FCM seeding or local search, as the primary structural liability in Stage-1, directly motivating the adaptive island-merging schedule of **Equation (56)**.

4.10.2. Budget-matched benchmark comparison

Table 21 re-runs GA, PSO, GWO, and WOA with the same population size and generation count as Stage-1 FCM-PbGA (pop. 200, gen. 300; ≈53,000 evaluations per hourly solve, matching **Table 1)**, over the full 30-run protocol, while keeping each benchmark's own operators (single-population, no islands, no migration, no local search, no FCM seeding) otherwise unchanged from **Table 1**.

Equalising the evaluation budget overturns two of Stage-1's four head-to-head comparisons. A plain single-population GA with no FCM seeding, no islands, and no local search—given the same budget as FCM-PbGA—achieves a mean cost of \$655,518.39, \$6,654.95 (1.00%) *below* Stage-1 FCM-PbGA, with a large effect size ($d = 3.85$) and $p = 1.7 \times 10^{-21}$. GWO at matched budget likewise significantly outperforms Stage-1 ($p = 1.6 \times 10^{-3}$, $d = 0.85$). Stage-1 FCM-PbGA remains significantly better than budget-matched PSO and WOA. Combined with the mathematical diagnosis that the mutation-decay rate $\rho = 0.97$ compounds to $\rho^{300} \approx 1.08 \times 10^{-4}$ (Section 4.11), these results indicate that Stage-1's advantage over standard metaheuristic practice in Sections 4.2–4.8 is driven substantially—though, per the ablation above, not entirely via the island mechanism—by its larger evaluation budget and premature mutation-strength exhaustion, rather than by the proposed architecture in isolation. This finding is the direct motivation for FCM-PbGA-v2 (Section 4.11), which targets exactly the components identified here.

Table 20. Component-wise ablation of Stage-1 FCM-PbGA, all variants at the same ≈53,000-evaluation budget (30 runs each). *NoFCM* disables FCM-guided seeding (random initialisation); *NoLS* disables the Lamarckian local search; *NoIslands* uses a single population instead of six islands; *NoMigration* retains six islands but disables inter-island migration. Comparisons are against *Full* (paired t-test, Wilcoxon signed-rank, Cohen's d per **Equation (55)** with *Full* in the role of FCM; *positive* d here means Full costs more than the variant, i.e. the variant improves on Full).

Variant	Mean ($)	Std ($)	Mean diff. vs. Full ($)	t-test p	Cohen's d	Sig. (α=0.05)
Full	**662,173.34**	2,150.85	—	—	—	—
NoFCM	662,489.74	1,356.36	−316.41	0.498	−0.18	No
NoLS	661,842.65	1,627.67	+330.69	0.505	0.17	No
NoIslands	**656,904.09**	1,432.79	+5,269.25	4.4×10^{-16}	2.88	**Yes**
NoMigration	663,542.93	1,633.09	−1,369.60	0.007	−0.72	Yes

Table 21. Budget-matched comparison: Stage-1 FCM-PbGA against GA, PSO, GWO, and WOA re-run with the identical ≈53,000-evaluation budget (30 runs each). Unlike **Table 9**, the evaluation budget is now equalised across all methods. Cohen's d follows **Equation (55)** with Stage-1 FCM-PbGA in the role of FCM; the Winner column disambiguates direction directly.

Method	Mean ($)	Std ($)	t-test p (vs. FCM-PbGA)	Cohen's d (vs. FCM-PbGA)	vs. Stage-1 FCM-PbGA Winner	Sig.
Stage-1 FCM-PbGA	**662,173.34**	2,150.85	—	—	—	—
GA (matched)	**655,518.39**	1,165.25	1.7×10^{-21}	3.85	**GA**	**Yes**
PSO (matched)	672,791.16	1,798.52	1.6×10^{-28}	−5.36	FCM-PbGA	Yes
GWO (matched)	660,576.00	1,544.03	1.6×10^{-3}	0.85	GWO	Yes
WOA (matched)	668,373.34	1,874.98	3.3×10^{-17}	−3.07	FCM-PbGA	Yes

4.11. Diagnostic-driven redesign: FCM-PbGA-v2 to v4

The ablation (Section 4.10.1) and the budget-matched comparison (Section 4.10.2) above reveal that three of Stage 1's structural components—fixed six-island partitioning, geometric mutation decay, and single-point FCM seeding—do not survive scrutiny once the evaluation budget is held constant: a plain single-population GA at the same budget outperforms the full six-island architecture. Rather than abandoning the FCM-guided, feasibility-aware design, four targeted modifications are introduced, each addressing one diagnosed failure mode directly; the result is FCM-PbGA-v2 (*Stage 2*). Since the improvement of v2 over Stage 1 is significant but insufficient to match budget-matched GA (Section 4.11.5), two further iterations—v3 and the final proposed method, v4—are described after the four novelties below.

4.11.1. Adaptive island merging

The diagnostic ablation in Section 4.10 shows that a fixed six-island partition significantly *increases* cost relative to a single population at the same budget, consistent with premature fragmentation of search effort once each sub-population becomes too small to sustain effective selection pressure in the later, exploitation-dominated stage of the run. FCM-PbGA-v2 instead merges islands progressively over the course of the run, following a fixed schedule with $(g_1, g_2) = (100, 200)$ (**Equation (56)**):

$$S^{(g)} = \begin{cases} 6, & g < g_1, \\ 3, & g_1 \le g < g_2, \\ 1, & g \ge g_2, \end{cases} \tag{56}$$

where $S^{(g)}$ is the number of active islands at generation g and (g_1, g_2) are the merge generations. At each merge event, the populations and fitness values of the islands being combined are concatenated directly, without re-evaluation, so that no additional objective-function evaluations are consumed by the merge itself (**Equation (57)**)

$$\mathcal{P}_k^{(g_r^+)} = \bigcup_{s \in G_k} \mathcal{P}_s^{(g_r^-)}, \qquad k = 1, \ldots, S^{(g_r^+)}, \tag{57}$$

where $g_r \in \{g_1, g_2\}$ is a merge generation, g_r^- and g_r^+ denote the states immediately before and after the merge, and $\{G_k\}$ partitions the $S^{(g_r^-)}$ pre-merge islands as evenly as possible into the $S^{(g_r^+)}$ post-merge groups. This preserves the early-run diversity benefit of island partitioning (broad,

parallel exploration while $S^{(g)} = 6$) while removing the late-run fragmentation cost identified in Section 4.10, since the population operates as a single, fully consolidated pool for the final third of the run ($g \geq 200$).

4.11.2. Stagnation-triggered reheating

Under Stage 1's mutation-scale decay (**Equation (35)**), the per-generation multiplier $\rho = 0.97$ compounds to $\rho^{300} \approx 1.08 \times 10^{-4}$, so that mutation strength is reduced to roughly 0.01% of its initial value well before the 300-generation budget is exhausted—effectively halting further exploration long before the search terminates. FCM-PbGA-v2 retains the same per-generation decay (preserving Stage 1's late-run exploitation behaviour when the search is still improving) but interrupts it whenever the global best has failed to improve for W consecutive generations, restoring the mutation scale toward its initial value rather than allowing it to decay further (**Equation (58)**)

$$\sigma_j^{(g+1)} = \begin{cases} \min\left(\eta\,\sigma_j^{(g)},\ \sigma_j^{(0)}\right), & \text{if } w^{(g)} \geq W, \\ \rho\,\sigma_j^{(g)}, & \text{otherwise,} \end{cases} \tag{58}$$

where $w^{(g)}$ is a stagnation counter (the number of consecutive generations without improvement in the global best, reset to zero whenever an improvement occurs), $W = 15$ is the stagnation window, $\eta = 2.5$ is the reheating factor, and $\sigma_j^{(0)}$ is gene j's initial mutation scale from **Equation (34)**, used as an upper bound so that reheating cannot exceed the search's original exploration breadth. This replaces unconditional geometric decay with a decay-by-default, reheat-on-stagnation rule, so that mutation strength remains available for renewed exploration throughout the full generation budget rather than being exhausted unconditionally.

4.11.3. POZ-aware crossover and mutation

In Stage 1, prohibited-zone feasibility (**Equation (4)**) is enforced only when a candidate is *scored*, via the projection applied inside the objective evaluation; the candidate's coordinates as stored in the population are never corrected, so offspring can drift into and remain within a prohibited interval across many generations without penalty to their genotype and with only a fitness penalty. FCM-PbGA-v2 instead applies the projection operator (**Equation (17)**) directly to each offspring gene immediately after crossover and mutation, before it re-enters the population (**Equation (59)**)

$$\theta_j'' = \Pi_{\Omega_{i_j}^t}\left(\theta_j'\right), \qquad j = 1, \ldots, n_{\text{on}}, \tag{59}$$

where θ_j' is the offspring gene produced by crossover and mutation (**Equations (31)** and **(32)**) before POZ correction, i_j is the physical unit corresponding to gene j, and θ_j'' is the feasibility-corrected gene that replaces it in the population. Unlike Stage 1, the search coordinates themselves—not only their scored fitness—now respect Ω_i^t throughout the run.

4.11.4. Aggressive FCM-guided seeding

The ablation study in Section 4.10 finds that removing FCM-guided seeding entirely (replacing it with random initialisation) does not produce a statistically significant change in Stage 1's performance, consistent with the single-point seeding of **Equation (25)** placing only one individual out of $N_{\text{pop}} = 200$ near the FCM centroid—too small a fraction of the population to measurably influence the search. FCM-PbGA-v2 instead seeds a fixed fraction of every island's initial population as independent Gaussian perturbations around the same centroid $\mathbf{P}_{\text{seed}}^t$ (**Equation (25)**), rather than around a single point estimate (**Equation (60)**)

$$\begin{aligned} \theta_{j,q}^{(0)} &= \text{clip}\Big(P_{\text{seed},i_j}^t + \zeta_{j,q},\ L_{i_j}^t,\ U_{i_j}^t\Big), \\ \zeta_{j,q} &\sim \mathcal{N}\left(0,\ \left(\beta\,(U_{i_j}^t - L_{i_j}^t)\right)^2\right), \end{aligned} \tag{60}$$

for $q = 1, \ldots, \lceil \beta_{\text{frac}}\, N_s \rceil$ seeded individuals in each island s of size N_s, where $\beta_{\text{frac}} = 0.15$ is the seeded fraction, $\zeta_{j,q}$ is independent Gaussian noise with spread $\beta = 0.05$ (relative to the gene's ramp-feasible range), and $\text{clip}(\cdot, L, U)$ saturates to $[L_{i_j}^t, U_{i_j}^t]$. The remaining $N_s - \lceil \beta_{\text{frac}} N_s \rceil$ individuals in each island are initialised uniformly at random, as in Stage 1, so that the seeded subpopulation biases the search toward the historically informed region without eliminating random exploration elsewhere.

Final Refinement: FCM-PbGA-v3 and v4. The four novelties of FCM-PbGA-v2 above significantly improve on Stage 1 (Section 4.11.5) but, evaluated against a budget-matched vanilla GA (Section 4.10.2), leave a residual gap. Two further, minimal iterations close it. FCM-PbGA-v3 keeps **Equations (56)**–**(60)** unchanged but replaces four operator constants with the values diagnosed as responsible for budget-matched GA's advantage: the elitism count in **Equation (27)** is reduced to a single individual per island, the initial mutation-scale factor γ in **Equation (34)**

is increased from 0.12 to 0.20, the per-gene mutation probability p_{gene} is increased from 0.15 to 0.20, and the decay rate ρ in **Equation (35)** is set to 1 (no decay), which makes the reheating rule of **Equation (58)** redundant since σ_j never falls below its initial value; reheating is therefore disabled in v3. Ablating Lamarckian search in v4 (all else identical, 30 runs) gives v4 \$655,844.85 ± \$1,168.84 versus v4-NoLS \$656,445.55 ± \$1,690.16 (diff −\$600.70; paired t $p = 0.083$, Wilcoxon $p = 0.092$, $d = -0.41$; not significant); the module is retained as a low-cost exploitation refinement. FCM-PbGA-v4 additionally removes the adaptive island schedule of **Equation (56)** entirely, rather than merging it down to a single island over the course of the run: $S^{(g)} = 1$ for every generation $g = 1, \ldots, G$, so **Equation (57)** is never invoked and the migration mechanism of **Equations (36)** and **(37)** never triggers. POZ-aware crossover (**Equation (59)**) and aggressive FCM seeding (**Equation (60)**) are retained unchanged from v2 and v3, so that FCM-guided initialisation remains a core component of the framework in every iteration. Lamarckian local search is enabled in FCM-PbGA-v4 ($L = 20$ iterations, perturbation scale $\sigma_{\text{ls}} = 0.03$ times each unit generation range, greedy single-gene acceptance per **Equations (14)** and **(15)**); v4 inherits these settings unchanged from Stage-1. The comparative outcome of all four stages is reported quantitatively in Section 4.11.5 below, where FCM-PbGA-v4 is carried forward as the final proposed method.

4.11.5. Iterative refinement: From FCM-PbGA-v2 to v4 (evaluation)

Section 4.11 introduces FCM-PbGA-v2, which significantly improves on Stage-1, but **Table 21** already shows that a plain single-population, budget-matched GA still posts a lower mean cost than v2. Rather than stop there, two further minimal design iterations were tested, each targeted at the specific diagnostic evidence available at that point. *FCM-PbGA-v3* keeps the v2 structure while adopting the operator hyperparameters diagnosed as responsible for the advantage of the budget-matched GA (full specification in Section 4.11). This narrows the gap to budget-matched GA from \$6,654.95 (Stage-1) to \$2,164.48 (v3)—a 67.5% reduction—but the residual gap remains statistically significant ($d = 1.53$, $p = 1.7 \times 10^{-7}$). *FCM-PbGA-v4* tests whether this residual gap originates in the adaptive island-merging schedule itself (**Equation (56)**) rather than in the operator hyperparameters already corrected in v3, motivated by the ablation finding in Section 4.10 (**Table 20**) that a single population outperforms six islands even before any operator changes. FCM-PbGA-v4 therefore removes the island mechanism entirely ($S^{(g)} = 1$ for every g; details in Section 4.11) while retaining v3's operators, POZ-aware crossover, and aggressive FCM seeding. **Table 22** summarises the resulting progression.

Table 23 evaluates FCM-PbGA-v4, the final proposed method, against all ten methods above—Stage-1 FCM-PbGA, and GA/PSO/GWO/WOA/Merit Order at both the vanilla (compact) and matched (equal-budget) settings—in a single 30-run comparison. FCM-PbGA-v4 uses the same ≈53,000-evaluation budget as Stage-1.

FCM-PbGA-v4 achieves a mean cost of \$655,073.62 (best run \$650,837.80), improving on Stage-1 FCM-PbGA by \$7,099.72 (1.07%) with a large effect size ($d = -3.85$, $p = 1.6 \times 10^{-21}$)—a substantially larger improvement than v2 or v3 achieved individually. FCM-PbGA-v4 remains significantly ahead of every vanilla-budget benchmark and of matched-budget PSO, GWO, and WOA, in every case with a large effect size, including GWO (matched), which had beaten both Stage-1 and v2 in earlier comparisons. Against budget-matched GA—the one comparison that neither Stage-1, v2, nor v3 could win—FCM-PbGA-v4 is numerically \$444.77 lower on average and its best run (\$650,837.80) is \$2,623.55 below budget-matched GA's best run (\$653,461.35); however, this difference does not reach statistical significance (t-test $p = 0.199$, Wilcoxon $p = 0.355$, Cohen's $d = -0.34$, small), so this comparison should be reported as a statistical tie rather than a win. Given the shared post-hoc repair pipeline used by every method evaluated in this paper (**Equations (14)–(46)**, Section 4.12), FCM-PbGA-v4's final adjusted schedule is expected to retain the zero power-balance and POZ violations already confirmed for Stage-1 (**Table 5**), though a dedicated constraint-satisfaction log for v4 specifically was not yet available at the time of writing and is reported as a verification item alongside the future work in Section 5. Taken as a whole, the three-step diagnostic-driven redesign—v2's structural novelties, v3's operator correction, and v4's removal of the island mechanism entirely—closes the full budget-confound gap identified in **Table 21**, arriving at a method that is statistically indistinguishable from the strongest budget-matched vanilla baseline while remaining significantly ahead of every other method tested, and while keeping FCM-guided seeding as a core, retained component throughout all three iterations.

Table 22. Progression of mean total cost and the gap to budget-matched GA across the four design iterations (30 runs each). Each iteration's design change is described in the text above and in Section 4.11.

Method	Mean Cost ($)	Gap to GA (matched)	Gap vs. Stage-1's Gap
Stage-1 FCM-PbGA (Full)	662,173.34	$6,654.95 (1.02%)	—
FCM-PbGA-v2	658,904.61	$3,386.22 (0.52%)	49.1% closed
FCM-PbGA-v3	657,682.87	$2,164.48 (0.33%)	67.5% closed
FCM-PbGA-v4	**655,073.62**	**−$444.77 (−0.07%)**	**107% closed (v4 numerically ahead)**
GA (matched, reference)	*655,518.39*	—	—

Table 23. Final comprehensive comparison: FCM-PbGA-v4 against Stage-1 FCM-PbGA and all benchmark methods at both vanilla and matched evaluation budgets (30 runs each). One-way ANOVA across all eleven methods: $p = 1.33 \times 10^{-226}$; Friedman test: $p = 7.40 \times 10^{-56}$ (both reject equal performance at $\alpha = 0.05$). Pairwise columns compare each method against FCM-PbGA-v4 (paired t-test, Cohen's d per **Equation (55)** with FCM-PbGA-v4 in the role of FCM; *negative* d favours FCM-PbGA-v4).

Method	Mean ($)	Std ($)	t-test p (vs. v4)	Cohen's d (vs. v4)	Sig.
FCM-PbGA-v4	**655,073.62**	1,469.06	—	—	—
Stage-1 FCM-PbGA (Full)	662,173.34	2,150.85	1.6×10^{-21}	−3.85	Yes
GA (vanilla)	663,692.91	1,733.43	1.5×10^{-28}	−5.36	Yes
PSO (vanilla)	676,324.83	2,562.30	1.4×10^{-43}	−10.18	Yes
GWO (vanilla)	666,045.56	2,238.93	2.8×10^{-30}	−5.79	Yes
WOA (vanilla)	674,320.01	2,239.87	1.6×10^{-43}	−10.16	Yes
Merit Order	688,819.61	0	2.1×10^{-72}	−32.49	Yes
GA (matched)	655,518.39	1,165.25	**0.199**	−0.34 **(small)**	**No**
PSO (matched)	672,791.16	1,798.52	5.4×10^{-45}	−10.79	Yes
GWO (matched)	660,576.00	1,544.03	1.9×10^{-20}	−3.65	Yes
WOA (matched)	668,373.34	1,874.98	1.8×10^{-37}	−7.90	Yes

The head-to-head numbers above close the diagnostic loop opened in Section 4.10: each redesign step is now traceable in cost (**Tables 22** and **23**), while the revision evidence—ablation, rematch, compliance, and generality checks—is consolidated in **Tables 24–28** and examined next.

Table 24. FCM ablation specific to FCM-PbGA-v4 (30 runs each, identical ≈53,000-evaluation budget and seeds; POZ sets correctly loaded). NoFCM disables FCM-guided seeding via the use_fcm_seed gate; FCMonly/Neither toggle POZ-aware crossover for the synergy test.

Variant	Mean ($)	Std ($)	Best ($)	Worst ($)	Median ($)
v4-Full (A)	655,844.85	1,168.84	653,793.82	657,598.73	656,014.66
v4-NoFCM (B)	656,882.43	1,882.18	653,908.63	662,989.82	656,819.35
FCMonly (C)	656,339.44	1,551.10	—	—	—
Neither (D)	656,827.52	1,661.65	—	—	—

A−B = −$1,037.58; paired t $p = 0.0113$; signed-rank $p = 0.0161$;
rank-sum $p = 0.0207$; ANOVA $p = 0.0129$ ($F = 6.58$); Friedman $p = 0.068$;
Cohen's $d = -0.66$ (medium); FCM×POZ interaction = −549.50.

Table 25. Lamarckian local-search ablation in v4 (30 runs each, all else identical)

Variant	Mean ($)	Std ($)	Best ($)	Worst ($)	Median ($)
v4 (LS enabled, $L = 20$)	655,844.85	1,168.84	653,793.82	657,598.73	656,014.66
v4-NoLS	656,445.55	1,690.16	653,231.43	659,978.28	656,420.83
Diff = −$600.70; paired t $p = 0.083$; signed-rank $p = 0.092$; rank-sum $p = 0.162$; ANOVA $p = 0.115$; Friedman $p = 0.273$; $d = -0.41$ (n.s.).					

Table 26. v4 versus budget-matched GA across loader/repair conditions (30 runs each). Black: previously reported. Rows marked Published give the submitted-manuscript numbers (old repair pipeline with empty POZ sets due to the loader bug); they are not directly comparable to revision rows.

Condition	v4 mean ($)	GA mean ($)	Gap ($)	Test
Published (zones empty)	655,073.62 ± 1,469	655,518.39 ± 1,165	−444.77	tie ($p = 0.199$, $d = -0.34$)
Rematch (POZ fix)	655,844.85 ± 1,169	657,199.90 ± 1,744	−1,355.05	$p = 1.5 \times 10^{-4}$, $d = -0.91$
Strict-final (POZ fix)	652,790.93 ± 969.95	654,200.33 ± 1,129.79	−1,409.41	$p = 5.2 \times 10^{-6}$, $d = -1.34$

Table 27. Constraint compliance: published v4 best run audited pre/post old repair (strict definitions), and strict-final schedules over 30 runs × 24 h × 15 units = 10,800 unit-hours per method

Constraint	Pre N / max (MW)	Post (old) N / max	Strict v4 (10,800 uh)	Strict GA (10,800 uh)
Power balance (shortfall)	0 / 0	0 / 0	0 / 0	0 / 0
Ramp-rate + startup	28 / 172.80	1 / 120.00	0 / 0	0 / 0
POZ interior	24 / 9.94	29 / 12.45	0 / 0	0 / 0
Generation limits	0 / 0	3 / 70.00	0 / 0	0 / 0

Table 28. Cost progression across design iterations and conditions (mean ± std, 30 runs). Black: previously reported. Rows marked (published) give the submitted-manuscript numbers (old repair pipeline with empty POZ sets due to the loader bug); they are not directly comparable to revision rows.

Method / condition	Mean ± std ($)
Stage-1 FCM-PbGA (published)	662,173.34 ± 2,150.85
FCM-PbGA-v4 (published)	655,073.62 ± 1,469.06
v4, POZ fix, old pipeline	655,844.85 ± 1,168.84
v4-strict (POZ fix + strict repair)	652,790.93 ± 969.95
Budget-matched GA (published)	655,518.39 ± 1,165.25
Matched GA, POZ fix, old pipeline	657,199.90 ± 1,744.42
Matched GA-strict	654,200.33 ± 1,129.79

Seeding that survives the budget test. When FCM-guided seeding is disabled with budget, seeds, and every other setting held constant, mean cost rises by $1,037.58 (**Table 24**), significant under four of five tests with a medium effect ($d = -0.66$); only the conservative Friedman test is marginal. The contrast with the Stage-1 ablation, where removing seeding changed nothing ($p = 0.498$), traces to dosage: v4 seeds 15% of every island around the centroid instead of a single individual. The same factorial design yields the synergy metric of **Equation (49)**: $I = -549.50$

shows the two mechanisms jointly deliver more than the sum of their separate gains, consistent with centroid seeding placing BLX offspring inside Ω_i^t where projection rarely snaps coordinates to zone edges. That cost-based reading matters because stored-coordinate violation counts are zero in all arms, and because the gap dissected here sets up the head-to-head rematch below.

A polish worth keeping. The $L = 20$ greedy polish (**Table 25**) improves the mean by \$600.70 at a small-to-medium effect ($d = -0.41$) without reaching significance on any test, mirroring the Stage-1 NoLS outcome at larger magnitude. It stays in the final method as a computationally cheap exploitation step, now documented in Table 1, while the decisive comparison against the strongest baseline follows.

From statistical tie to significant lead. **Table 26** reconciles the two readings of the v4-versus-GA contest: the published tie ($p = 0.199$) was measured while the loader bug left all zones empty and the POZ-aware crossover dormant; with zones correctly loaded the rematch is significant ($p = 1.5 \times 10^{-4}$, $d = -0.91$), because matched GA degrades by \$1,681.51 against v4's \$771.23 shift. The strict-final mapping widens the lead ($p = 5.2 \times 10^{-6}$, $d = -1.34$) while buying feasibility, as the compliance accounts show next. GA's better best-run in the old pipeline is reported descriptively, without significance claims.

Counting violations, honestly. **Table 27** keeps the published best run on record—28 ramp/startup and 24 POZ violations pre-repair, with the old repair leaving 1 ramp, 29 POZ, and 3 gen-limit violations—while power balance never breaks. The strict mapping of **Equations (50)–(52)**, anchored by the projection-based balance restoration in **Equation (51)**, then drives all four classes to zero across the full 10,800 unit-hours for both methods, enabled by the two documented corrections (string-forced zone loading; startup bound $[\, P_i^{\min}, \min(P_i^{\max}, P_i^{\min}+R_i^U)\,]$). Feasibility established, the remaining question is generality, tested next.

The full arc. **Table 28** reads the redesign as a monotone descent from Stage-1 (\$662,173.34) to v4-strict (\$652,790.93), with strictness itself paying economically for both methods. Cost and feasibility improve together rather than trading off, which reframes the operational discussion that follows.

4.12. Discussion: Constraint handling as the central practical contribution

The diagnostics in Section 4.10 establish that, once the evaluation budget is equalised, the cost ranking among evolutionary methods is close: earlier iterations (Stage-1, v2, v3) trailed budget-matched GA significantly, and even the final method, FCM-PbGA-v4, only reaches a statistical tie rather than a decisive win (Section 4.11.5). This raises a natural question for a power engineering audience: if the proposed framework does not unconditionally minimise cost, what is its practical contribution? This section argues that the answer lies not in the cost ranking itself but in *how* feasibility is achieved, and connects this distinction to concrete operational consequences in thermal power system dispatch.

4.12.1. Feasibility is shared among repair-equipped methods, but not with existing practice

Every evolutionary method evaluated in this paper—Stage-1 FCM-PbGA, FCM-PbGA-v2 through v4, and GA, PSO, GWO, and WOA at both the vanilla and matched budgets—is coupled to the same external repair pipeline (ramp-feasible clipping via **Equations (14)–(15)**, POZ projection via **Equation (17)**, and headroom-proportional demand restoration via **Equations (45)–(46)**) before its hourly schedule is finalised. Consequently, the zero-violation outcome directly confirmed for Stage-1 through v3 (and expected, though not yet separately verified, for the final method v4; Section 4.11.5) is *not* unique among the evolutionary methods compared: every metaheuristic in **Tables 20–23**, including the budget-matched GA that remains statistically tied with v4 on cost, likewise produces a fully feasible final schedule once the same repair is applied. The meaningful feasibility contrast in this paper is therefore not FCM-PbGA against other metaheuristics, but the entire repair-equipped evolutionary framework against the Merit Order baseline that represents existing dispatch practice (Section 4.4, **Table 6**): Merit Order's raw schedule carries more than three times as many generation-limit violations (11 versus 3) and three times as many ramp-rate violations (3 versus 1) as any evolutionary method's raw output, and—critically—Merit Order as practiced has no analogue of **Equations (14)–(46)** to correct them; the correction step exists in this study only as a common simulation control applied uniformly for fair comparison among optimisation methods, not as a component of Merit Order dispatch itself. What FCM-PbGA contributes, across all four stages, beyond matching this shared feasibility floor is that constraint awareness is embedded inside the search rather than left entirely to post-

hoc correction: the commitment-consistent chromosome encoding (Section 3) never represents an offline unit's output as a free variable, and POZ-aware crossover (**Equation (59)**, retained in v2, v3, and the final method v4) keeps the population's search coordinates inside Ω_i^t throughout the run rather than relying on the final projection alone. Whether this distinction affects solution *quality* is answered empirically by **Table 23**; whether it affects *feasibility*, it does not, because feasibility is already guaranteed for every repair-equipped competitor.

4.12.2. Operational significance of the constraints enforced

The three constraint classes handled throughout this paper are not abstract modelling conveniences; each corresponds to a specific physical failure mode in thermal power plants, and violating any of them has consequences that a percentage-point cost comparison does not capture.

Ramp-rate violations (**Equation (5)**) correspond to commanding a boiler-turbine unit to change output faster than its thermal and mechanical inertia allow. Repeated or severe ramp-rate excursions accelerate fatigue in high-temperature components (boiler tubes, turbine rotors, header welds) through thermal cycling stress, shortening component life and increasing unplanned outage risk; this is precisely why ramp-rate limits are treated as hard operating constraints—not soft preferences—by transmission system operators and reliability coordinators[60] See also the Policy and Regulatory Implications discussion following the Conclusion, which examines this point from a grid-reliability compliance perspective.. A dispatch method that only satisfies ramp limits after the fact, via an external correction layer with no visibility into why the uncorrected trajectory was infeasible, offers no assurance that the corrected schedule remains economically sensible; embedding ramp-feasible bounds directly into the optimisation variables (**Equations (14)–(15)**), as every FCM-PbGA stage does, ensures the search never wastes evaluation budget on trajectories that would later require heavy correction.

Prohibited-operating-zone violations (**Equation (4)**) correspond to sustained operation within a generator's mechanical resonance band, where shaft vibration amplitudes rise sharply and can contribute to bearing and coupling wear.[12,61] Unlike ramp-rate limits, POZ constraints are unit-specific and can fragment a single unit's otherwise-continuous feasible range into several disjoint intervals (**Equation (16)**), which is exactly the source of the non-convexity this paper's chromosome encoding and POZ-aware projection (Section 3.10) are designed to handle without resorting to a large penalty coefficient that would need re-tuning for every unit's resonance geometry.

Power-balance violations (**Equation (2)**) are the most immediately consequential of the three: a sustained generation shortfall manifests as system frequency decline, and if uncorrected can trigger automatic under-frequency load shedding or, in the worst case, cascading protection trips. All methods evaluated in this paper enforce power balance in the strict-inequality sense (generation at least meeting demand, **Equation (46)**), and the resulting surplus is small (below 2% in all but three transition hours, **Table 3**)—but the more relevant point for this discussion is that Merit Order, evaluated without the repair pipeline that every evolutionary method in this study receives for fair comparison, provides no structural guarantee that a real-world implementation of the same priority-order logic would achieve this without manual operator intervention at each hour.

4.12.3. Reframing the cost comparison in operational context

Given the discussion above, the fact that FCM-PbGA-v4's \$444.77 (0.07%) numerical advantage over budget-matched GA (**Table 23**) does not reach statistical significance ($p = 0.199$) should be read as a secondary finding relative to the primary, practice-facing result of this paper: replacing Merit Order with *any* of the evolutionary, repair-equipped methods evaluated here delivers both a substantial, statistically robust cost reduction (Section 4.6) and—unlike Merit Order as practiced—a structural mechanism for honouring ramp-rate, POZ, and power-balance constraints without per-hour manual correction. For a system operator deciding whether to adopt a computational-intelligence dispatch tool, the choice of which specific evolutionary method to deploy is a secondary optimisation over an already large expected gain; the choice of whether to adopt structural feasibility handling at all is the more consequential decision, and it is this second choice that the Merit Order comparison in Sections 4.2–4.7 and **Table 6** most directly informs. This reframing does not diminish the value of the iterative redesign documented in Section 4.11.5, which took FCM-PbGA from a significant loss against budget-matched GA (Stage-1, $d = 3.85$) to a statistical tie (v4, $d = -0.34$) while preserving FCM-guided seeding throughout; establishing

a statistically significant advantage, not merely parity, over budget-matched GA remains identified as future work (Section 5).

5. Conclusion and future work

This paper presented FCM-PbGA, a computational intelligence framework that combines fuzzy clustering, evolutionary search, and constraint-projection repair into a single feasibility-aware architecture, and demonstrated it on non-convex economic dispatch with unit commitment, valve-point effects, prohibited operating zones, and ramp-rate constraints—a representative electrical power engineering problem for which classical numerical solvers are unsuitable. Unlike conventional penalty-based approaches, operational constraints are embedded directly into the chromosome representation, initialization, and repair mechanisms. On the IEEE 15-unit 24-hour benchmark, the initial island-based architecture (Stage 1) achieved a mean operating cost of \$662,173 ± \$2,151 over 30 runs, with zero ramp and POZ violations, and outperformed vanilla GA, GWO, WOA, PSO, and the Merit Order baseline, with the benchmark methods evaluated under standard-practice (compact) settings. The best solution, with a cost of \$657,914, consisted of a base fuel cost of \$646,377 and a valve-point effect (VPE) cost of \$11,537, 32.2% lower than the Merit Order baseline, and Stage 1 reduced operating cost relative to Merit Order by \$30,905 (4.49%) in the best case and \$26,646 (3.87%) on average. Because Stage 1 also carries a ∼10–11× larger evaluation budget than these compact baselines, Section 4.10 subjected the architecture to a component-wise ablation and a budget-matched re-run of every benchmark. Both diagnostics agreed that Stage 1's fixed six-island partitioning is counterproductive once the budget is equalised, and that budget-matched vanilla GA significantly outperforms Stage 1. This finding motivated three successive, minimal redesign iterations (Section 4.11). FCM-PbGA-v2 replaced fixed island partitioning with adaptive island merging, geometric mutation decay with stagnation-triggered reheating, objective-only POZ correction with POZ-aware crossover, and single-point FCM seeding with fractional population seeding, closing 49% of the gap to budget-matched GA. FCM-PbGA-v3 additionally adopted budget-matched GA's operator hyperparameters, closing a further 36% of the remaining gap. FCM-PbGA-v4 removed the island mechanism entirely while retaining POZ-aware crossover, fractional FCM seeding, and v3's operators, arriving at a mean cost of \$655,073.62 (best run \$650,837.80) that is no longer statistically distinguishable from budget-matched GA ($p = 0.199$, Cohen's $d = -0.34$) while remaining significantly ahead of Stage 1 ($d = -3.85$, $p = 1.6 \times 10^{-21}$) and every other benchmark tested (Section 4.11.5). FCM-guided seeding was retained as a core architectural component throughout all three redesign iterations, consistent with the project's original research proposal. Sensitivity analysis of the Stage-1 architecture showed that population size had the greatest influence on performance, followed by generation depth, island diversity, and cluster granularity, and the complete 24-hour schedule was obtained in under 14 s of wall-clock computation across all four stages, confirming that the proposed computational architecture is practical, not just accurate. More broadly, because FCM-based initialisation and projection-based repair are formulated independently of the economic-dispatch cost function, the framework—together with the diagnose-then-redesign methodology used to refine it across four iterations—offers a template for applying and improving the same computational intelligence architecture on other constrained scheduling and resource allocation problems in electrical and computer engineering.

Policy and regulatory implications

Beyond its algorithmic contribution, the results carry practical implications for the system operators and regulators who oversee day-to-day thermal dispatch. First, the documented savings are not merely an average-case improvement: the *worst-case* FCM-PbGA run across 30 trials still beat the Merit Order baseline by \$21,665 (3.15%) per day (**Table 9**), and the 95% confidence interval for the mean daily saving reported in Section 4.7 corresponds to an estimated \$9.4–10.0 million in annual savings. For regulated utilities under cost-of-service or performance-based ratemaking,[62] a guaranteed lower bound—rather than an average claim—is the more defensible basis for utilities to justify capital or software expenditure to a regulator, and for regulators to evaluate whether such expenditure should be recovered through tariffs. Second, the raw-violation comparison in **Table 6** and the constraint-satisfaction summary in **Table 5** show that the existing Merit Order practice can generate ramp-rate and generation-limit violations that are, by construction, never corrected within the dispatch procedure itself and must instead be caught downstream by human operators or protection systems. This is directly relevant to grid-reliability compliance regimes (e.g.

mandatory ramp-rate and reserve-margin standards enforced by transmission system operators and reliability coordinators):[60] an optimisation-embedded feasibility guarantee, of the kind FCM-PbGA provides through **Equations (14)–(17)**, reduces reliance on after-the-fact manual correction and could inform procurement specifications or certification criteria that regulators or grid codes place on dispatch-optimisation software before it is approved for operational use. Third, because the complete 24-hour schedule is obtained in under 14 s of wall-clock computation (Section 4.5.1), the approach is computationally compatible with existing day-ahead and rolling-horizon scheduling cycles, lowering one practical barrier (computational cost) that utilities and regulators often cite when evaluating whether to modernise legacy dispatch software. Finally, because FCM-PbGA—like any stochastic metaheuristic—produces run-to-run variability that a deterministic method such as Merit Order does not, its adoption in a regulated setting would benefit from an agreed statistical validation protocol[63] (analogous to the 30-run, multi-test procedure used in Section 4.7) that vendors and operators could use to demonstrate a guaranteed worst-case performance bound prior to deployment, rather than relying on best-case or average-case claims alone. These implications are offered as directions for policy and procurement discussion; they are grounded in a single 15-unit benchmark and are not a substitute for a dedicated market-impact or regulatory-cost study, which we identify as future work below.

Limitations

- FCM-PbGA-v4's parity with budget-matched GA is a statistical tie, not a significant win. Section 4.11.5 shows FCM-PbGA-v4 is numerically \$444.77 (0.07%) below budget-matched GA on average and \$2,623.55 below on the best run, but this difference does not reach statistical significance ($p = 0.199$, Cohen's $d = -0.34$, small). The correct reading is that the three-step redesign closes the budget-confound gap rather than reversing it decisively; a larger sample of runs, or a further targeted design change, would be needed to establish a significant advantage over the single strongest budget-matched baseline.
- A dedicated constraint-satisfaction log for FCM-PbGA-v4 specifically was not available at the time of writing. Zero power-balance and POZ violations are directly confirmed for Stage 1 through v3 (**Table 5** and analogous checks); v4 is expected to inherit the same result because every method in this study, including v4, is coupled to the identical post-hoc repair pipeline (Section 4.12), but this has not yet been verified empirically for v4's specific final schedules and is reported as an open verification item below.
- The ablation and budget-matched diagnostics in Section 4.10 were performed for the Stage-1 architecture and for each redesign iteration (v2, v3, v4) as a complete architecture; a further component-wise ablation of v4's own two retained mechanisms (POZ-aware crossover and fractional FCM seeding) individually, isolated from the operator-hyperparameter and island-removal changes, was not performed, so their relative contributions to the final Stage-1-to-v4 improvement are not yet separately quantified.
- The study is restricted to a single 15-unit benchmark with a synthetic load profile; generalisability to larger systems is not established.
- Minimum up/down times, start-up and shut-down costs, and reserve requirements are not modelled; these omissions limit the realism of the UC representation.
- Network constraints (line flows, voltage limits) and renewable generation are absent.
- The claim of superiority rests on one benchmark configuration; performance on larger, differently structured systems may differ.

Future work

Future directions include: (i) extension to 100+ unit systems with transmission network and spinning-reserve constraints; (ii) hybrid renewable–thermal scheduling under forecast uncertainty using chance-constrained formulations; (iii) multi-objective cost–emission optimisation; (iv) minimum up/down time and start-up cost modelling; (v) rolling-horizon real-time implementation with warm-start transfer learning across successive dispatch intervals; (vi) integration with smart-grid and microgrid energy management systems;[4] and (vii) porting the FCM-guided, repair-based computational architecture itself to other constrained combinatorial and mixed discrete-continuous optimisation problems in electrical and computer engineering, such as communication-network resource allocation and signal-processing hardware/software co-design, to assess how much of its performance advantage is problem-specific versus architectural; and (viii) a dedicated market-impact and regulatory-cost

study, conducted with system operators or regulators on production-scale data, to quantify the policy implications outlined above beyond the single-benchmark evidence presented in this paper; and (ix) generating and verifying the constraint-satisfaction log for FCM-PbGA-v4's final adjusted schedules specifically, and running a larger-sample (e.g. 100-run) comparison against budget-matched GA to determine whether v4's numerical advantage becomes statistically significant with additional statistical power, or whether a further targeted design change is required to move beyond the current statistical tie identified in Section 4.11.5.

Acknowledgments

The authors gratefully acknowledge the support provided by the Directorate of Research and Community Service at Telkom University.

Fundings

This research was funded by the Directorate of Research and Community Service at Telkom University under Agreement/Contract Letter No. 009/LIT06/PPM-LIT/2025.

Conflict of interest

The authors declare that they have no competing interests.

Author contributions

Conceptualization:Rifki Rahman Nur Ikhsan, Jangkung Raharjo
Data curation: Rifki Rahman Nur Ikhsan, Jangkung Raharjo
Formal analysis: Rifki Rahman Nur Ikhsan
Funding acquisition: Efri Suhartono
Investigation: Rifki Rahman Nur Ikhsan
Methodology: Rifki Rahman Nur Ikhsan, Efri Suhartono
Project administration: Lindiasari Martha Yustika, Gede Putu Oka Indra Wijaya
Software: Rifki Rahman Nur Ikhsan
Supervision: Jangkung Raharjo
Validation: Rifki Rahman Nur Ikhsan, Efri Suhartono, Lindiasari Martha Yustika
Writing–original draft: Rifki Rahman Nur Ikhsan, Jangkung Raharjo
Writing–review & editing: All authors

Availability of data

The benchmark test systems used in this study were obtained from the Economic Load Dispatch Test Systems Repository.[64]

AI tools statement

All authors confirm that no AI tools were used in the preparation of this manuscript.

References

1. Wang H, Xu Y, Yi Z, Xu J, Xie Y, Li Z. A review on economic dispatch of power system considering atmospheric pollutant emissions. *Energies*. 2024;17(8):1878. https://doi.org/10.3390/en17081878
2. Yang Z, Yong P, Xiang M. Revisit power system dispatch: Concepts, models, and solutions. *iEnergy*. 2023;2(1):43-62. https://doi.org/10.23919/IEN.2023.0010
3. Kunya AB, Abubakar AS, Yusuf SS. Review of economic dispatch in multi-area power system: State-of-the-art and future prospective. *Electric Power Syst Res*. 2023;217:109089. https://doi.org/10.1016/j.epsr.2022.109089
4. Rizki A, Touil A, Echchatbi A, Oucheikh R. A reinforcement learning approach based on group relative policy optimization for economic dispatch in smart grids. *Electricity*. 2025;6(3):49. https://doi.org/10.3390/electricity6030049
5. Luo Q, Garcia-Menendez F, Yang H, Deshmukh R, He G, Lin J, *et al*. The health and climate benefits of economic dispatch in China's power system. *Environ Sci Technol*. 2023;57(7):2898-2906. https://doi.org/10.1021/acs.est.2c05663
6. Wu Y, Shi Z, Lin Z, Zhao X, Xue T, Shao J. Low-carbon economic dispatch for integrated energy system through the dynamic reward and penalty carbon emission pricing mechanism. *Front Energy Res*. 2022;10:843993. https://doi.org/10.3389/fenrg.2022.843993
7. Hernández K, Barrera-Singaña C, Tipán L. Hydrothermal economic dispatch incorporating the valve point effect in thermal units solved by heuristic techniques. *Energies*. 2025;18(11):2789. https://doi.org/10.3390/en18112789
8. Ramli NF, Kamari NAM, Halim SA, Zulkifley MA, Sahri MSM, Musirin I. A non-convex economic dispatch problem with point-valve effect using a wind-driven optimisation approach. *J Electr Eng Technol*. 2022;17(1):85-95. https://doi.org/10.1007/s42835-021-00859-6
9. Van Hoorebeeck L, Absil PA, Papavasiliou A. Solving non-convex economic dispatch with valve-point effects and losses with guaranteed accuracy. *Int J Electr Power Energy Syst*. 2022;134:107143. https://doi.org/10.1016/j.ijepes.2021.107143
10. Abbas G, Khan IA, Ashraf N, Raza MT, Rashad M, Muzzammel R. On employing a constrained nonlinear optimizer to constrained economic dispatch problems. *Sustainability*. 2023;15(13):9924. https://doi.org/10.3390/SU15139924
11. Zein H, Raharjo J, Mardiyanto IR. A method for completing economic load dispatch using the technique of narrowing down area. *IEEE Access*.

2022;10:30822-30831.
https://doi.org/10.1109/ACCESS.2022.3158928

12. Torchani A, Boudjemline A, Gasmi H, Bouazzi Y, Guesmi T. Dynamic economic/environmental dispatch problem considering prohibited operating zones. *Eng Technol Appl Sci Res.* 2019;9(5):4586-4590.
https://doi.org/10.48084/etasr.2904
13. Manieri L, Falsone A, Prandini M. A dual bisection approach to economic dispatch of generators with prohibited operating zones. *IEEE Control Syst Lett.* 2024;8:1547-1552.
https://doi.org/10.1109/LCSYS.2024.3409230
14. Pan S, Jian J, Yang L. Solution to dynamic economic dispatch with prohibited operating zones via MILP. *Math Biosci Eng.* 2022;19(7):6455-6468.
https://doi.org/10.3934/mbe.2022303
15. Martins D, Mendes M. A hybrid evolutionary approach applied to the economic dispatch problem with prohibited operating zones and uncertainties. *IEEE Lat Am Trans.* 2021;19(7):1225-1232.
https://doi.org/10.1109/TLA.2021.9461852
16. Raharjo J, Ikhsan RRN, Yustika LM. A greedy adaptive and backtracking framework for reducing emission costs in generator scheduling. *Int J Technol.* 2025;16(6):1888-1893.
https://doi.org/10.14716/ijtech.v16i6.7588
17. Ikhsan RRN, Raharjo J, Rahmat B. Vectorized lambda iteration method for swift economic dispatch analysis. *Evergreen.* 2024;11(1):435-447.
https://doi.org/10.5109/7172306
18. Spea SR. Cost-effective economic dispatch in large-scale power systems using enhanced manta ray foraging optimization. *Neural Comput Appl.* 2025;37(18):12487-12524.
https://doi.org/10.1007/s00521-025-11086-9
19. Asif M, Amin A, Jamil U, Mahmood A, Ahmed U, Razzaq S, *et al.* Combined emission economic dispatch using quantum-inspired particle swarm optimization and its variants. *Energy Explor Exploit.* 2024;42(5):1602-1644.
https://doi.org/10.1177/01445987241235419
20. Kamboj VK, Kumari CL, Bath SK, Prashar D, Rashid M, Alshamrani SS, *et al.* A cost-effective solution for non-convex economic load dispatch problems in power systems using slime mould algorithm. *Sustainability.* 2022;14(5):2586.
https://doi.org/10.3390/su14052586
21. Secui DC, Secui ML. Social small group optimization algorithm for large-scale economic dispatch problem with valve-point effects and multi-fuel sources. *Appl Intell.* 2024;54(17):8296-8346.
https://doi.org/10.1007/S10489-024-05517-8
22. Basetti V, Rangarajan SS, Shiva CK, Pulluri H, Kumar R, Collins RE, *et al.* Economic emission load dispatch problem with valve-point loading using a novel quasi-oppositional-based political optimizer. *Electronics.* 2021;10(21):2596.
https://doi.org/10.3390/electronics10212596
23. Fahim KE, De Silva LC, Andiappan V, Shezan SA, Yassin H. A novel hybrid algorithm for solving economic load dispatch in power systems. *Int J Energy Res.* 2024;2024(1):8420107.
https://doi.org/10.1155/2024/8420107
24. Madhale K, Jamadar N, Jadhav HT. Economic emission dispatch of power plants considering ramp rate, prohibited zones, and spinning reserve using a hybrid MOABC-NSGA-II algorithm. *Life Cycle Reliab Saf Eng.* 2025;14(4):645-663.
https://doi.org/10.1007/s41872-025-00339-0
25. Katre R, Khadse C, Bhadade R. Metaheuristic optimization and machine learning-based unit commitment strategies in smart grids. *Eng Technol Appl Sci Res.* 2025;15(5):27293-27299.
https://doi.org/10.48084/etasr.12853
26. Widayanti L, Afandi AN, Herwanto HW, Fitria VA. Metaheuristic optimization in dynamic unit commitment of power systems: Current trends and future prospects. *JOIV Int J Inform Vis.* 2024;8(4):2261-2272.
https://doi.org/10.62527/joiv.8.4.2423
27. Kaewpasuk S, Intiyot B, Jeenanunta C. A fuzzy unit commitment model for enhancing stability and sustainability in renewable energy-integrated power systems. *Sustainability.* 2025;17(15):6800.
https://doi.org/10.3390/su17156800
28. Chen Y, Xiao C, Yang S, Yang Y, Wang W. Research on long term power load grey combination forecasting based on fuzzy support vector machine. *Comput Electr Eng.* 2024;116:109205.
https://doi.org/10.1016/j.compeleceng.2024.109205
29. Senen A, Putri TWO, Jamian JJ, Supriyanto E, Anggaini D. Fuzzy C-means clustering based on micro-spatial analysis for electricity load profile characterization. *Indones J Electr Eng Comput Sci.* 2023;30(1):33-45.
https://doi.org/10.11591/ijeecs.v30.i1.pp33-45
30. He W, Tang X, Ji W, Meng L, Wei J, Cao D, *et al.* An improved multi-island genetic algorithm and its utilization in the optimal design of a micropositioning stage. *Expert Syst Appl.* 2024;257:125029.
https://doi.org/10.1016/j.eswa.2024.125029
31. da Silveira LA, de Lima TA, Barros JB, Soncco-Álvarez JL, Llanos CH, Ayala-Rincón M. On the behavior of parallel island models. *Appl Soft Comput.* 2023;148:110880.
https://doi.org/10.1016/j.asoc.2023.110880
32. Dedić A, Konjić T, Ćalasan M, Dedić Z. Fuzzy C-means clustering applied to load profiling of industrial customers. *Electric Power Compon Syst.* 2021;49(11-12):1068-1084.
https://doi.org/10.1080/15325008.2022.2049660
33. Soni R, Mehta B. Diagnosis and prognosis of incipient faults and insulation status for asset management of power transformer using fuzzy logic controller and fuzzy clustering means. *Electric Power Syst Res.* 2023;220:109256.
https://doi.org/10.1016/j.epsr.2023.109256

34. Silva Filho TM, Pimentel BA, Souza RMCR, Oliveira ALI. Hybrid methods for fuzzy clustering based on fuzzy c-means and improved particle swarm optimization. *Expert Syst Appl.* 2015;42(17-18):7357-7369. https://doi.org/10.1016/j.eswa.2015.04.032
35. Zhang J, *et al.* An improved fuzzy c-means clustering algorithm based on PSO and shadowed sets. *Comput Intell Neurosci.* 2014;2014:1-10. https://doi.org/10.1155/2014/368628
36. Niu Q, Huang X. An enhanced fuzzy c-means clustering algorithm based on PSO. *J Softw.* 2011;6(5). https://doi.org/10.4304/jsw.6.5.873-879
37. Kuo RJ, *et al.* A hybrid metaheuristic and kernel intuitionistic fuzzy c-means algorithm. *Appl Soft Comput.* 2018;67:299-308. https://doi.org/10.1016/j.asoc.2018.02.039
38. Ezugwu AE, *et al.* Nature-inspired metaheuristic techniques for automatic clustering: a survey. *SN Appl Sci.* 2020;2(2). https://doi.org/10.1007/s42452-020-2073-0
39. Duan P, *et al.* Short-term load forecasting for electric power systems using PSO-SVR and FCM clustering techniques. *Energies.* 2011;4(1):173-184. https://doi.org/10.3390/en4010173
40. Ahmad FA, *et al.* Short-term load forecasting utilizing a combination model with PSO-based kernel fuzzy c-means clustering. *Int J Technol.* 2018.
41. Chan KY, *et al.* Fuzzy clustering-based deep learning for short-term load forecasting. *Sensors.* 2024;24(5):1391. https://doi.org/10.3390/s24051391
42. Niu S, *et al.* Power short-term load forecasting based on fuzzy c-means clustering and improved locally weighted linear regression. *Trans Inst Meas Control.* 2025;47(2):278-290. https://doi.org/10.1177/01423312241239229
43. Lin J, Magnago F, Alemany JM. Optimization methods applied to power systems: Current practices and challenges. *Classical and Recent Aspects of Power System Optimization.* 2018:1-18. https://doi.org/10.1016/B978-0-12-812441-3.00001-X
44. Marzbani F, Abdelfatah A. Economic dispatch optimization strategies and problem formulation: A comprehensive review. *Energies.* 2024;17(3):550. https://doi.org/10.3390/en17030550
45. Lotfi H. Multi-area economic dispatch under renewable integration: Optimization challenges and research perspectives. *Processes.* 2025;13(12):3766. https://doi.org/10.3390/pr13123766
46. Habib S, Kamarposhti MA, Shokouhandeh H, Colak I, Barhoumi EM. Economic dispatch optimization considering operation cost and environmental constraints using the HBMO method. *Energy Rep.* 2023;10(3):1718-1725. https://doi.org/10.1016/j.egyr.2023.08.032
47. Alkoffash MS, Awadallah MA, Alweshah M, Zitar RA, Assaleh K, Al-Betar MA. A non-convex economic load dispatch using hybrid salp swarm algorithm. *Arab J Sci Eng.* 2021;46(9):8721-8740. https://doi.org/10.1007/s13369-021-05646-z
48. Wood AJ, Wollenberg BF, Sheblé GB. *Power Generation, Operation, and Control.* 3rd ed. Hoboken, NJ: Wiley; 2013.
49. Rastgou A. Enhancing unit commitment in power systems: A piecewise-linear approach for optimal scheduling and environmental sustainability. *Results Eng.* 2025;28(1):107387. https://doi.org/10.1016/j.rineng.2025.107387
50. Sharma H, Galván E, Mooney P. Multi-Path island-based genetic algorithm for the K-most diverse near-shortest paths. *Inf Sci.* 2025;719:122495. https://doi.org/10.1016/j.ins.2025.122495
51. Hussain A, Riaz S, Amjad MS, Haq EU. Genetic algorithm with a new round-robin based tournament selection: Statistical properties analysis. *PLoS One.* 2022;17(9):e0274456. https://doi.org/10.1371/journal.pone.0274456
52. Alkafaween E, Hassanat A, Hassanat A, Essa E, Elmougy S. CSGA: A dual population genetic algorithm based on Mexican cavefish genetic diversity. *Jordan J Comput Inf Technol.* 2024;1. https://doi.org/10.5455/jjcit.71-1707664207
53. Santos GO, Vieira LS, Rossetti G, Ferreira CH, Moreira GJ. A high-performance evolutionary multiobjective community detection algorithm. *Soc Netw Anal Min.* 2025;15(1):110. https://doi.org/10.1007/s13278-025-01519-7
54. Zhang C, Wang L, He K. Cloud service composition optimization based on service association impact and improved NSGA-II algorithm. *Sci Rep.* 2025;15(1):26001. https://doi.org/10.1038/s41598-025-10040-y
55. Xing W, Zhang J, Zou Q, Lin J. Application of Gauss mutation genetic algorithm to optimize neural network in image painting art teaching. *Comput Intell Neurosci.* 2021;2021(1):3302617. https://doi.org/10.1155/2021/3302617
56. Jamshidi M, Alirezanejad M, Motameni H, Enayatifar R. A parallel hybrid genetic search for solving the capacitated vehicle routing problem. *Int J Comput Intell Syst.* 2025;18(1):306. https://doi.org/10.1007/s44196-025-01059-0
57. Zhang X, Gong Y, Cao Z, Zhang J. Island-based evolutionary computation with diverse surrogates and adaptive knowledge transfer for high-dimensional data-driven optimization. *ACM Trans Evol Learn Optim.* 2025;5(4):31. https://doi.org/10.1145/3700886
58. Suganthan PN, Kumar A. Problem definitions and evaluation criteria for the CEC 2020 special session and competition on single objective bound constrained numerical optimization. Technical Report. Singapore: Nanyang Technological

University; 2019. Accessed September 20, 2026. https://github.com/P-N-Suganthan/2020-Bound-Constrained-Opt-Benchmark
59. Wu F, Zhang J, Li S, Lv D, Li M. An enhanced differential evolution algorithm with Bernstein operator and refracted oppositional-mutual learning strategy. *Entropy*. 2022;24(9):1205. https://doi.org/10.3390/e24091205
60. North American Electric Reliability Corporation. *Reliability Standards for the Bulk Electric Systems of North America*. Atlanta, GA: NERC; 2021. Accessed September 20, 2026. https://www.nerc.com/pa/Stand/Reliability%20Standards%20Complete%20Set/RSCompleteSet.pdf
61. Alshammari BM. Dynamic environmental/economic power dispatch with prohibited zones using improved multi-objective PSO algorithm. *Int Rev Electr Eng*. 2016;11(4):411-419. https://doi.org/10.15866/iree.v11i4.8976
62. Comnes GA, Stoft S, Greene N, Hill LJ. *Performance-Based Ratemaking for Electric Utilities: Review of Plans and Analysis of Economic and Resource-Planning Issues, Volume I*. Berkeley, CA: Lawrence Berkeley National Laboratory; 1995. Accessed September 20, 2026. https://eta-publications.lbl.gov/sites/default/files/report-lbnl-37577.pdf
63. Stiasny J, Chevalier S, Nellikkath R, Sævarsson B, Chatzivasileiadis S. Closing the loop: A framework for trustworthy machine learning in power systems. In: *11th Bulk Power Systems Dynamics and Control Symposium (IREP)*. 2022. Accessed September 20, 2026. https://arxiv.org/abs/2203.07505
64. Al-Roomi AR. Economic load dispatch test systems repository. 2016. Accessed September 20, 2026. https://www.al-roomi.org/economic-dispatch

Rifki Rahman Nur Ikhsan *was born and raised in Bogor, Indonesia. He received his Bachelor's degree in electrical engineering (electronics) in 2022 and his Master's degree in electrical engineering, specializing in sustainable energy systems, in 2023, both from Telkom University, Bandung, Indonesia. He is currently a lecturer in the Energy Systems Engineering undergraduate programme at Telkom University and serves as head of external affairs at the university's SDGs Center. His research interests include power system optimization, economic dispatch, and generator scheduling, combining artificial-intelligence techniques such as machine learning and neural networks with calculus-based analytical methods. He has published journal and conference papers on power system operation and co-authored the book Regulasi dan Keekonomian Energi (Regulation and Economics of Energy, 2024).*
https://orcid.org/0009-0002-7873-0049

Efri Suhartono *received the Insinyur (Ir.) degree from Universitas Brawijaya, Malang, Indonesia, in 1998, and the Magister Teknik (M.T.) degree from Bandung Institute of Technology (ITB), Bandung, Indonesia, in 2004. He has been a lecturer at Telkom University, Bandung, Indonesia, since 2000, and he is currently pursuing his doctoral studies in electrical engineering at Telkom University. His research interests include power systems, renewable energy, power quality, and information systems. His published work in journals and IEEE conference proceedings covers optimal energy management, the placement and sizing of distributed generation, and the techno-economic analysis of solar PV. He also authored the textbook Pengantar Teknik Sistem Energi (Introduction to Energy Systems Engineering) and is a certified renewable-energy competency assessor.*
https://orcid.org/0009-0005-7958-2449

Jangkung Raharjo *was born in Sukoharjo, Central Java, Indonesia, on January 19, 1966. He obtained his bachelor's degree in electrical engineering from Bandung Institute of Technology (ITB), Indonesia, in 1991. He obtained his master's degree, with a concentration in electrical information systems, from ITB in 2001. He obtained his doctoral degree in the field of power systems at the Electrical Engineering Department, Institut Teknologi Sepuluh Nopember (ITS), Surabaya, Indonesia, in 2018. Since 1991, he has been a lecturer at Telkom University and received a professorship in smart grid and energy on June 1, 2023. Several papers related to power system optimization and renewable energy have been published. Currently, he is the head of the Centre of Excellence of Sustainable Energy and Climate Change, Telkom University. His research interests are energy, power systems, and electrical engineering.*
https://orcid.org/0000-0001-7082-5583

Lindiasari Martha Yustika *was born in Malang, Indonesia, in 1996. She received the Applied Bachelor's degree in Electrical Engineering (Power Systems) and the Applied Master's degree in Power Conversion and Renewable Energy, both from Politeknik Negeri Malang, graduating cum laude and as the top institution-level graduate of the Applied Master's programme in 2020. She was a volunteer researcher on renewable energy at PT Lentera Bumi Nusantara (2019–2021) and worked in automotive manufacturing at PT Indospring Tbk (2021–2022) before joining academia. She is currently a lecturer in the Energy Systems Engineering undergraduate programme, Faculty of Electrical Engineering, Telkom University, and serves as an Assessor at LSP-Renewable Energy and as a member of the Indonesian Association of Professional Assessors (IASPRO).*
https://orcid.org/0009-0004-4281-048X

I Gede Putu Oka Indra Wijaya *was born in Tampaksiring, Bali, on February 29, 2000. He earned his Bachelor's degree in 2022 and his Master's degree*

in 2023, majoring in electrical engineering at Telkom University in Bandung, Indonesia. In 2022, he served as a research assistant, investigating the characteristics of 3T load potential in collaboration with PLN Puslitbang. In the same year, he also presented at the 5th International Conference on Power Engineering and Renewable Energy, focusing on the optimal hybrid power plant planning in Grid 3 Nusa, Nusa Penida Island, Bali. His research interests include renewable energy and power systems. Currently, Mr. Wijaya is a lecturer at the Study Program of Energy Systems Engineering at Telkom University.

https://orcid.org/0009-0006-2455-9070

An International Journal of Optimization and Control: Theories & Applications (https://accscience.com/journal/ijocta)



This work is licensed under a Creative Commons Attribution 4.0 International License. The authors retain ownership of the copyright for their article, but they allow anyone to download, reuse, reprint, modify, distribute, and/or copy articles in IJOCTA, so long as the original authors and source are credited. To see the complete license contents, please visit http://creativecommons.org/licenses/by/4.0/.