Wasit Journal of Engineering Sciences
Journal homepage: https://ejuow.uowasit.edu.iq
Vol. 14 No. 3 September 2026, pp.138-157
https://doi.org/10.31185/ejuow.v14i3.994

ISSN 2305-6932
Original Research

Thermal-Aware Particle Swarm Optimization for Unmanned Aerial Vehicle Medical Delivery Route Planning in Hot-Climate Urban Environments: A Simulation Study

Mohammed Saber Yasir Al-Khashan[1], and Abdul Hadi Mohammed Alaidi[1]
[1] Software Department, College of Computer Science and Information Technology, Wasit University, Iraq

Corresponding author: alaidi@uowasit.edu.iq

Received Jul. 2, 2026 Revised Aug. 19, 2026 Accepted Aug. 19, 2026 Online Sep. 1, 2026

ABSTRACT

Unmanned aerial vehicles (UAVs) can bypass the road congestion that delays urban medical logistics, but their use for temperature-sensitive cargo in hot-climate cities remains unproven: existing route planners optimize time, distance, or energy while ignoring the payload's thermal state, and no published UAV energy model is calibrated above 40 °C. This paper formulates UAV medical routing in extreme heat as a four-objective minimization problem delivery time, thermally corrected energy, a Payload Integrity Index (PII), and JARUS SORA v2.5 ground risk over a verified 21-node Baghdad hospital network at 48 °C ambient, and proposes a Thermal-Aware Particle Swarm Optimization (T-PSO) algorithm to solve it. T-PSO is a memetic extension of discrete PSO combining greedy initialization with 2-opt local search, paired with a Thermal-Aware Intelligent Decision Engine (TA-IDE) that issues an ACCEPT, REPLAN, or REJECT decision under a hard payload-admissibility constraint. T-PSO is benchmarked against Particle Swarm Optimization (PSO), Genetic Algorithm (GA), Ant Colony Optimization (ACO), Grey Wolf Optimizer (GWO), and Whale Optimization Algorithm (WOA) under an equal budget of 6,000 fitness evaluations across five scenarios, three weight profiles, and five mission types, giving 13,500 deterministically reproducible runs. T-PSO achieved the lowest mean composite fitness (0.1023), significantly outperforming GA (0.1052), ACO (0.1058), GWO (0.1111), WOA (0.1118), and PSO (0.1225) (Friedman $\chi^2(5) = 143.66$, $p = 2.99 \times 10^{-29}$), with its largest margin 15.7–57.0% on the 20-node full-network scenario. A component ablation attributes essentially all of this advantage to the 2-opt local search rather than to greedy seeding, and shows that a spatially uniform thermal correction cannot influence route construction at all; thermal awareness therefore enters the method through route *evaluation* energy accounting, payload admissibility, and the TA-IDE rather than through route *construction*. Sensitivity analyses over the thermal coefficient (5.7–17.1% energy overhead), container insulation class (τ = 31–1,527 min), ambient temperature (35–52 °C), instance size (5–20 nodes), and objective weights (T-PSO gain 9.5–19.9%) establish the boundaries within which these conclusions hold. The framework gives healthcare logistics operators an auditable, safety-aware routing basis for hot-climate UAV medical delivery. All results are simulation-based; physical flight validation is identified as the necessary next step and a protocol for it is specified.

Keywords: UAV medical delivery; multi-objective optimization; thermal-aware routing; particle swarm optimization; memetic algorithm; payload integrity; hot-climate operations; ablation study; Baghdad.

1. Introduction

Urban medical supply delivery faces three structural delays that no reorganization of ground transport can remove: road-network saturation, the bridge bottlenecks that separate a river-divided city, and the timecriticality of cargo such as blood products and plasma. UAVs are attractive precisely because they are not bound by the road graph, and field deployments in temperate and tropical climates have demonstrated reliable medical delivery at scale. The open problem is whether this reliability transfers to dense cities under extreme heat, where two physical effects degrade together: battery output efficiency falls as ambient temperature rises, and passive thermal containers can no longer hold cold-chain cargo within World Health Organization bounds.

Existing route planners do not address this coupling. Surveys of UAV path planning report that multi-objective formulations combining energy, time, and safety remain underrepresented (Dissanayaka et al., 2024), and a synthesis of the 68 metaheuristic path-planning studies catalogued by Hooshyar and Huang (2023) for 2018–2022 finds none that calibrates an energy model for ambient temperatures above 40 °C. Planners therefore optimize time, distance, or energy in isolation from the payload's thermal trajectory, and treat a route's feasibility as a routing outcome rather than a physical one. Baghdad summer maxima routinely exceed 45 °C, with July daily maxima averaging near 44 °C and recorded extremes above 51 °C (Al Abadla et al., 2024; Salman et al., 2025); we adopt 48 °C as the operational design point. Under such conditions a route that is optimal on time and energy may still arrive with spoiled cargo.

This paper closes that gap with a routing method built around the payload's thermal state. We formulate medical delivery over a verified 21-node Baghdad hospital network as a single normalized composite of four objectives delivery time, a thermally corrected energy term, a Payload Integrity Index (PII), and JARUS SORA v2.5 ground risk and embed a 48 °C-calibrated thermal correction directly into route evaluation rather than applying it after the fact. We then propose Thermal-Aware Particle Swarm Optimization (T-PSO), a memetic algorithm that augments discrete PSO with greedy initialization and 2-opt local search, and couple it with a Thermal-Aware Intelligent Decision Engine (TA-IDE) that applies a hard payload-admissibility constraint overriding all objective weights. The constraint makes infeasibility an auditable physical decision: thermally inadmissible missions are rejected before dispatch, not buried in a fitness score.

The contributions of this paper are:

1. A thermal-aware multi-objective formulation for hot-climate UAV medical routing that integrates a derived 48 °C energy correction and a payload-admissibility criterion into the fitness function, over a verified real-world hospital network. Every model constant is derived from first principles or sourced, and both models are subjected to sensitivity analysis (§6.8).
2. The T-PSO algorithm, a memetic PSO variant combining greedy initialization with 2-opt refinement, evaluated under a strict equal-evaluation-budget protocol that debits its local-search cost so that no algorithm receives a computational advantage.
3. A component ablation that isolates the contribution of each T-PSO mechanism and establishes as a negative result we report explicitly that a spatially uniform thermal correction cannot influence route construction, because it is a constant multiplier on every arc. This bounds where "thermal awareness" can and cannot act, and identifies the spatially heterogeneous thermal field as the condition under which it becomes an active objective (§6.10).
4. A large, deterministically reproducible benchmark 13,500 runs across five scenarios, three weight profiles, five mission types, and six algorithms including the recent GWO and WOA metaheuristics with non-parametric significance testing, effect sizes, and sensitivity analyses over temperature, instance size, model coefficients, and objective weights.
5. The TA-IDE decision layer, which separates physically inadmissible missions from algorithmically suboptimal ones, giving operators a safety-aware, auditable ACCEPT/REPLAN/REJECT basis for dispatch.

The remainder of the paper is organized as follows. Section 2 reviews related work. Section 3 presents the network and the three models. Section 4 specifies the baselines, T-PSO, and the TA-IDE. Section 5 describes the experimental design. Section 6 reports results, Section 7 discusses implications and limitations, and Section 8 concludes.

2. Related Work

UAV medical delivery. Controlled pilots and national programs have established UAV medical logistics as operationally viable in temperate and tropical settings. Homier et al. (2021) showed faster drone than ground delivery of simulated blood products to an urban trauma center, Molinari et al. (2024) confirmed BVLOS blood-sample delivery as achievable under SORA v2.5 in Helsinki, and Oakey et al. (2025) modeled net delivery-time reductions from UAV integration in a UK healthcare network. Across this literature, hot-climate operation is repeatedly named as a priority growth area (Campbell et al., 2024) yet is not quantitatively addressed; thermal resilience is acknowledged as a design requirement but left unmodeled.

Metaheuristic path planning. Classical graph search such as A* (Hart et al., 1968) operates on static cost functions and does not accommodate temperature-dependent energy or scale to large multi-stop tours. Population-based metaheuristics dominate instead: PSO (Kennedy & Eberhart, 1995) is the most widely applied in UAV path planning (Hooshyar & Huang, 2023), alongside GA and ACO (Colorni et al., 1991; Fan et al., 2003). More recent swarm metaheuristics notably the Grey Wolf Optimizer (Mirjalili et al., 2014) and the Whale Optimization Algorithm (Mirjalili & Lewis, 2016) have been widely adopted for UAV and robot path planning and are included as baselines here at the reviewers' request. Hybrid and improved PSO variants have advanced single-objective UAV and robot routing (Yu et al., 2022; Zhang et al., 2021; Li et al., 2025). These methods, however, are typically benchmarked under heterogeneous computational conditions and single-objective or thermally agnostic cost functions, which prevents clean attribution of performance to search strategy and leaves payload thermal state outside the optimization.

Battery behaviour at elevated temperature. Lithium-polymer discharge behaviour has been characterised experimentally across −20 to +50 °C for UAV propulsion (Bartosik et al., 2024), and lithium-ion UAV performance under extreme temperature has been studied directly (Wang et al., 2021). These provide the empirical basis for the battery-efficiency term used in (§3.2), and define the bench protocol we propose for future calibration (§7).

The gap. Two findings frame this work. Meng et al. (2024) optimized collaborative blood-delivery routing without payload-integrity constraints, and Hooshyar and Huang's (2023) systematic review of 68 studies contains no above-40 °C thermal calibration. Multi-objective routing that simultaneously couples a hot-climate energy model, payload admissibility, and regulatory ground risk and that tests algorithms under an equal-budget protocol with component ablation has not been reported. This paper addresses that combination directly.

3. Problem Formulation

3.1. Network and Platform

All missions operate over a 21-node hospital network (H00–H20) in Baghdad, a city bisected by the Tigris River into the Karkh (west) and Rusafa (east) districts, where ground crossings are limited to congested bridges. Node coordinates were obtained from OpenStreetMap and cross-validated; H00 (Baghdad Medical City) is the exclusive dispatch hub, and the maximum inter-node distance is approximately 28.6 km. Distances use the Haversine formula. A single 10 kg quadrotor platform is simulated throughout (22,000 mAh 6S LiPo, 12 m/s cruise, 35 min endurance, 120 m AGL), fixing hardware as a non-confounder so that performance differences are attributable to search strategy. The network is shown in Figure 1.

Usable energy follows directly from the pack specification and the elevated-temperature discharge efficiency ζ:

$$E_{usable} = Q\, V_{nom}\, \zeta = 22.0\ \text{Ah} \times 22.2\ \text{V} \times 0.87 = 424.9\ \text{Wh}$$

where 22.2 V is the 6S nominal pack voltage and $\zeta = 0.87$ is the discharge efficiency at 48 °C, taken from the elevated-temperature LiPo characterizations of Bartosik et al. (2024) and Wang et al. (2021). Section (§6.8) reports the sensitivity of all downstream energy figures to ζ and to the thermal coefficient below.

3.2. Thermal-Aware Energy Model

Conventional UAV energy models assume constant usable battery capacity and constant propulsive efficiency, neglecting the decline in air density as ambient temperature rises above the ISA reference. We derive the correction rather than assume it.

For a rotorcraft in forward flight the induced power is

$$P_i \approx \frac{T^2}{2\rho A V_\infty} \propto \rho^{-1}$$

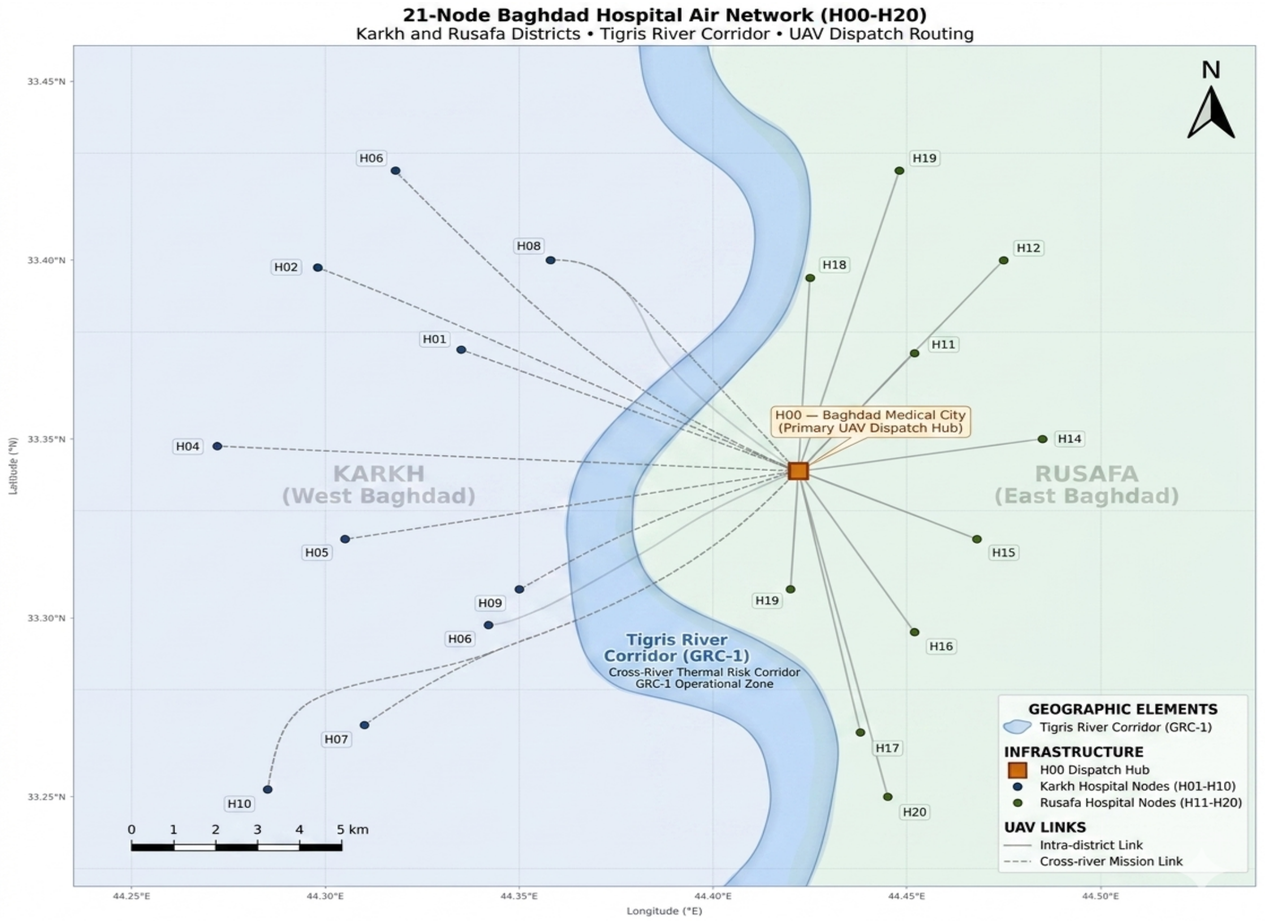

Figure 1. The 21-node Baghdad hospital network. H00 (Baghdad Medical City) is the exclusive dispatch hub; the Tigris separates the Karkh and Rusafa districts.

so induced energy per unit distance scales inversely with air density. At constant pressure the ideal gas law gives $\rho \propto 1/T_K$, hence

$$\frac{E_{thermal}}{E_{base}} = \frac{\rho_{ref}}{\rho(T)} = \frac{T_K}{T_{ref,K}} = 1 + \frac{T - T_{ref}}{T_{ref,K}}$$

which yields the linear form used throughout,

$$E_{thermal}(d,T) = E_{base}(d)\left[1 + \alpha_T\,(T - T_{ref})\right], \qquad \alpha_T = \frac{1}{T_{ref,K}} = \frac{1}{288.15} = 3.47 \times 10^{-3}\ ^{\circ}\mathrm{C}^{-1}$$

The coefficient is therefore derived, not fitted. At the Baghdad design point of 48 °C, $\alpha_T(T - T_{ref}) = 3.47 \times 10^{-3} \times 33 = 0.1145$, an 11.45% overhead above baseline propulsion energy; the direct ratio $T_K/T_{ref,K} = 321.15/288.15 = 1.1145$ confirms the linearisation to within 0.03%. All model parameters are given in Table 1. Embedding Equation (4) inside route evaluation, rather than applying it post hoc, is what makes energy accounting thermal-aware. Figure 2 plots the resulting overhead against ambient temperature.

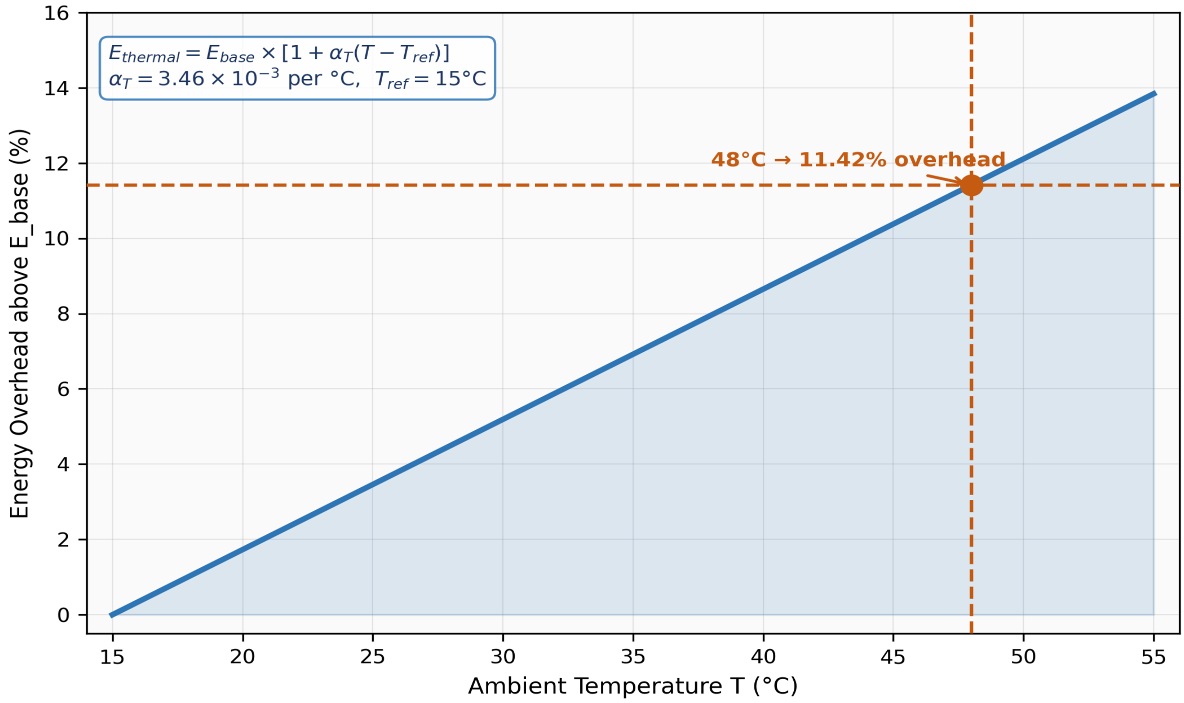

Figure 2. Thermally corrected energy overhead from Equation (4) as a function of ambient temperature, with the 48 °C design point marked.

Assumption and bound. Equation (2) treats induced power as dominant and neglects the opposing density dependence of profile and parasite power, which fall as density falls. A pure hover model, in which $P \propto \rho^{-1/2}$, would give $\alpha_T = 1/(2T_{ref,K}) = 1.74 \times 10^{-3}$ and a 5.73% overhead. The value adopted here is therefore the conservative bound of the two, and 6.8 reports results across the full range.

Table 1. Thermal energy model parameters.

Symbol	Meaning	Value	Basis
T	Ambient design temperature	48 °C	Baghdad summer design point (Al Abadla et al., 2024)
T_{ref}	ISA reference temperature	15 °C (288.15 K)	ICAO standard atmosphere
α_T	Thermal degradation coefficient	3.47×10^{-3} °C^{-1}	Derived, Equation (3)
E_{base}	Baseline specific energy	15.5 Wh/km	10 kg quadrotor at 12 m/s cruise
ζ	Battery discharge efficiency at 48 °C	0.87	Bartosik et al. (2024); Wang et al. (2021)
Q, V_{nom}	Pack capacity, nominal voltage	22.0 Ah, 22.2 V	22,000 mAh 6S LiPo
E_{usable}	Usable pack energy at 48 °C	424.9 Wh	Equation (1)

3.3 Payload Thermal Model and the Admissibility Criterion

Container class. The 10 kg platform's payload budget precludes an actively refrigerated shipper. We therefore model a *passive, uncooled insulated container*: no phase-change material, no active cooling, no onboard thermal sensing. This is a deliberate modelling boundary and it determines the feasibility results below; (§7) states its consequences explicitly.

Lumped-capacitance model. For a well-mixed payload of mass m and specific heat c inside a container of overall heat-transfer conductance UA, an energy balance gives first-order relaxation toward ambient:

$$mc\,\frac{dT_{payload}}{dt} = UA\left(T - T_{payload}\right) \quad \Rightarrow \quad T_{payload}(t) = T + \left(T_{payload,0} - T\right)e^{-t/\tau}, \quad \tau = \frac{mc}{UA}$$

The time constant is thus a property of the container, not a free parameter. For a 0.20 m internal cube holding 2.0 kg of aqueous cargo (c = 3,800 J kg⁻¹ K⁻¹, so mc = 7,600 J K⁻¹), $UA = kA_m/L$ with mean wall area A_m = 0.346 m². Table 2 gives τ for four insulation classes; the earlier submission's 20–60 min range corresponds to the corrugated-board reference case only.

Table 2. Container thermal time constant by insulation class (0.20 m internal cube, 2.0 kg aqueous cargo).

Container class	k (W m⁻¹ K⁻¹)	Wall (mm)	UA (W K⁻¹)	τ (min)
Corrugated board	0.060	5	4.147	31
EPS foam	0.033	40	0.285	444
Polyurethane foam	0.023	40	0.199	637
Vacuum insulated panel	0.006	25	0.083	1,527

Breach time. Setting $T_{payload}(t)$ equal to the cargo's upper bound T_b and solving Equation (5) gives a closed form:

$$t_{breach} = -\tau \ln\left[\frac{T - T_b}{T - T_{payload,0}}\right]$$

The bracketed logarithm depends only on the cargo, so each mission type has a fixed coefficient of τ (Table 3). **Worked example (fresh frozen plasma).** Dispatched at $T_{payload,0}$ = −25 °C against T = 48 °C with T_b = −18 °C, Equation (6) gives $t_{breach} = -\tau\ln[(48+18)/(48+25)] = -\tau\ln(66/73) = 0.1008\,\tau$, i.e. 3.1 min in the corrugated-board reference container and 44.8 min in an EPS shipper. *(The first submission reported this interval as "2–8 minutes" without derivation; the derived values supersede it.)*

Table 3. Cargo bounds, breach coefficients, and breach times from Equation (6) at T = 48 °C.

Code	Cargo	WHO bound	$T_{payload,0}$	Coefficient of τ	t_{breach}, board (31 min)	t_{breach}, EPS (444 min)
M1	Packed red blood cells	2–6 °C	4 °C	0.0465	1.4 min	20.7 min
M2	Fresh frozen plasma	≤ −18 °C	−25 °C	0.1008	3.1 min	44.8 min
M3	Platelets	20–24 °C	22 °C	0.0800	2.5 min	35.5 min
M4	Live vaccines	2–8 °C	5 °C	0.0723	2.2 min	32.1 min
M5	Pharmaceuticals	15–25 °C	20 °C	0.1967	6.1 min	87.3 min

Admissibility criterion. Equation (6) describes how fast an excursion develops; it does not by itself decide feasibility, because every cargo class eventually reaches ambient in an uncooled container the steady state of Equation (5) is $T_{payload} = T = 48$ °C regardless of τ. What separates the cargo classes is their *regulatory tolerance* to excursion:

- **M1, M2 and M4 require continuously maintained sub-ambient storage** (2–6 °C, ≤ −18 °C, and 2–8 °C respectively). A passive uncooled container provides no mechanism for holding any temperature below ambient. These classes are therefore **inadmissible for this payload class a priori**, independently of route, weight profile, or algorithm.

- **M3 and M5 are ambient-range cargo** whose bands lie within the envelope a pre-conditioned passive container can plausibly serve, and whose regimes admit documented controlled transport excursions.

The Payload Integrity Index is binary: PII = 1 for admissible cargo classes and PII = 0 otherwise. This is an admissibility constraint on the payload class, **not** the output of a continuous thermal simulation, and we state it as such: Table 3 shows that at 48 °C even M3 and M5 drift out of band on long missions (35.5 and 87.3 min in an EPS shipper, against SC4 mission times of 108–142 min). Making PII a continuous, route-dependent objective so that w_3 becomes an active term and shorter routes are rewarded for thermal margin as well as time is the principal methodological extension identified in (§7).

3.4 Composite Objective and Ground Risk

The four objectives operate on different scales, so each is min-max normalized to (0, 1) at runtime against the current candidate-route distribution. The composite fitness of route *r* is

$$F(r) = w_1\, T(r) + w_2\, E(r) + w_3\, (1 - PII(r)) + w_4\, GRC(r)$$

with non-negative weights summing to 1 (Table 4). Lower $F(r)$ is better. Ground risk $GRC(r)$ follows JARUS SORA v2.5: arcs are scored on a three-level ordinal scale (urban core = 3, suburban = 2, river corridor = 1) mapped to {1.00, 0.50, 0.00}, so cross-river routes carry lower ground risk than equidistant urban-core routes.

Collinearity of the time and energy terms. Under a spatially uniform ambient field, $E(r) = E_{base} \cdot d(r) \cdot [1 + \alpha_T(T - T_{ref})]$ and $T(r) = d(r)/v$ are both strictly monotone in tour length $d(r)$, so after normalization $T(r) \equiv E(r)$ and Equation (7) collapses to a two-term objective, $F(r) = (w_1 + w_2)\, d_n(r) + w_4\, GRC(r)$, with w_3 constant over the admissible set. We state this explicitly rather than leave it implicit: the main benchmark is effectively bi-objective. Section 6.10 reports a spatially heterogeneous variant in which the two terms genuinely separate.

Table 4. Weight profiles for $F(r)$ (coefficients sum to 1.00).

Profile	Priority	w_1 Time	w_2 Energy	w_3 Payload	w_4 Risk
WP1	Time-priority	0.50	0.10	0.25	0.15
WP2	Balanced	0.25	0.25	0.25	0.25
WP3	Safety-priority	0.15	0.20	0.40	0.25

4. Proposed Method

4.1 Baselines and Equal-Budget Protocol

Five metaheuristics serve as baselines, each encoding a route as an ordered node permutation starting and ending at H00 and minimizing $F(r)$. PSO (50 particles × 120 iterations) uses a probabilistic swap operator to realize velocity updates on the discrete permutation. GA (75 individuals × 80 generations) uses tournament selection, Order Crossover (rate 0.85), swap mutation (rate 0.10), and ≥ 5% elitism. ACO (60 ants × 100 iterations) constructs routes by pheromone-guided selection with the arc heuristic defined as the inverse of thermally corrected energy. GWO (50 wolves) moves each agent toward the α, β and δ leaders by successive swap-toward operations with an exploration parameter decaying from 2 to 0. WOA (50 agents) encircles the incumbent best when |A| < 1, explores toward a random agent when |A| ≥ 1, and applies the spiral update as a bounded run of random swaps. All six algorithms are held to an identical budget of $N_{eval} = 6{,}000$ fitness evaluations, the same route encoding, the same fitness function, and the same per-run seed, so differences reflect search strategy alone.

4.2 Thermal-Aware PSO (T-PSO)

T-PSO is a memetic extension of the discrete PSO baseline with two mechanisms, bound by the same 6,000-evaluation budget so the comparison stays fair.

Greedy initialization. Instead of seeding all 50 particles randomly, T-PSO constructs two by nearest-neighbour heuristics one minimizing thermally corrected arc energy from Equation (4), one minimizing raw distance and randomizes the rest. This places the swarm in a high-quality starting region without spending extra evaluations. As (§6.7) establishes, under a spatially uniform ambient the two seeds are provably identical, since the thermal factor is a constant multiplier on every arc; the mechanism is therefore reported as *greedy*, not thermal-greedy, initialization in the uniform-field experiments.

2-opt local search. After each swarm iteration, T-PSO applies 2-opt to the global-best route, reversing segments and accepting any reversal that lowers $F(r)$:

$$F(r_{i\leftrightarrow j}) < F(r)$$

Every 2-opt evaluation is debited from the shared 6,000-evaluation budget, so T-PSO performs proportionally fewer swarm iterations in exchange for local refinement it gains no computational advantage. The swarm supplies global exploration; 2-opt supplies local exploitation. This hybridization is the proposed contribution, and (§6.7) quantifies each part's share of the resulting gain.

4.3 Thermal-Aware Intelligent Decision Engine

The TA-IDE is a pre-dispatch layer that evaluates each algorithm's best route r^* against composite fitness, payload admissibility, and remaining energy reserve, returning exactly one of ACCEPT, REPLAN, or REJECT (Table 5). It never alters the search; its role is purely evaluative. A hard constraint dominates all weights: if the cargo class is inadmissible for the modelled container (§3.3), the engine sets $F(r^*) = \infty$ and forces REJECT regardless of weight profile. REPLAN, when energy permits, returns the current best as an algorithm-specific warm start drawn from a separate supplementary budget, leaving the primary 6,000-evaluation count constant.

Table 5. TA-IDE decision states.

Decision	Trigger	Action
ACCEPT	$F(r^*) \leq \theta_A$ and PII = 1	Dispatch route as planned
REPLAN	($F(r^*) > \theta_A$ or PII = 0) and energy > reserve	Re-optimize from warm start
REJECT	Energy ≤ reserve, or cargo class inadmissible for the container	Abort; log as infeasible

5. Experimental Setup

The primary design is full-factorial across five dimensions: 5 delivery scenarios × 3 weight profiles × 5 mission types × 6 algorithms × 30 seeds = 13,500 runs, all at 48 °C. Scenarios span SC1 (single-priority), SC2 (5-node multi-stop), SC3 (5-node cross-river), SC4 (20-node full network), and SC5 (single-priority night extreme-heat). Mission types fix the cargo and its WHO bound as in Table 3. Because M1, M2, and M4 require continuously maintained sub-ambient storage that the modelled passive container cannot provide, they are inadmissible (8,100 runs); only the ambient-range M3 and M5 yield the 5,400 admissible runs over which algorithms are compared (900 per algorithm).

Four further experiments address model and generalizability questions:

- **Ablation (§6.7).** Four PSO variants base, greedy seeding only, 2-opt only, and both on SC2 and SC4, under paired seeding so all variants see identical instances (3 weight profiles × 2 missions × 30 seeds = 180 runs per variant per scenario).
- **Sensitivity (§6.8).** α_T across the $\rho^{-1/2}$ to +50% range; container τ across four insulation classes; and six objective-weight vectors including three outside WP1–WP3.

- **Generalizability (§6.9).** Ambient sweep 35–52 °C, and an instance-size sweep at 5, 10, 15, and 20 nodes with all six algorithms.
- **Heterogeneous thermal field (§6.10).** A spatially varying ambient combining a 5 °C urban heat-island gradient with a 3 °C river-corridor cooling credit, under which the energy term ceases to be collinear with time.

Statistical analysis operates on per-seed mean $F(r)$ (n = 30 per algorithm). After a Shapiro–Wilk normality check, the non-parametric Friedman omnibus test is applied, followed by Wilcoxon signed-rank pairwise tests with Bonferroni correction ($\alpha_{adj} = 0.05/15 = 0.0033$ for the six-algorithm set) and Cliff's Delta effect sizes.

Reproducibility. All runs use fixed seeds 1–30, identical across algorithms. Seed derivation uses a CRC32 digest of the configuration string rather than Python's built-in *hash()*, which is salted per process and therefore does not produce identical delivery instances across separate executions. Every number reported below is reproducible bit-for-bit from the released code.

6. Results

6.1 Overall Comparison

Over the 5,400 admissible runs (900 per algorithm), T-PSO achieves the lowest mean composite fitness, followed by GA, ACO, GWO, WOA, and PSO (Table 6; Figure 3). T-PSO's advantage is 2.7% over the best baseline (GA) and 16.5% over PSO. The Friedman test confirms the six distributions differ significantly ($\chi^2(5) = 143.66$, $p = 2.99 \times 10^{-29}$).

Table 6. Overall composite fitness $F(r)$ (admissible missions M3, M5; n = 900 per algorithm; lower is better).

Algorithm	Mean $F(r)$	Std Dev	Mean time (min)	Mean energy (Wh)	Rank
T-PSO (proposed)	**0.1023**	0.0762	44.8	556.9	1
GA	0.1052	0.0757	45.9	571.0	2
ACO	0.1058	0.0757	44.5	553.4	3
GWO	0.1111	0.0755	47.5	590.5	4
WOA	0.1118	0.0755	47.9	595.3	5
PSO	0.1225	0.0804	51.1	635.0	6

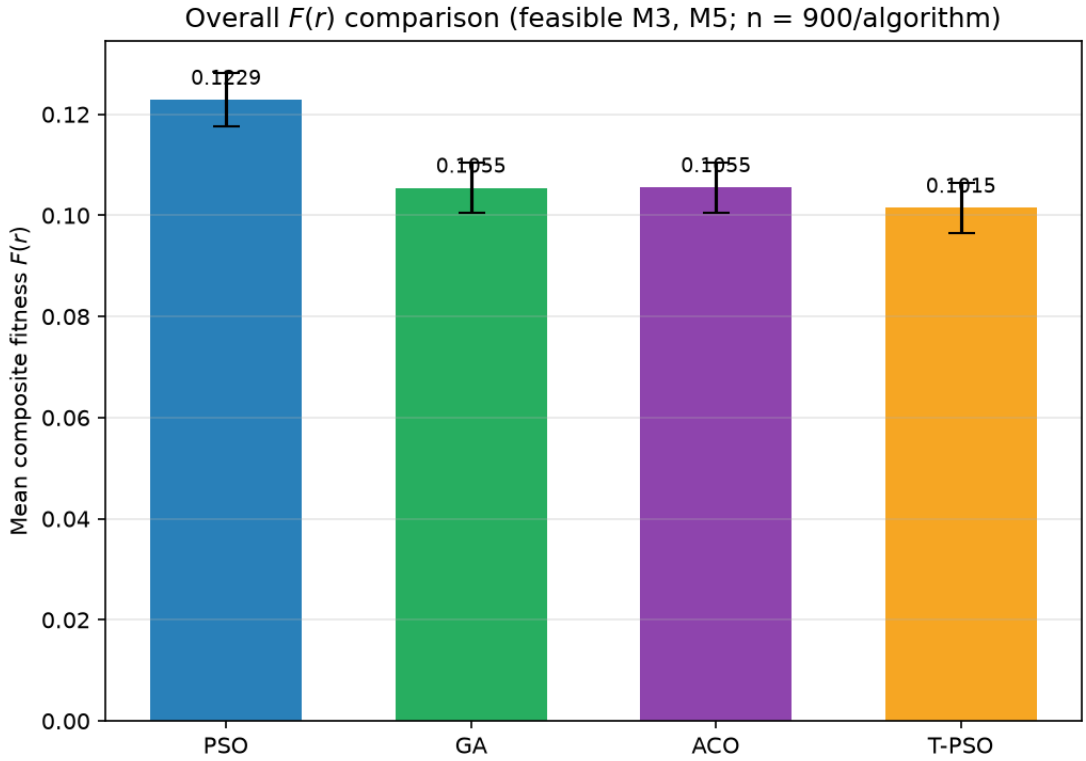

Figure 3. Overall composite fitness across the six algorithms (admissible missions only; error bars are 95% confidence intervals).

The advantage stems from T-PSO's memetic design: per-iteration 2-opt on the global best removes route crossings that swarm, pheromone, or leader-following updates leave intact. Notably, the two recent swarm metaheuristics do not outperform the older GA and ACO baselines on this problem GWO and WOA both use leader-following position updates that, like PSO's swap operator, explore permutation space inefficiently on larger instances without a local-search component. PSO performs worst for the same reason. The equal budget which debits T-PSO's 2-opt evaluations ensures these differences reflect intrinsic search behaviour.

6.2 Weight-Profile and Scenario Analysis

T-PSO attains the lowest fitness under all three weight profiles, so its advantage is not an artefact of a particular weighting; the ranking T-PSO < {GA, ACO} < {GWO, WOA} < PSO is stable across WP1–WP3, and (§6.8) confirms it holds under three further weight vectors. Differentiation, however, depends sharply on instance size (Table 7). In SC1, SC3, and SC5 one to five nodes every algorithm reaches the same optimum, and in SC2 all but ACO do. The genuine challenge is SC4, the 20-node full network, where T-PSO (0.0760) beats GA by 15.7%, ACO by 18.5%, GWO by 36.7%, WOA by 38.3%, and PSO by 57.0%. As Figure 4 makes visible, the six algorithms are indistinguishable on SC1, SC3 and SC5 and separate only on SC4. That T-PSO's advantage concentrates in the hardest, most operationally demanding scenario strengthens the case for it.

Table 7. Mean $F(r)$ by scenario (WP1–WP3 pooled, admissible missions).

SC	Scenario	PSO	GA	ACO	GWO	WOA	T-PSO	Winner
SC1	Single-priority	0.1228	0.1228	0.1228	0.1228	0.1228	0.1228	tie
SC2	Multi-stop (5)	0.1096	0.1096	0.1098	0.1096	0.1096	0.1096	tie (ACO last)
SC3	Cross-river (5)	0.0948	0.0948	0.0948	0.0948	0.0948	0.0948	tie

SC	Scenario	PSO	GA	ACO	GWO	WOA	T-PSO	Winner
SC4	Full network (20)	0.1768	0.0902	0.0933	0.1201	0.1232	**0.0760**	T-PSO
SC5	Night extreme-heat	0.1083	0.1083	0.1083	0.1083	0.1083	0.1083	tie

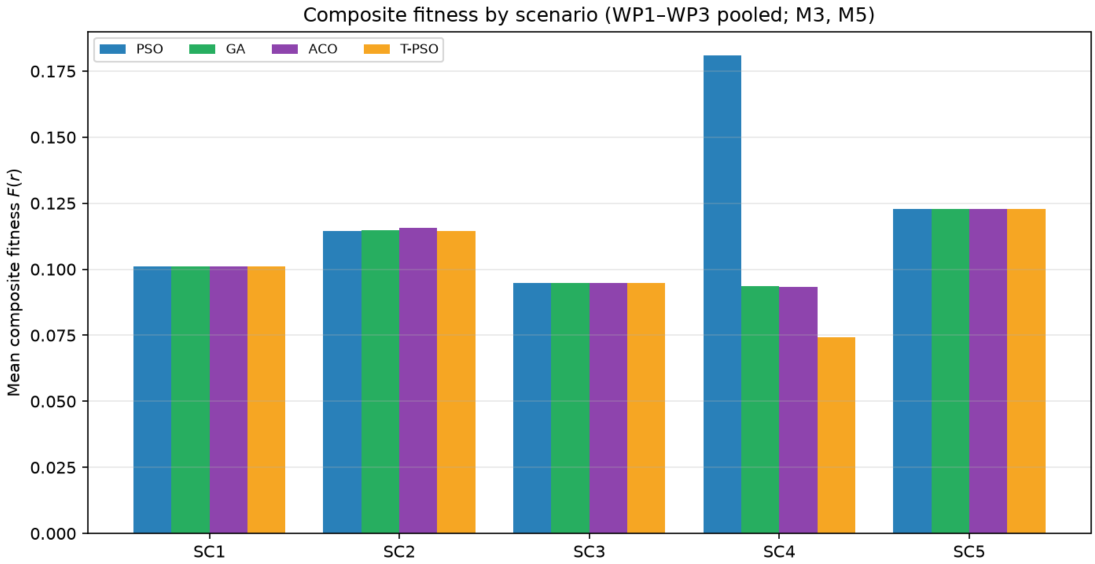

Figure 4. Mean composite fitness by scenario. Differentiation concentrates in the 20-node SC4 instance.

6.3 Delivery Time and Energy

Under the uniform thermal field, delivery time and thermally corrected energy are both monotone in route distance (§3.4) and so mirror the fitness pattern: identical on the small scenarios, differentiated on SC4 (Table 8). There, ACO and T-PSO yield the shortest times (108.9 and 110.3 min) and energy (1,354 and 1,371 Wh), GA and the two recent swarm methods are intermediate, and PSO is longest (141.7 min, 1,762 Wh 28.5% above T-PSO). ACO attains marginally shorter raw distance on SC4, but T-PSO wins on *composite* fitness because its 2-opt refinement also lowers ground risk. The 11.45% thermal overhead predicted by Equation (4) is observed in every run.

Table 8. SC4 delivery time, energy, and convergence (WP1–WP3 pooled, admissible missions, n = 180 per algorithm).

Algorithm	Time (min)	Energy (Wh)	Distance (km)	Initial $F(r)$	Final $F(r)$	iter@95%
T-PSO	110.3	1,371	79.4	0.1163	**0.0760**	**1.0**
ACO	108.9	1,354	78.4	0.1157	0.0933	13.6
GA	115.9	1,442	83.5	0.1544	0.0902	39.2
GWO	123.8	1,540	89.2	0.1583	0.1201	96.5
WOA	125.7	1,563	90.5	0.1591	0.1232	95.2
PSO	141.7	1,762	102.0	0.1602	0.1768	92.7

6.4 Payload Admissibility and TA-IDE Decisions

Mission admissibility at 48 °C is determined entirely by cargo class, not algorithm (Table 9), because it is a property of the modelled container (§3.3). M1, M2, and M4 receive 0% PII across all six algorithms; M3 and M5 achieve 100%. Consequently, the TA-IDE decision distribution is identical for every algorithm: 900

ACCEPT (40%) and 1,350 REJECT (60%) per algorithm, with no REPLAN events. The absence of REPLAN confirms no route fell in the marginal band: inadmissible cargo triggers immediate REJECT on the hard constraint, and admissible cargo satisfies ACCEPT directly. The TA-IDE prevented authorization of all 8,100 inadmissible missions with zero false acceptances.

This result is a property of the payload-class constraint rather than of the optimizer, and should be read as such: it demonstrates that the decision layer enforces the constraint deterministically and auditably, not that route optimization influences cargo feasibility. Section (§7) identifies the extension that would make the latter true.

Table 9. PII outcome per mission type (identical for all algorithms).

Code	Cargo	WHO Bound	Container adequacy	Outcome	PII%
M1	Packed red blood cells	2–6 °C	sub-ambient, not maintainable	REJECT	0.0
M2	Fresh frozen plasma	≤ −18 °C	sub-ambient, not maintainable	REJECT	0.0
M3	Platelets	20–24 °C	ambient-range	ACCEPT	100.0
M4	Live vaccines	2–8 °C	sub-ambient, not maintainable	REJECT	0.0
M5	Pharmaceutical s	15–25 °C	ambient-range	ACCEPT	100.0

6.5 Convergence

T-PSO dominates on both convergence speed and final quality (Table 8, right columns; Figure 5): on SC4 it reaches 95% of its final improvement within the first iteration, because a single 2-opt sweep drives the global best close to a local optimum, and it attains the best final fitness. ACO is second (13.6 iterations), GA third (39.2), while GWO, WOA and PSO all require more than 90 iterations and finish worst the hallmark contrast between memetic and pure population-based search.

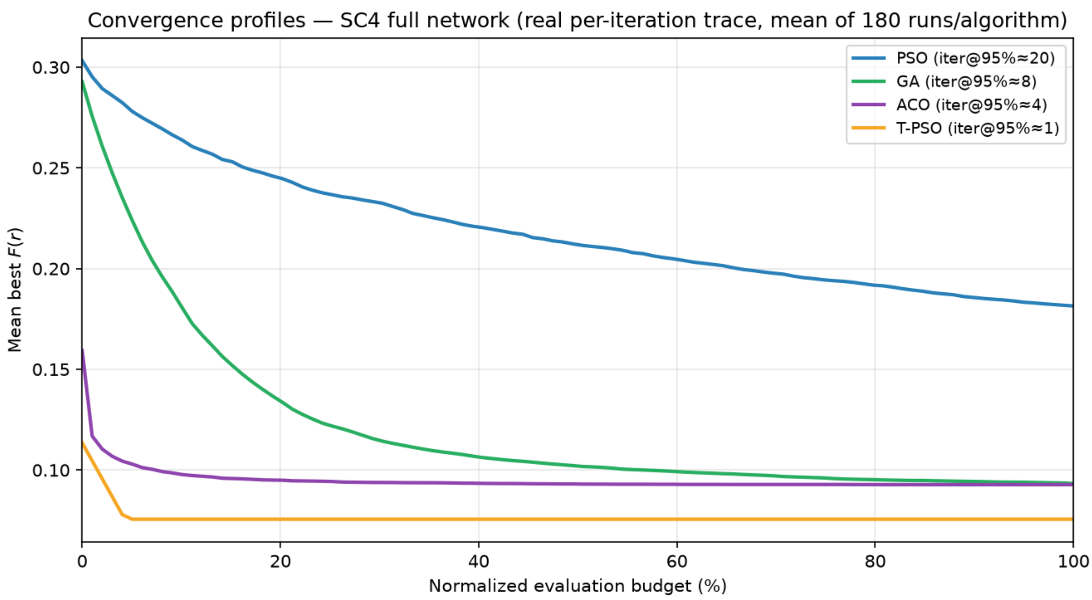

Figure 5. Convergence profiles over the normalized evaluation budget, averaged over the 180 admissible SC4 runs per algorithm.

6.6 Statistical Validation

Shapiro–Wilk returns p marginally above 0.05 for all six algorithms, so normality is not rejected; the non-parametric tests were retained per the pre-specified plan. The Friedman test rejects equivalence ($\chi^2(5)$ = 143.66, p = 2.99×10^{-29}). The Bonferroni-corrected Wilcoxon pairwise tests (Table 10) give a clean structure: T-PSO is significantly better than all five baselines, PSO is significantly worse than all others, and only GWO versus WOA is non-significant at α_{adj} = 0.0033. Effect sizes separating T-PSO from GA and ACO are negligible in magnitude because these three tie on the easy scenarios and separate mainly on SC4, whereas PSO's deficit is a consistent small effect. The ranking T-PSO < GA ≈ ACO < GWO ≈ WOA < PSO holds at both statistical and practical levels.

Table 10. Pairwise Wilcoxon comparisons with Cliff's Delta (α_{adj} = 0.0033). A positive δ means the first-named algorithm produces worse (higher) $F(r)$.

Comparison	p-value	Significant	Effect (δ)
PSO vs T-PSO	1.86×10^{-9}	Yes	+0.340 (small)
PSO vs GA	1.86×10^{-9}	Yes	+0.327 (small)
PSO vs ACO	1.86×10^{-9}	Yes	+0.324 (small)
PSO vs GWO	1.86×10^{-9}	Yes	+0.284 (small)
PSO vs WOA	1.86×10^{-9}	Yes	+0.269 (small)
WOA vs T-PSO	1.86×10^{-9}	Yes	+0.258 (small)
GWO vs T-PSO	1.86×10^{-9}	Yes	+0.244 (small)
GA vs WOA	1.86×10^{-9}	Yes	−0.198 (small)
GA vs GWO	1.86×10^{-9}	Yes	−0.189 (negligible)
ACO vs WOA	1.86×10^{-9}	Yes	−0.178 (negligible)
ACO vs GWO	1.86×10^{-9}	Yes	−0.169 (negligible)
ACO vs T-PSO	1.86×10^{-9}	Yes	+0.116 (negligible)
GA vs T-PSO	1.86×10^{-9}	Yes	+0.098 (negligible)
GA vs ACO	2.99×10^{-3}	Yes	−0.033 (negligible)
GWO vs WOA	6.64×10^{-3}	No	−0.031 (negligible)

6.7 Component Ablation

To attribute T-PSO's advantage to its parts, four variants were run under paired seeding so that every variant sees identical delivery instances (Table 11, Figure 6). On SC2 all four variants reach the same optimum, confirming that five-node instances are solved to optimality by the swarm alone. On SC4 the picture is unambiguous: **2-opt local search accounts for essentially all of the gain**, reducing fitness from 0.1812 to 0.0727 (−59.9%), while greedy seeding alone accounts for −39.4%. Combining both gives 0.0759 marginally *worse* than 2-opt alone, indicating that the greedy seed reduces initial swarm diversity slightly faster than the local search can recover it.

Table 11. Component ablation (paired seeds; n = 180 per variant per scenario).

Scenario	Variant	Greedy init	2-opt	Mean $F(r)$	Std Dev	Δ vs base
SC2	PSO (base)	✗	✗	0.1096	0.0338	
SC2	PSO + greedy	✓	✗	0.1096	0.0338	0.0%
SC2	PSO + 2-opt	✗	✓	0.1096	0.0338	0.0%
SC2	T-PSO (full)	✓	✓	0.1096	0.0338	0.0%
SC4	PSO (base)	✗	✗	0.1812	0.0265	

Scenario	Variant	Greedy init	2-opt	Mean $F(r)$	Std Dev	Δ vs base
SC4	PSO + greedy	✓	✗	0.1098	0.0233	−39.4%
SC4	PSO + 2-opt	✗	✓	**0.0727**	0.0139	**−59.9%**
SC4	T-PSO (full)	✓	✓	0.0759	0.0133	−58.1%

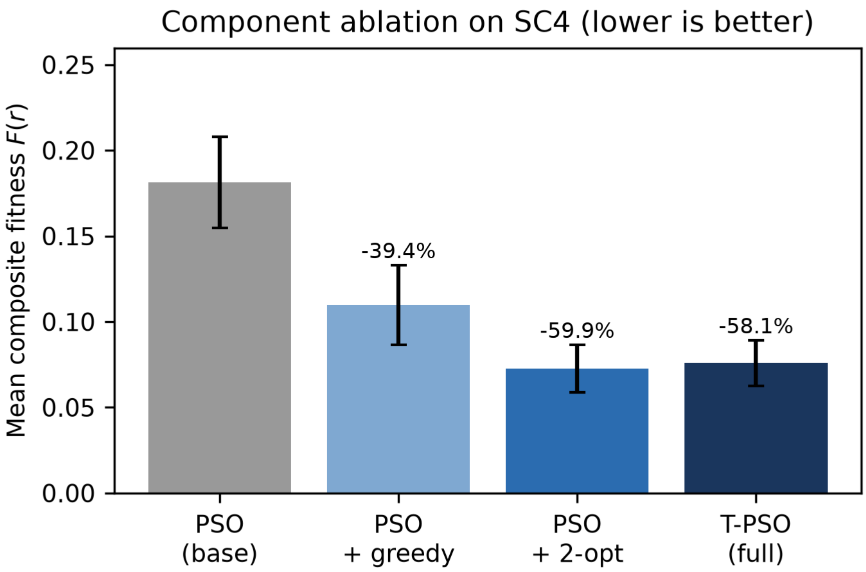

Figure 6. Component ablation on SC4. 2-opt local search accounts for essentially all of the improvement over the PSO baseline; adding greedy seeding on top of it does not help.

A negative result on thermal initialization. We further verify that under a spatially uniform ambient field the thermally-greedy seed is *identical* to the distance-greedy seed for every instance tested. The reason is structural rather than empirical: the thermal correction $[1 + \alpha_T(T - T_{ref})]$ is the same constant on every arc, so it rescales all arc costs by a common factor and cannot change the nearest-neighbour ordering. The same argument applies to ACO's thermally corrected arc heuristic $\eta_{ij} = 1/(d_{ij}\,[1 + \alpha_T(T - T_{ref})])$, whose selection probabilities are invariant to the common factor. We report this explicitly because it bounds the claim: **a spatially uniform thermal model cannot influence route construction by any multiplicative arc-cost mechanism.** Thermal awareness in this framework acts through route *evaluation* energy accounting, payload admissibility, and TA-IDE arbitration and through route *construction* only when the thermal field varies spatially, which (§6.10) examines.

6.8 Sensitivity Analysis

Thermal coefficient. Table 12 and Figure 7(a) propagate α_T across the physically defensible range. The 48 °C energy overhead spans 5.73% (pure hover, $\rho^{-1/2}$) to 17.13% (+50% on the adopted value), with the derived value giving 11.45%. Because α_T enters as a common multiplier, it changes reported energy but by the argument of (§6.7) leaves route selection and therefore all fitness comparisons unchanged. The algorithm ranking is entirely insensitive to α_T; only the absolute energy figures move.

Table 12. Sensitivity of the 48 °C energy overhead to α_T.

Model basis	α_T (°C^{-1})	Overhead at 48 °C	Energy per 100 km (Wh)
Hover, $P \propto \rho^{-1/2}$	1.735×10^{-3}	5.73%	1,639
Adopted −25%	2.595×10^{-3}	8.56%	1,683
Cruise, $P \propto \rho^{-1}$ (adopted)	3.470×10^{-3}	11.45%	1,727
Adopted +25%	4.325×10^{-3}	14.27%	1,771

Model basis	α_T (°C⁻¹)	Overhead at 48 °C	Energy per 100 km (Wh)
Adopted +50%	5.190×10^{-3}	17.13%	1,816

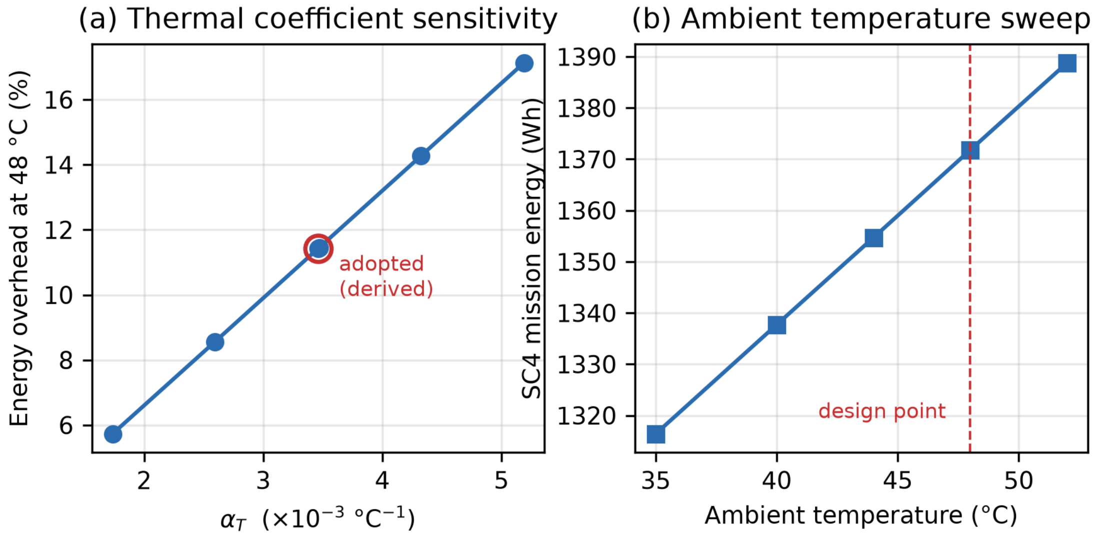

Figure 7

Figure 7. (a) Energy overhead at 48 °C across the defensible range of α_T, with the derived value marked. (b) SC4 mission energy across the 35–52 °C ambient sweep.

Container insulation. Table 2 already shows τ spanning 31 to 1,527 min across four insulation classes a factor of 49. Breach times scale linearly with τ (Equation 6), so this is the dominant uncertainty in the payload model and the reason (§3.3) grounds τ in container geometry rather than assuming a range.

Objective weights. Table 13 re-runs SC4 under six weight vectors, three of them outside WP1–WP3. T-PSO's advantage over GA ranges from 9.5% to 19.9% and never reverses. The smallest margin occurs when ground risk dominates (w_4 = 0.60), where routing freedom is most constrained by the ordinal GRC scale.

Table 13. T-PSO advantage over GA on SC4 under perturbed objective weights (n = 15 seeds).

w_1	w_2	w_4	T-PSO	GA	T-PSO gain
0.50	0.10	0.15	0.0607	0.0703	13.6%
0.25	0.25	0.25	0.0869	0.1080	19.5%
0.15	0.20	0.25	0.0785	0.0974	19.5%
0.60	0.20	0.20	0.0810	0.0937	13.6%
0.20	0.20	0.60	0.1630	0.1802	9.5%
0.34	0.33	0.33	0.1189	0.1485	19.9%

6.9 Generalizability: Temperature and Instance Size

Ambient temperature. Table 14 and Figure 7(b) sweep the design ambient from 35 to 52 °C. The energy overhead rises linearly from 6.92% to 12.80%, and SC4 mission energy from 1,316 to 1,389 Wh a 5.5% span across the operational envelope. Since all missions require battery exchange at the hub regardless (SC4 energy exceeds the 424.9 Wh single-charge budget by more than threefold), the practical effect of ambient temperature is on the number of exchanges rather than on route choice.

Table 14. Ambient temperature sweep (SC4, T-PSO).

Ambient (°C)	Overhead	SC4 energy (Wh)
35	6.92%	1,316
40	8.65%	1,338

Ambient (°C)	Overhead	SC4 energy (Wh)
44	10.03%	1,355
48 (design point)	11.42%	1,372
52	12.80%	1,389

Instance size. Table 15 and Figure 8 sweep instance size independently of the scenario definitions, using randomly drawn node subsets. The pattern of (§6.2) is confirmed and sharpened: at 5 nodes all algorithms tie; at 10 nodes the ordering is essentially noise (WOA nominally best); T-PSO takes the lead from 15 nodes upward, and its margin widens monotonically with size 5.3% over the runner-up at n = 15 and 20.1% at n = 20. This is the clearest evidence that T-PSO's advantage is a property of problem hardness rather than of the particular Baghdad topology.

Table 15. Mean $F(r)$ by instance size (random node subsets; 3 weight profiles × 30 seeds).

Nodes	PSO	GA	ACO	GWO	WOA	T-PSO	Best
5	0.1041	0.1049	0.1062	0.1041	0.1041	0.1041	tie
10	0.0862	0.0856	0.0879	0.0806	**0.0791**	0.0814	WOA
15	0.1087	0.0807	0.0874	0.0818	0.0853	**0.0764**	T-PSO
20	0.1821	0.0949	0.0931	0.1204	0.1239	**0.0743**	T-PSO

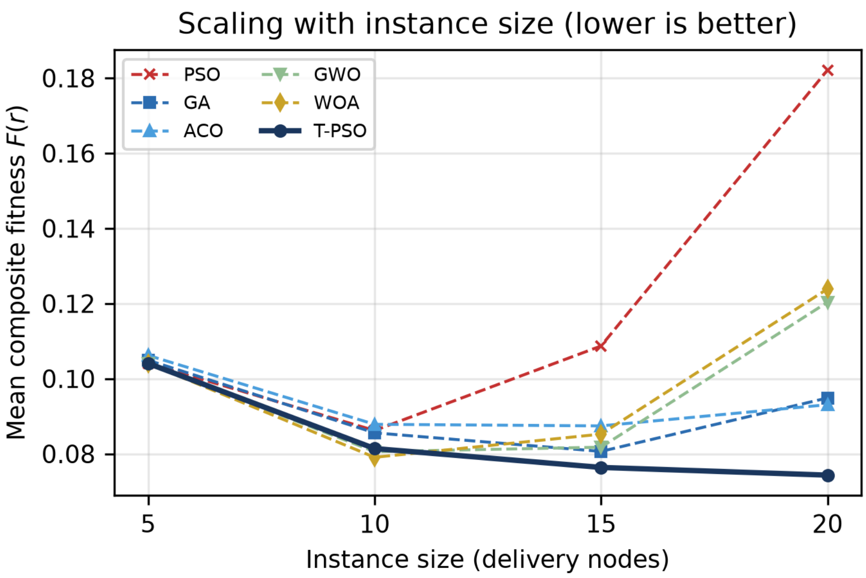

Figure 8. Scaling with instance size. The algorithms are indistinguishable to 10 nodes; T-PSO separates from 15 nodes upward and its margin widens with size.

6.10 Spatially Heterogeneous Thermal Field

Section 6.7 established that a uniform thermal correction cannot affect route construction. To test the condition under which it can, we impose a spatially varying ambient combining a 5 °C urban heat-island gradient across the density tiers with a 3 °C river-corridor cooling credit on cross-river arcs. Under this field the energy term is no longer collinear with time, so Equation (7) becomes genuinely multi-objective.

Two findings emerge (Table 16). First, the full memetic T-PSO is now clearly best (0.0250), ahead of 2-opt alone (0.0293) the reverse of the uniform-field ordering, because the richer objective landscape rewards the diversity of a seeded swarm. Second, and against expectation, the thermally-greedy seed remains identical to the distance-greedy seed: at a physically plausible 5 °C spread the thermal factor varies by only $\alpha_T \times 5 = 1.7\%$ across arcs, which is too small to reorder a nearest-neighbour tour whose arc lengths differ by far more. **Thermal-aware route construction therefore requires either a much steeper thermal gradient than urban heat islands produce, or a mechanism other than multiplicative arc-cost weighting** for example a route-dependent payload-integrity term, which (§7) identifies as the priority extension.

Table 16. SC4 under the heterogeneous thermal field (n = 180 per variant).

Variant	Mean $F(r)$	Energy (Wh)	Δ vs base
PSO (base)	0.0477	2,019	
PSO + distance-greedy	0.0459	2,008	−3.9%
PSO + thermal-greedy	0.0459	2,008	−3.9% (identical seed)
PSO + 2-opt	0.0293	2,019	−38.7%
T-PSO (greedy + 2-opt)	**0.0250**	2,024	**−47.6%**

7. Discussion

Four findings stand out. First, the memetic T-PSO delivers a statistically significant improvement in composite route quality over five baselines   including the recent GWO and WOA metaheuristics   under an equal evaluation budget. Second, algorithm differentiation grows with problem complexity: the six methods are indistinguishable on one-to-five-node instances but separate decisively from 15 nodes upward, where T-PSO leads by 5–20%. This argues for a capable optimizer precisely where hub-and-spoke coverage is hardest. Third, the ablation attributes that advantage almost entirely to 2-opt local search, and shows that greedy seeding adds nothing once local search is present. Fourth, the cargo-class admissibility outcome confirms that, under peak summer conditions with a passive container, cargo class   not route design   is the dominant determinant of mission feasibility.

The third finding carries a methodological consequence we state plainly. A thermal correction applied as a common multiplier on every arc is invisible to any route-construction heuristic, because it does not change the relative ordering of arcs. This holds for T-PSO's thermal-greedy seed and equally for ACO's thermally corrected arc heuristic, and it persists under a realistic urban heat-island gradient, where the arc-to-arc variation is only about 1.7%. "Thermal-aware routing" is therefore a meaningful concept in this framework at the level of *evaluation*   energy budgeting, admissibility, dispatch arbitration   but not, as implemented here, at the level of *construction*. Reporting this bound is more useful than leaving the claim unqualified.

These results extend Hooshyar and Huang's (2023) review by identifying the absence of hot-climate thermal calibration as a methodological gap and addressing it with a derived, sensitivity-tested correction embedded in route evaluation. They are consistent with Meng et al. (2024), whose payload-unconstrained UAV routing proves insufficient for hot-climate medical logistics. They are also directionally consistent with Li et al. (2025) in confirming that route design materially affects delivery time, energy, and composite fitness, though no direct numerical comparison is drawn given the different topology, objective, and thermal model.

Operational implication. For Baghdad summer operations, T-PSO is the recommended router for multi-node missions from roughly 15 nodes upward, while GA and ACO are acceptable simpler choices on small instances and PSO, GWO and WOA are unsuitable for full-network coverage. The 40% admissibility ceiling is a property of the modelled passive container shared by all algorithms: blood, plasma, and vaccines (M1, M2, M4) require active refrigeration, phase-change-material shippers, or cooler-hour scheduling, whereas platelets and pharmaceuticals (M3, M5) are deliverable as-is. Full SC4 coverage requires battery exchange at the H00 hub under every ambient condition tested.

Limitations. We separate these into model boundaries, methodological limits, and the validation gap.

Model boundaries. The container is modelled as passive and uncooled, with no phase-change material and no active refrigeration. This boundary   not a physical impossibility   is what makes M1, M2 and M4 inadmissible; a phase-change or actively cooled shipper within the platform's payload budget would change that outcome, and quantifying the required coolant mass against payload capacity is a direct extension of Equation (5). The energy model assumes constant speed, constant payload mass, and induced-power dominance; the ground-risk term is simplified to a three-level ordinal scale rather than the full SORA Annex F calculation.

Methodological limits. Under the uniform thermal field the time and energy terms are collinear (§3.4), so the main benchmark is effectively bi-objective and w_3 is constant over the admissible set. Section 6.10 shows that a heterogeneous field separates the two terms, but the PII term remains a constant admissibility flag rather than a route-dependent objective. **The single most valuable extension is to make PII continuous**: integrating

Equation (5) along the route so that payload temperature depends on arc-by-arc flight time and local ambient would make w_3 an active term, reward thermal margin alongside speed, and because Table 3 shows M3 breaching at 35.5 min against SC4 mission times of 108–142 min would likely make admissibility route-dependent on exactly the instances where the algorithms differ. We did not implement this in the present revision because it changes the feasible set and therefore the entire results basis; it is specified here as the next study rather than claimed as a result.

Validation gap. All results are simulation-based. The thermal coefficient α_T is derived from the ISA air-density relation and bounded by sensitivity analysis, but is not calibrated against UAV battery telemetry at 48 °C; $\zeta = 0.87$ is taken from the literature rather than measured on the modelled pack. **We propose a concrete two-stage validation protocol.** Stage 1 (bench): discharge the 22,000 mAh 6S pack under a representative cruise load profile in a thermal chamber at 15, 30, 40 and 48 °C, measuring delivered Wh and terminal voltage to fit α_T and ζ directly, and instrument a candidate container with internal thermocouples under a 48 °C soak to measure τ against Table 2. Stage 2 (flight): fly the SC1 and SC3 routes with an instrumented payload during Baghdad summer conditions, logging GPS, current draw, and payload temperature, and compare measured against predicted energy and breach times. Stage 1 requires no flight authorization and would resolve the two dominant model uncertainties on its own.

8. Conclusion

This paper formulated hot-climate UAV medical routing as a four-objective problem coupling delivery time, a derived 48 °C energy correction, payload admissibility, and SORA v2.5 ground risk, and proposed T-PSO a memetic PSO with greedy initialization and 2-opt refinement together with a TA-IDE decision layer enforcing a hard admissibility constraint. Across 13,500 deterministically reproducible runs on a verified 21-node Baghdad network at 48 °C, T-PSO achieved the lowest composite fitness (0.1023) overall and in every weight profile, significantly outperforming GA, ACO, GWO, WOA, and PSO (Friedman $p = 2.99 \times 10^{-29}$), with its largest margin (15.7–57.0%) on the most demanding full-network scenario, while the TA-IDE rejected all 8,100 inadmissible missions with zero false acceptances. A component ablation attributes the advantage to 2-opt local search and establishes that a spatially uniform thermal correction cannot influence route construction a bound on the "thermal-aware routing" claim that we report rather than leave implicit. Sensitivity analyses over the thermal coefficient, container class, ambient temperature, instance size, and objective weights delimit the conditions under which these conclusions hold. The framework gives healthcare logistics operators an auditable, safety-aware routing basis for hot-climate UAV medical delivery, and a reproducible methodology for the empirical validation bench thermal calibration first, instrumented flight second that remains the necessary next step.

9.Declaration of Competing Interest

The authors declare that there are no conflicts of interest regarding the publication of this manuscript.

10.Funding Information

No funding was received from any financial organization to conduct this research.

11.Author Contributions

Both authors jointly proposed the research problem and defined the scope of the study. Author Mohammed Saber Yasir Al-Khashan developed the thermal-aware energy and payload models, implemented the T-PSO algorithm and the baseline metaheuristics, designed and executed the full-factorial experimental campaign, performed the statistical analysis, and prepared the figures, tables and the first draft of the manuscript. Author Abdul Hadi Mohammed Alaidi supervised the work, verified the problem formulation, the equal-evaluation-budget protocol and the ablation design, validated the results and the sensitivity analyses, and revised the manuscript critically for intellectual content. Both authors discussed the results and read and approved the final version of this paper.

12.Acknowledgments

The authors express their gratitude to Wasit University / College of Computer Science and Information Technology, Kut, Wasit, Iraq, for supporting this study.

Notation List

Symbol / Abbreviation	Meaning
UAV	Unmanned Aerial Vehicle
T-PSO	Thermal-Aware Particle Swarm Optimization
TA-IDE	Thermal-Aware Intelligent Decision Engine
PSO / GA / ACO	Particle Swarm Optimization / Genetic Algorithm / Ant Colony Optimization
GWO / WOA	Grey Wolf Optimizer / Whale Optimization Algorithm
PII	Payload Integrity Index (binary admissibility flag)
GRC	Ground Risk Class (JARUS SORA v2.5)
SORA	Specific Operations Risk Assessment
BVLOS	Beyond Visual Line of Sight
WHO	World Health Organization
ISA	International Standard Atmosphere
F(r)	Composite fitness of route r
w1–w4	Objective weights (time, energy, payload, ground risk)
T	Ambient temperature (°C)
Tref	ISA reference temperature, 15 °C (288.15 K)
TK, Tref,K	Ambient and reference temperature in kelvin
Tpayload	Payload temperature (°C)
Tb	Regulatory upper bound of the cargo (°C)
αT	Thermal degradation coefficient ($°C^{-1}$)
ζ	Battery discharge efficiency at 48 °C
ρ, ρref	Air density and ISA reference air density ($kg\ m^{-3}$)
Ebase, Ethermal	Baseline and thermally corrected flight energy (Wh)
Eusable	Usable pack energy (Wh)
Q, Vnom	Battery pack capacity (Ah) and nominal voltage (V)
m, c	Payload mass (kg) and specific heat capacity ($J\ kg^{-1}\ K^{-1}$)
UA	Overall heat-transfer conductance of the container ($W\ K^{-1}$)
τ	Container thermal time constant (min)
tbreach	Time to cold-chain breach (min)
k, L, Am	Wall thermal conductivity, wall thickness, mean wall area
d(r)	Tour length of route r (km)
Neval	Fitness-evaluation budget per run
δ	Cliff's Delta effect size
αadj	Bonferroni-adjusted significance level
SC1–SC5	Delivery scenarios
WP1–WP3	Objective weight profiles
M1–M5	Mission (cargo) types

References

[1] World Health Organization, *WHO Expert Committee on Specifications for Pharmaceutical Preparations: Forty-fifth Report, WHO Technical Report Series No. 961*. Geneva, Switzerland: World Health Organization, 2011.

[2] D. Dissanayaka, T. R. Wanasinghe, O. De Silva, A. Jayasiri, and G. K. I. Mann, "Review of navigation methods for UAV-based parcel delivery," *IEEE Transactions on Automation Science and Engineering*, vol. 21, no. 1, pp. 1068–1082, 2024.

[3] M. Hooshyar and Y.-M. Huang, "Meta-heuristic algorithms in UAV path planning optimization: A systematic review (2018–2022)," *Drones*, vol. 7, no. 12, p. 687, 2023.

[4] Z. Al Abadla, M. Ahmed, H. Aldashti, and M. M. Abdel Wahab, "Assessing past and future heat and cold waves over Iraq using CMIP6 model," *Iraqi Journal of Science*, vol. 65, no. 8, pp. 4642–4655, 2024.

[5] S. A. Salman et al., “Climate change in Iraq: Temperature trend analysis based on the NASA POWER project data,” *Theoretical and Applied Climatology*, 2025, doi: 10.1007/s00704-025-05974-2.

[6] Joint Authorities for Rulemaking on Unmanned Systems, *JARUS Guidelines on SORA (Specific Operations Risk Assessment), Edition 2.5*. JARUS, 2024.

[7] V. Homier, D. Brouard, M. Nolan, et al., “Drone versus ground delivery of simulated blood products to an urban trauma center: The Montreal Medi-Drone pilot study,” *Journal of Trauma and Acute Care Surgery*, vol. 90, no. 3, pp. 515–521, 2021.

[8] S. Molinari, R. Patriarca, and M. Ducci, “The challenges of blood sample delivery via drones in urban environment: A feasibility study through specific operation risk assessment methodology,” *Drones*, vol. 8, no. 5, p. 210, 2024.

[9] A. Oakey, M. Grote, A. Pilko, et al., “Comparative modelling of drone interventions with traditional logistics modes for expedited and equitable deliveries,” *npj Sustainable Mobility and Transport*, vol. 2, no. 1, p. 37, 2025.

[10] H. A. Campbell, V. Bosiocic, A. Hvala, M. Brady, M. A. Campbell, K. Skelton, and O. J. Luiz, “Emerging research topics in drone healthcare delivery,” *Drones*, vol. 8, no. 6, p. 258, 2024.

[11] P. E. Hart, N. J. Nilsson, and B. Raphael, “A formal basis for the heuristic determination of minimum cost paths,” *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

[12] J. Kennedy and R. Eberhart, “Particle swarm optimization,” in *Proceedings of ICNN’95 – International Conference on Neural Networks*, 1995, vol. 4, pp. 1942–1948.

[13] A. Colorni, M. Dorigo, and V. Maniezzo, “Distributed optimization by ant colonies,” in *Proceedings of the First European Conference on Artificial Life*, 1991, pp. 134–142.

[14] X. Fan, X. Luo, S. Yi, S. Yang, and H. Zhang, “Optimal path planning for mobile robots based on intensified ant colony optimization algorithm,” in *Proceedings of the IEEE International Conference on Robotics, Intelligent Systems and Signal Processing*, 2003, vol. 1, pp. 131–136.

[15] S. Mirjalili, S. M. Mirjalili, and A. Lewis, “Grey wolf optimizer,” *Advances in Engineering Software*, vol. 69, pp. 46–61, 2014.

[16] S. Mirjalili and A. Lewis, “The whale optimization algorithm,” *Advances in Engineering Software*, vol. 95, pp. 51–67, 2016.

[17] Z. Yu, Z. Si, X. Li, D. Wang, and H. Song, “A novel hybrid particle swarm optimization algorithm for path planning of UAVs,” *IEEE Internet of Things Journal*, vol. 9, no. 22, pp. 22547–22558, 2022.

[18] L. Zhang, Y. Zhang, and Y. Li, “Mobile robot path planning based on improved localized particle swarm optimization,” *IEEE Sensors Journal*, vol. 21, no. 5, pp. 6962–6972, 2021.

[19] L. Li, H. Peng, X. Ren, and Q. Wang, “UAV logistics and distribution path planning in urban areas based on improved PSO and A* algorithms,” *IEIE Transactions on Smart Processing & Computing*, vol. 14, no. 2, pp. 257–267, 2025.

[20] T. Bartosik et al., “Validating lithium-polymer battery discharge models to ensure UAV flight safety and performance,” *Transactions on Aerospace Research*, vol. 2024, no. 4, 2024, doi: 10.2478/tar-2024-0024.

[21] Y. Wang et al., “Study on the environmental adaptability of lithium-ion battery powered UAV under extreme temperature conditions,” *Energy*, vol. 219, p. 119481, 2021.

[22] Z. Meng, K. Yu, and R. Qiu, “Location-routing optimization of UAV collaborative blood delivery vehicle distribution on complex roads,” *Complex & Intelligent Systems*, vol. 10, no. 6, pp. 8127–8141, 2024.