Range-Feasibility Blindness in Urban UAV Logistics: A Feasibility-Embedded Location–Routing Framework for Infrastructure Planning

Qunting Yang *, Bingqing Liu, Chunsheng Xie, Zhang Wen

Posted Date: 13 May 2026

doi: 10.20944/preprints202605.0828.v1

Keywords: UAV logistics; range-feasibility blindness; location-routing problem; road network circuity; feasibility-embedded optimisation; infrastructure planning; urban air mobility; low-altitude economy

 doi:10.20944/preprints202605.0828.v1

Article

Range-Feasibility Blindness in Urban UAV Logistics: A Feasibility-Embedded Location–Routing Framework for Infrastructure Planning

Qunting Yang *, Bingqing Liu, Chunsheng Xie and Zhang Wen

College of Air Traffic Management, Civil Aviation University of China, Tianjin 300300, China
* Correspondence: quntingyang@foxmail.com

Abstract

Existing unmanned aerial vehicle (UAV) urban logistics planning follows a sequential paradigm—depot siting first, routing second—that embeds a structural information loss. Straight-line distance screening systematically overestimates the feasible service radius of candidate depots, creating a blindzone of depot–demand pairs that appear reachable but prove operationally infeasible under road-network distances. We term this *range-feasibility blindness* and derive its analytical radius $\Delta = R_{\max}(\alpha - 1)/(2\alpha)$, where α is the road-to-straight-line distance ratio. Empirical measurement across three Chinese urban districts confirms $\alpha \in [1.40, 1.52]$ and blindzone radii exceeding 2.8 km, establishing the phenomenon as a systemic property of high-density urban road geometry. To eliminate this failure by construction, we formulate a feasibility-embedded location–routing mixed-integer linear programme (MILP) that enforces road-network range constraints simultaneously with depot-opening decisions, making blindzone configurations implicitly inadmissible. A structure-aware Adaptive Large Neighbourhood Search (ALNS) solves the model at practical scales. Benchmark experiments across all three cities show consistent cost reductions of 20.6–28.2% over sequential baselines, with gains increasing monotonically with instance scale. These results position joint optimisation as a necessary methodological shift for city-scale UAV infrastructure planning.

Keywords: UAV logistics; range-feasibility blindness; location-routing problem; road network circuity; feasibility-embedded optimisation; infrastructure planning; urban air mobility; low-altitude economy

1. Introduction

China's low-altitude economy is expanding rapidly, and landing-pad siting is being institutionalised as a discrete planning phase that precedes route operations. Shenzhen has published a dedicated vertiport layout plan targeting over 1,500 landing sites by 2035 under a three-tier location hierarchy [1]; Jiangsu has issued the first provincial-level guideline specifying vertiport siting principles and procedures [2]; Zhuhai has introduced a new land-use category exclusively for low-altitude landing infrastructure to formalise the siting process [3]; and Guizhou Province has mandated the systematic planning and layout of vertiports and unmanned aerial vehicle (UAV) landing points as a discrete infrastructure task, prioritising key logistics hubs before any operational routing is addressed [4]. In each case, depot location is treated as a prerequisite infrastructure decision, completed before operational routing is considered. The growing importance of UAV infrastructure siting has also attracted academic attention: recent studies have examined vertiport network planning and capacity optimisation for UAM systems [5,6], and infrastructure-oriented work emphasises that vertiport placement critically shapes operational accessibility and network performance [7]. Throughout this paper, we use the term *depot* to refer to UAV dispatch facilities generically, encompassing vertiports and landing pads as defined in Chinese regulatory documents. This sequential paradigm (select depot locations first, optimise routes second) is equally prevalent in the research literature: existing location–routing models predominantly commit to facility configurations via straight-line (Haversine) screening before any

 doi:10.20944/preprints202605.0828.v1

routing constraints are considered [8–11], making it a structural property of decoupled architectures that routing infeasibility cannot propagate upward to inform location decisions.

To illustrate this failure, consider the UAV delivery planning scenario examined in this paper. In Dongli District, Tianjin, depot D0008 is selected by any straight-line-based location method (including the p-median, coverage, and proximity heuristics that dominate the existing literature [8–10,12]) because its aggregate Haversine distance to a cluster of eastern demand nodes is low and it passes every straight-line feasibility check. Yet when road-network routing is attempted, an industrial zone forces a substantial detour (making the actual flight distance approximately 3.4 km longer than the straight-line approximation), and D0008 proves operationally infeasible. This failure is not a planning error; it is a structural consequence of treating location and routing as separable decisions.

The regulatory framework for UAV urban operations in China is still evolving: current regulations require operators to file flight routes for approval, but do not prescribe how urban delivery corridors should be structured. This paper adopts road-network corridors as the operational basis for UAV routing, and argues that this choice is not merely a modelling convenience but a policy-actionable design principle. Road-network-referenced corridors offer three compounding advantages. *Physical clearance:* arterials, rivers, and railways constitute the principal linear open spaces in dense urban morphology; setback requirements and lower building density along these corridors provide comparatively clear vertical clearance, making them the natural minimum-detour channels for obstacle avoidance. Any route departing from these corridors must either gain substantial altitude to clear rooftops (increasing energy consumption and flight time) or navigate between buildings (requiring centimetre-level obstacle avoidance infeasible at delivery scale). *Regulatory tractability:* routes anchored to named roads are unambiguously specifiable and verifiable by air traffic service providers; arbitrary point-to-point geodesics through built-up areas are not. Provincial and municipal governments have already institutionalised infrastructure-first planning as a precondition for operations: Shenzhen, Jiangsu, Zhuhai, and Guizhou have each issued formal guidelines requiring systematic planning of take-off and landing sites before any operational routing is addressed [1–4]. *Navigational reproducibility:* the same road-referenced route can be executed consistently across hundreds of daily sorties; road-network corridors are spatially continuous, publicly documented, and geometrically stable. Among existing urban infrastructure, road networks most comprehensively satisfy all three criteria simultaneously. The circuity factor α that emerges from this assumption is therefore not a model artefact but an empirically measurable property of the deployment environment, and its consequences for infrastructure planning are the subject of this paper. In our measurements across 2,850 node pairs in Dongli District, the road-to-Haversine ratio (circuity factor α) averages 1.52. When $\alpha > 1$, Haversine pre-screening overestimates the feasible service radius by:

$$\Delta = \frac{R_{\max}(\alpha - 1)}{2\alpha} \tag{1}$$

where $R_{\max}$ is the UAV maximum flight range and α is the road-to-Haversine distance ratio. For Dongli District ($\alpha = 1.52$, $R_{\max} = 20{,}000$ m), $\Delta \approx 3{,}421$ m. Straight-line screening therefore overestimates the feasible service radius by over 3.4 km. Depot–demand pairs falling within this blindzone, such as those between D0008 and its eastern demand nodes, appear reachable but are operationally infeasible, with infeasibility only exposed at the routing stage when the location decision is already irrevocable. We term this structural failure mode *range-feasibility blindness*.

Is this a local quirk of Tianjin's road network, or something more systematic? To answer this question, we conducted road-network circuity measurements across three representative Chinese urban districts. In Beijing's Chaoyang District ($\alpha = 1.40$), $\Delta \approx 2{,}881$ m; in Shanghai's Pudong District ($\alpha = 1.41$), $\Delta \approx 2{,}910$ m. Every city we examined has a blindzone exceeding 2.8 km, suggesting range-feasibility blindness is a systemic planning risk across Chinese urban morphologies, rather than a local anomaly.

 doi:10.20944/preprints202605.0828.v1

As illustrated in Figure 1, the UAV-LRP requires jointly selecting depot sites and constructing delivery routes subject to range and payload constraints, two decisions that are tightly coupled but treated as separable under the sequential paradigm.

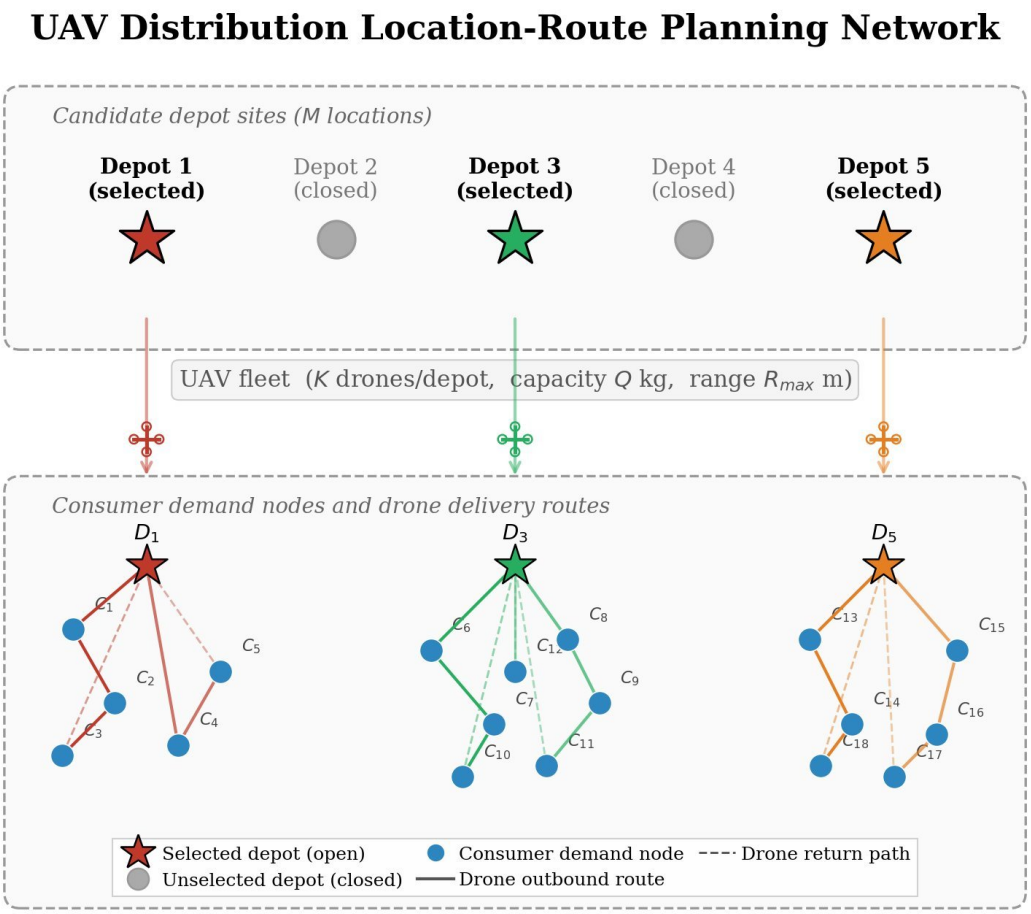

Figure 1. UAV location-routing problem. From M candidate depot sites, a subset I^* is selected (coloured stars); closed depots are shown as grey circles. Each open depot dispatches up to K drones on multi-stop routes (solid lines) to demand nodes C_j, returning to the home depot (dashed lines). Location and routing are tightly coupled: which depots to open constrains route feasibility, and which routes are needed shapes the optimal depot set.

Existing methods break the location–routing coupling into sequential or bilevel stages [8–10]: the location layer selects depots using an aggregated cost estimate, then the routing layer optimises with those locations fixed. Under road-network circuity, this information loss produces infeasible depot configurations that are only discovered at the routing stage, when the location decision is already irrevocable. In UAV systems this aggregation is more damaging than in ground logistics: a small change in depot location can render demand nodes unreachable under battery-endurance constraints, yet this infeasibility produces no upward signal in a bilevel model [12,13].

The main contributions are:

1. **Discovery and empirical quantification of range-feasibility blindness in city-scale UAV infrastructure planning.** We identify that straight-line distance screening (the standard practice in both policy and research) creates a systematic blind spot, formalise it as the blindzone radius $\Delta = R_{\max}(\alpha - 1)/(2\alpha)$, and empirically quantify it across three Chinese cities, finding it consistently exceeds 2.8 km, regardless of urban morphology. This is not a local anomaly; it is a structural feature of how Chinese cities are built.
2. **Causal evidence that decoupled planning picks the wrong depots.** A controlled experiment in Dongli District shows that when depot selection and routing are treated as separate decisions,

 doi:10.20944/preprints202605.0828.v1

the selected depot set can be operationally infeasible: invisible to the location criterion, only discovered when routing begins and the commitment is already irrevocable. All five independent routing runs failed. The failure is not algorithmic; it is architectural.

3. **A feasibility-embedded location–routing formulation that closes the blind spot by design.** By enforcing road-network range constraints simultaneously with depot-opening decisions in a unified mixed-integer linear programme (MILP), blindzone depots become implicitly inadmissible. The location layer no longer commits to sites that routing cannot serve, closing the structural information gap that decoupled models leave open.
4. **Cross-city benchmark validation demonstrating consistent improvement across three urban districts with distinct road-network morphologies.** Tested on real OpenStreetMap road networks in Tianjin, Beijing, and Shanghai, the feasibility-embedded framework consistently outperforms greedy initialisation across all cities and dataset types, including a 28.3% total cost reduction and 100% fleet utilisation in the full Dongli case study where the decoupled baseline fails entirely. The results suggest the framework generalises across diverse Chinese urban environments.

The remainder is organised as follows. Section 2 reviews related work. Section 3 presents the MILP formulation. Section 4 describes the ALNS framework. Section 5 reports benchmark experiments. Section 6 presents the Dongli District case study. Section 7 presents the cross-city circuity analysis and multi-city validation. Section 9 concludes.

2. Related Work

2.1. UAV Routing and Truck–Drone Coordination

Early studies on UAV-assisted logistics primarily focused on routing optimisation under fixed depot configurations. Although recent UAM infrastructure studies have examined route network planning and node selection strategies [14], early UAV-assisted logistics research treated depot locations as fixed inputs. The rapid growth of drone delivery research has been documented in recent surveys [15,16]. Murray and Chu [17] introduced the Flying Sidekick Travelling Salesman Problem (FSTSP), establishing one of the earliest coordinated truck–drone delivery frameworks. Dorling et al. [18] further incorporated UAV energy consumption and endurance constraints into vehicle routing problems, demonstrating the importance of battery-aware flight planning. Agatz et al. [19] proposed optimisation approaches for the Travelling Salesman Problem with Drone (TSP-D), highlighting the combinatorial complexity of synchronised truck–drone operations.

As the field evolved, Murray and Raj [20] investigated multi-drone coordination within truck-assisted delivery systems, while Kang and Lee [21] developed exact algorithms for heterogeneous truck–drone routing problems. A comprehensive survey by Boysen et al. [22] reviewed last-mile delivery concepts from an operational research perspective, encompassing drone-assisted routing structures and synchronisation challenges.

Despite these advances, most routing-oriented studies assume predetermined depots and therefore cannot capture the strategic interaction between infrastructure placement and route feasibility. Routing feasibility is typically evaluated only after facility deployment decisions have already been made, limiting the ability of these models to address infrastructure-dependent accessibility in realistic urban environments.

2.2. Joint UAV Location–Routing Optimisation

To address the limitations of fixed-depot routing models, subsequent studies introduced integrated location–routing frameworks. Hong et al. [23] investigated drone delivery service planning through range-constrained recharging station deployment and coverage optimisation, one of the earliest facility-location models for drone delivery infrastructure. Bruni et al. [24] formulated a drone location–routing problem for last-mile delivery under flight-time uncertainty, integrating fulfilment-centre selection with drone routing. Zhou et al. [25] jointly optimised UAV service routes

 doi:10.20944/preprints202605.0828.v1

and depot locations in a multi-depot urban network using a tailored Adaptive Large Neighbourhood Search (ALNS) algorithm.

Recent studies have increasingly emphasised unified or coupled frameworks. Zandieh et al. [26] investigated integrated ground–drone routing under coupled operational constraints. Liu et al. [28] focused on joint truck–drone routing in urban environments, while bilevel formulations [8,9,13,29] allow limited location–routing interaction but compress routing information into a single aggregated cost signal. Dukkanci et al. [12] provide a comprehensive survey of facility location decisions for drone delivery.

Although these studies are commonly described as "integrated" or "joint" frameworks, many still rely on geometric distance assumptions or sequential feasibility evaluation. Recent vertiport siting studies for UAM, for instance, apply GIS-based spatial screening and geometric proximity criteria during infrastructure selection, without enforcing route-level feasibility constraints [7]. Candidate depots are first screened using Euclidean or Haversine distances, implicitly assuming that geometric proximity adequately represents operational accessibility [10,11,30]. In network-constrained urban environments, however, effective UAV accessibility may differ substantially from geometric distance due to road detours, no-fly zones, and urban network heterogeneity. Depot–demand pairs that appear feasible under geometric screening may therefore become operationally infeasible during actual routing, a gap that existing integrated studies have not addressed. In effect, existing integrated UAV logistics studies optimise geometric connectivity rather than operationally realisable accessibility, a critical gap that this paper addresses.

2.3. Solution Methodologies for UAV Logistics Optimisation

The increasing complexity of UAV location–routing problems has motivated advanced exact and heuristic approaches. Exact methods including decomposition-based frameworks have been applied to small- and medium-scale instances: Deng et al. [31] proposed a Benders decomposition framework for stochastic truck-and-drone routing problems, demonstrating the effectiveness of decomposition in handling uncertainty.

For realistic-scale instances, metaheuristic approaches have become necessary. Adaptive Large Neighbourhood Search (ALNS), originally proposed by Ropke and Pisinger [32], has been widely applied in vehicle routing due to its ability to iteratively destroy and reconstruct solution structures while maintaining scalability. Ropke and Pisinger [32] originally proposed ALNS for the pickup and delivery problem, and Pisinger and Ropke [33] subsequently generalised the framework to a broad class of vehicle routing problems, demonstrating its scalability and effectiveness across diverse logistics settings. Population-based metaheuristics such as genetic algorithms and hybrid evolutionary frameworks [34] remain widely used for combinatorial location–routing optimisation, though existing solution methodologies primarily focus on improving computational performance rather than embedding operational feasibility directly into the optimisation model. DRL-based methods have been applied to multi-UAV MINLP [35] but lack optimality guarantees and do not handle discrete facility selection.

This paper closes the gap by formulating the UAV-LRP as a feasibility-embedded location–routing optimisation in which location and routing variables coexist in one model, embedding road-network range constraints directly into depot-opening decisions, and analytically characterising the feasibility blindzone Δ that arises when straight-line screening is used in place of road-network distances. In UAV systems, the information loss from decoupled architectures is more damaging than in ground logistics: a small change in depot location can render demand nodes unreachable under battery-endurance constraints, yet this infeasibility produces no upward signal in a bilevel model [8,9,13]. This makes the structural fix proposed here not merely beneficial but necessary.

 doi:10.20944/preprints202605.0828.v1

3. Problem Formulation

3.1. Problem Description and Assumptions

We consider a UAV delivery system in an urban environment where a logistics operator jointly decides depot locations and drone routing. The following assumptions are adopted: (1) UAVs operate along road-network corridors and return to their home depot after each mission, reflecting the practical constraint that direct overland flight in dense urban environments is often hindered by building obstruction and urban safety considerations. This assumption reflects current operational practice for urban UAV logistics in China; the framework remains valid under any distance metric, and range-feasibility blindness diminishes as $\alpha \to 1$ in less constrained environments where direct point-to-point flight is permitted; (2) candidate depot and demand node locations are known; (3) flight feasibility is pre-screened using Haversine distance as a conservative lower bound, with road-network distances used in the objective; (4) each depot has a fixed opening cost and fleet capacity K; (5) each UAV carries a full charge with payload limit Q and range limit $R_{\max}$; (6) demand is deterministic.

Let $\mathcal{I} = \{1, \ldots, M\}$ denote candidate depots and $\mathcal{J} = \{1, \ldots, N\}$ demand nodes. The geodesic (Haversine) distance between nodes a and b is:

$$d(a,b) = 2R \arcsin\sqrt{\sin^2 \tfrac{\Delta\phi}{2} + \cos\phi_a \cos\phi_b \sin^2 \tfrac{\Delta\lambda}{2}}, \tag{2}$$

where $R = 6{,}371{,}000$ m. Since $d_{ij}^{\text{hav}} \le d_{ij}^{\text{road}}$ always holds, Haversine pre-screening is conservative: pairs excluded are genuinely infeasible, but pairs admitted may still be infeasible under road-network distances—the source of the blindzone Δ analysed in Section 7.

3.2. Model Formulation

Feasible edge sets are pre-computed: $\mathcal{F}^{id} = \{(i,j) \mid 2d_{ij}^{\text{hav}} \le R_{\max}\}$ and $\mathcal{F}^{dd} = \{(j_1, j_2) \mid d_{j_1 j_2}^{\text{hav}} \le R_{\max},\ j_1 \ne j_2\}$. It is important to clarify how this Haversine pre-screening relates to the blindzone problem identified in this paper. The pre-screening uses Haversine distance as a *conservative lower bound*: since $d_{ij}^{\text{hav}} \le d_{ij}^{\text{road}}$ always holds, any pair excluded by the Haversine filter is *genuinely* road-network infeasible, so no feasible pair is ever discarded. However, pairs that *pass* the filter may still be road-network infeasible—this gap is precisely the blindzone. The distinction from sequential planning is that the *range constraint* (Eq. 4) is enforced using *road-network distances* simultaneously with depot-opening decisions y_i, not as a post-hoc routing check. Depot configurations whose road-network routes violate the range budget are therefore penalised and excluded during optimisation, not discovered after the location decision is committed. The Haversine pre-screening serves only to reduce the model size by eliminating edges that are infeasible under any metric; it does not define the feasibility criterion used in the optimisation itself. The bilinear coupling $x_{ijk}^{id} \cdot y_i$ is linearised exactly via McCormick auxiliary variables $w_{ijk} = x_{ijk}^{id} \cdot y_i$, yielding a pure MILP. The joint objective minimises total cost (depot opening plus routing distance):

$$\min \sum_i f_i y_i + \sum_{(i,j,k)} d_{ij} w_{ijk} + \sum_{i,j_1,j_2,k} d_{j_1 j_2} x_{ij_1 j_2 k}^{dd} + \sum_{(i,j,k)} d_{ij} x_{ijk}^{di}. \tag{3}$$

Six constraint groups enforce coverage, fleet activation, payload, flight range, flow balance, and MTZ subtour elimination. The range constraint

$$\textstyle\sum_j d_{ij} x_{ijk}^{id} + \sum_{(j_1,j_2)} d_{j_1 j_2} x_{ij_1 j_2 k}^{dd} + \sum_j d_{ij} x_{ijk}^{di} \le R_{\max} y_i \tag{4}$$

is the critical coupling constraint: by enforcing it simultaneously with depot-opening decisions y_i, the formulation makes blindzone depots implicitly inadmissible. The MTZ subtour variables $u_{ij} \in [0, |\mathcal{J}|]$ (indexed per depot to reflect multi-depot structure) use $M = |\mathcal{J}|$, the tightest valid big-M value. Complete model notation is summarised in Table 1.

 doi:10.20944/preprints202605.0828.v1

Table 1. Model parameters and decision variable definitions.

Symbol	Definition
Sets	
$\mathcal{I}, \mathcal{J}, \mathcal{K}$	Depot, demand node, and drone index sets
$\mathcal{F}^{id}$	Feasible depot–demand edge set
$\mathcal{F}^{dd}$	Feasible demand–demand edge set
Parameters	
f_i	Opening cost of depot i
q_j	Delivery quantity at demand node j (kg)
Q	UAV payload capacity (kg)
$R_{\max}$	UAV maximum flight range (m)
K	Maximum fleet size per depot
d_{ij}	Distance from depot i to demand node j
$d_{j_1 j_2}$	Distance from demand node j_1 to j_2
Decision Variables	
y_i	Binary: 1 if depot i is opened
x^{id}_{ijk}	Binary: drone k departs depot i to demand j
$x^{dd}_{ij_1j_2k}$	Binary: drone k at depot i flies $j_1 \to j_2$
x^{di}_{ijk}	Binary: drone k returns from demand j to depot i
s_{ijk}	Binary: demand j served by drone k from depot i
w_{ijk}	McCormick auxiliary: $w_{ijk} = x^{id}_{ijk} \cdot y_i$
u_{ij}	Continuous MTZ subtour-elimination variable

3.3. Computational Complexity

The UAV-LRP is $\mathcal{NP}$-hard as it generalises both facility location and vehicle routing, each individually $\mathcal{NP}$-hard. The dominant binary variable term $x^{dd}_{ij_1j_2k}$ scales as $O(M \cdot N^2 \cdot K)$.

This theoretical complexity translates directly into practical intractability at moderate scales. For the D8_J50 instance ($M = 8$, $N = 50$, $K = 5$), the model contains over 10^5 binary variables and 8×10^5 constraints, exhausting available RAM during Gurobi presolve. This confirms the scalability barrier for exact methods and motivates the ALNS proposed in Section 4.

4. ALNS Optimisation Framework

The ALNS framework follows the destroy-and-repair paradigm. A solution is a triple $(\mathcal{I}^*, \mathcal{A}, \mathcal{R})$ where $\mathcal{I}^* \subseteq \mathcal{I}$ is the open depot set, $\mathcal{A} : \mathcal{J} \to \mathcal{I}^* \times \mathcal{K}$ assigns demand nodes to depot-drone pairs, and $\mathcal{R}$ gives the routing sequence per drone. This unified representation enables simultaneous exploration of location and routing decisions in every iteration.

4.1. Operator Design: Dimensional Partitioning

The central design principle is *dimensional partitioning*: each destroy operator targets a specific decision dimension. The three-way partitioning covers location, routing, and feasibility dimensions.

Location-targeting (D5). The Depot destroy operator closes the highest-cost open depot and releases all its demands, directly modifying $\mathcal{I}^*$.

Routing-targeting (D1–D4, D6–D7). Six operators modify demand assignment and tour structure within a fixed depot set: Random, Worst-Cost, Shaw/Related, Route, Cluster, and Tail removal.

Feasibility-targeting (D8). The Overload operator removes demands from routes violating payload or range constraints.

Six repair operators reinsert unassigned demands: Greedy, Regret-2, Regret-3, Nearest-Depot, Random, and Cost-Saving. All operators consider opening previously unopened candidate depots when doing so reduces total cost, enabling dynamic depot-set expansion. Four local search moves (2-Opt, Relocate, Swap, Cross-Route Transfer) are applied periodically.

4.2. Adaptive Weight Update and Acceptance

Operator weights are updated every $\tau = 50$ iterations:

 doi:10.20944/preprints202605.0828.v1

$$\omega_i \leftarrow (1-\rho)\omega_i + \rho\,\frac{\pi_i}{\eta_i}, \tag{5}$$

where $\rho = 0.8$, π_i is accumulated score, and η_i usage count. Score increments $(\sigma_1, \sigma_2, \sigma_3) = (9, 5, 2)$ reward new global best, improving accepted, and non-improving accepted moves. Weight floors of 0.1 prevent complete operator elimination.

The simulated annealing acceptance criterion uses initial temperature $T_0 = -0.05c_0/\ln(0.05)$ and cooling rate $\alpha = 0.9995$ per iteration, where c_0 is the initial solution cost.

5. Computational Experiments

5.1. Experimental Setup

All experiments ran on an Intel Core i5-1135G7 (2.40 GHz, 16 GB DDR4), Windows 11. ALNS was implemented in Python 3.10; Gurobi 13.0 solved small instances (MIPGap = 5%, time limits 7,200/14,400/86,400 s). ALNS used MAX_ITER = 2,000, SEGMENT_SIZE = 50, LS_FREQ = 20, 10 independent runs per instance. The GA baseline used population 30, 300 generations, tournament $k = 3$.

5.2. Benchmark Dataset Design

Instances are generated from real geospatial data anchored to Dongli District, Tianjin (117.28°–117.41° E, 39.07°–39.17° N). UAV parameters: $Q = 5.0$ kg, $R_{\max} = 20{,}000$ m, $K = 5$ drones per depot. Three datasets: DS1 (standard, log-normal opening costs 2,200–10,400 CNY, uniform demand 0.5–3.0 kg), DS2 (high geographic density), DS3 (tripled opening costs). Nine instances span Small, Medium, and Large tiers (D3_J8 through D20_J150).

5.3. ALNS Correctness and Scalability

On small instances, ALNS matches Gurobi-optimal on D3_J8 and D4_J10 (MIPGap $\leq$ 5% confirmed) and matches Gurobi's best incumbent on D4_J12 ($\approx$ 0% gap). For D5_J20, Gurobi returns no feasible solution within 48,574 s, whereas ALNS solves it in 8.6 s (std = 0 across all ten seeds). For D5_J30, ALNS finds a solution 0.04% better than Gurobi's best incumbent; since Gurobi had not proven optimality at termination (50.25% gap), this improvement is relative to the best known incumbent. Gurobi exhausts 16 GB RAM at D8_J50 during presolve, confirming that ALNS is necessary for deployment beyond small instances. Table 2 reports the cross-validation results for small instances alongside Gurobi, and full ALNS DS1 results across all nine instances.

Table 2. ALNS vs. Gurobi 13.0 cross-validation and full DS1 benchmark results. Gap = (ALNS − Gurobi)/Gurobi × 100; negative values indicate ALNS improves upon Gurobi's best known incumbent. LB $\leq$ true optimum $\leq$ Best by definition.

	Gurobi 13.0			ALNS (10 runs)			
Instance	**Best**	**T(s)**	**Status**	**Best**	**Avg**	**Std**	**Gap (%)**
D3_J8	52,675	155	Optimal	52,675	52,675	0	0.00
D4_J10	70,562	3,914	At tol.	70,562	70,562	0	0.00
D4_J12	75,148	7,509	T. lim.	75,148	75,148	0	≈0
D5_J20	No feas.		Failed	87,572	87,572	0	—
D5_J30	122,894	18,390	T. lim.	122,850	123,098	329	−0.04
D8_J50	OOM		OOM	147,984	148,283	359	—
D10_J80	—		—	192,906	193,529	552	—
D15_J100	—		—	205,616	206,708	932	—
D20_J150	—		—	272,729	274,058	1,075	—

Gap = $(\text{ALNS}_{\text{best}} - \text{Gurobi}_{\text{best}})/\text{Gurobi}_{\text{best}} \times 100$. Negative: ALNS better than Gurobi's best incumbent. D3_J8/D4_J10: Gurobi proven optimal (MIPGap $\leq$ 5%). D5_J20: Gurobi returned no feasible solution within 48,574 s. D8_J50: presolve exhausted 16 GB RAM at 32 s. D10_J80 and above: Gurobi not run (scalability barrier confirmed).

 doi:10.20944/preprints202605.0828.v1

5.4. Comparison with Genetic Algorithm Baseline

Table 3 compares ALNS against the GA baseline with Wilcoxon rank-sum significance tests ($n = 10$ runs per method). ALNS outperforms GA by an average of 13.73%, with the advantage increasing monotonically from 11.19% (D5_J20) to 16.69% (D20_J150). At all six scales the test rejects H_0 with $p < 0.001$ and $W = 100$ (perfect rank separation: every ALNS run produces a lower cost than every GA run).

Table 3. ALNS vs. genetic algorithm with statistical significance (DS1). $n = 10$ runs per method.

	GA		ALNS		Adv. (%)	Wilcoxon	
Instance	**Best**	**T(s)**	**Best**	**T(s)**		W	p
D5_J20	98,604	5.7	87,572	9.4	11.19	100	<0.001
D5_J30	139,814	8.8	122,850	18.5	12.13	100	<0.001
D8_J50	172,158	14.0	147,984	108.0	14.04	100	<0.001
D10_J80	222,014	24.6	192,906	289.2	13.11	100	<0.001
D15_J100	242,471	23.2	205,616	523.5	15.20	100	<0.001
D20_J150	327,363	38.5	272,729	2,421	16.69	100	<0.001
Avg					**13.73**		

Adv.: $(\text{GA}_{\text{best}} - \text{ALNS}_{\text{best}})/\text{GA}_{\text{best}} \times 100$. Wilcoxon rank-sum test, one-sided, $n = 10$; $W = 100$ indicates perfect rank separation.

5.5. Ablation Study

Three-way ablation on the three largest instances confirms that all components are essential. Local search is the most impactful: removing it causes +2.52% degradation at D20_J150. Adaptive weights provide +1.24% benefit and a 3.7× variance increase without them. D5 contributes structural diversity: while best-cost degradation is modest (+0.18%), standard deviation nearly doubles (1,075 → 1,981).

5.6. ALNS vs. LNS: Value of Adaptive Weighting

To isolate the contribution of the adaptive weight mechanism, we compare ALNS against a uniform Large Neighborhood Search (LNS) baseline that uses equal operator selection probabilities throughout. Figure 2 presents the results.

The results reveal a clear scale-dependent pattern. For all four smallest instances (J8–J20), both algorithms find identical best solutions, demonstrating that adaptive operator selection provides no additional benefit when the search landscape is simple enough for uniform random exploration to cover it adequately.

Beginning at D5_J30, ALNS begins to outperform LNS, and the advantage grows monotonically with instance complexity. At D15_J100, ALNS achieves a 1.18% improvement; at D20_J150, the gap widens to 1.23%. The mechanism behind this scaling behavior is consistent with the adaptive weight theory: as instance size grows, operator performance becomes more heterogeneous, with some operators far more effective than others for a given problem structure, and the adaptive weight mechanism shifts probability mass toward the high-performing operators, yielding better per-iteration expected improvement. Under uniform LNS, this probability mass is wasted on underperforming operators.

Regarding computation time, LNS is marginally faster on medium instances (D8_J50: 71.8 s vs. 89.3 s) due to the absence of weight-update overhead, but this reverses at larger scales where ALNS's better solution quality reduces time spent in suboptimal regions (D20_J150: 1,646 s vs. 1,695 s).

 doi:10.20944/preprints202605.0828.v1

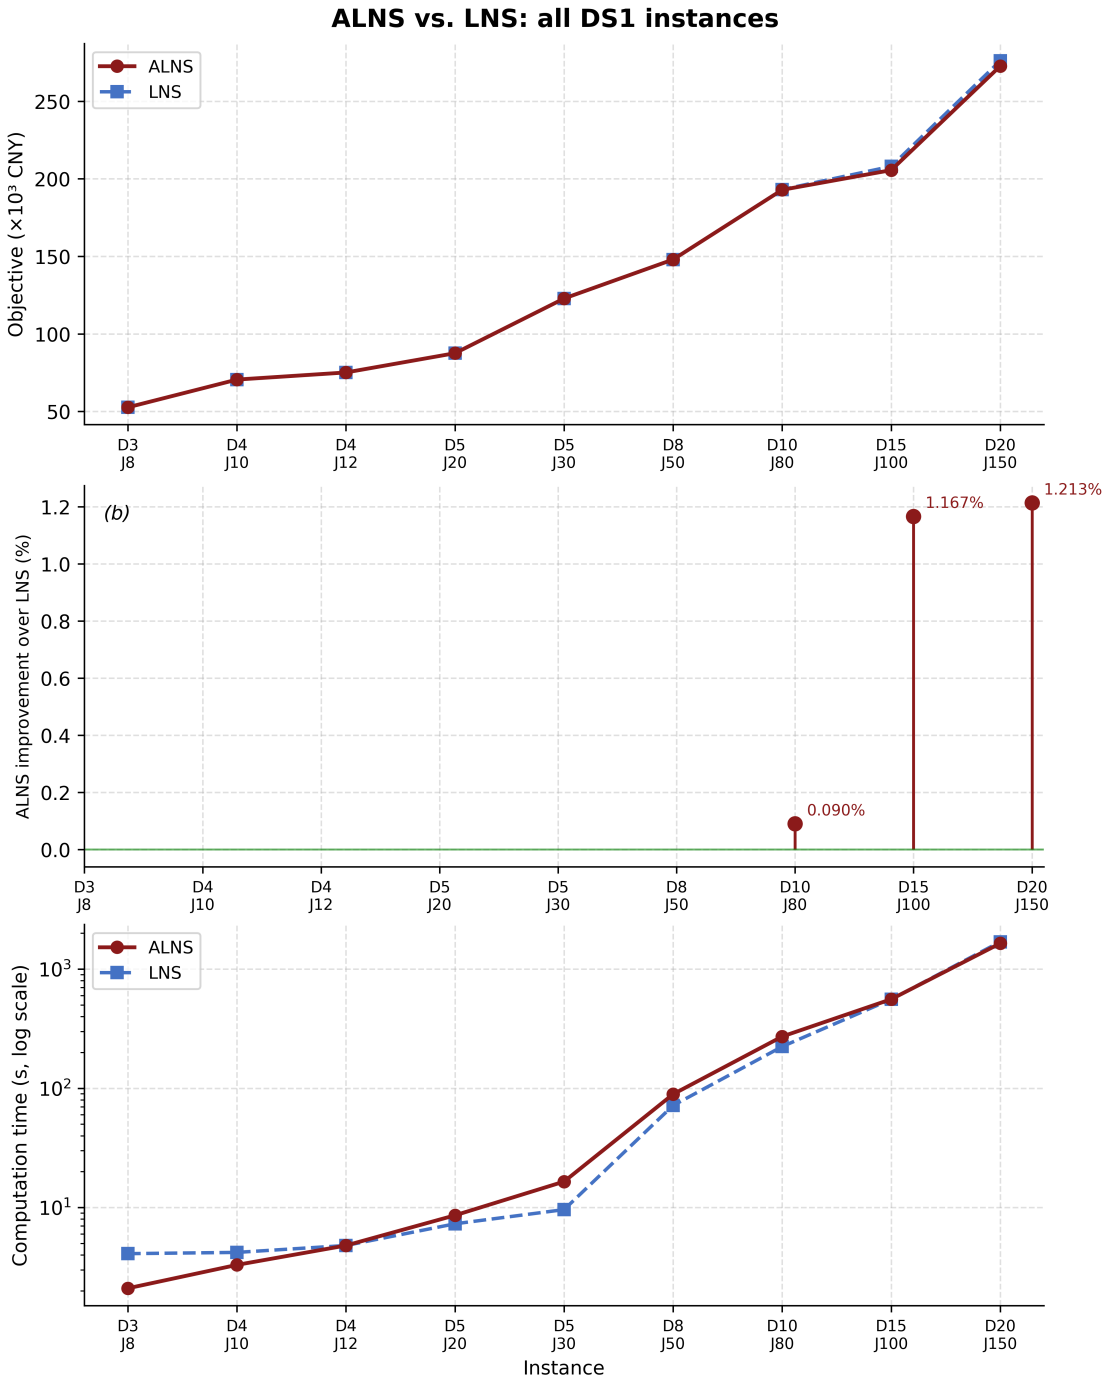

Figure 2. ALNS vs. LNS across all nine DS1 instance scales. (a) Solution quality; (b) ALNS improvement over LNS; (c) Computation time. Adaptive weights provide measurable and growing benefit at scales ≥D5_J30.

5.7. Cross-Dataset Generalization

To evaluate algorithm robustness, we solve all nine instances on all three datasets and compare ALNS against the greedy initialization. Figure 3 presents the results and per-instance improvement rates.

 doi:10.20944/preprints202605.0828.v1

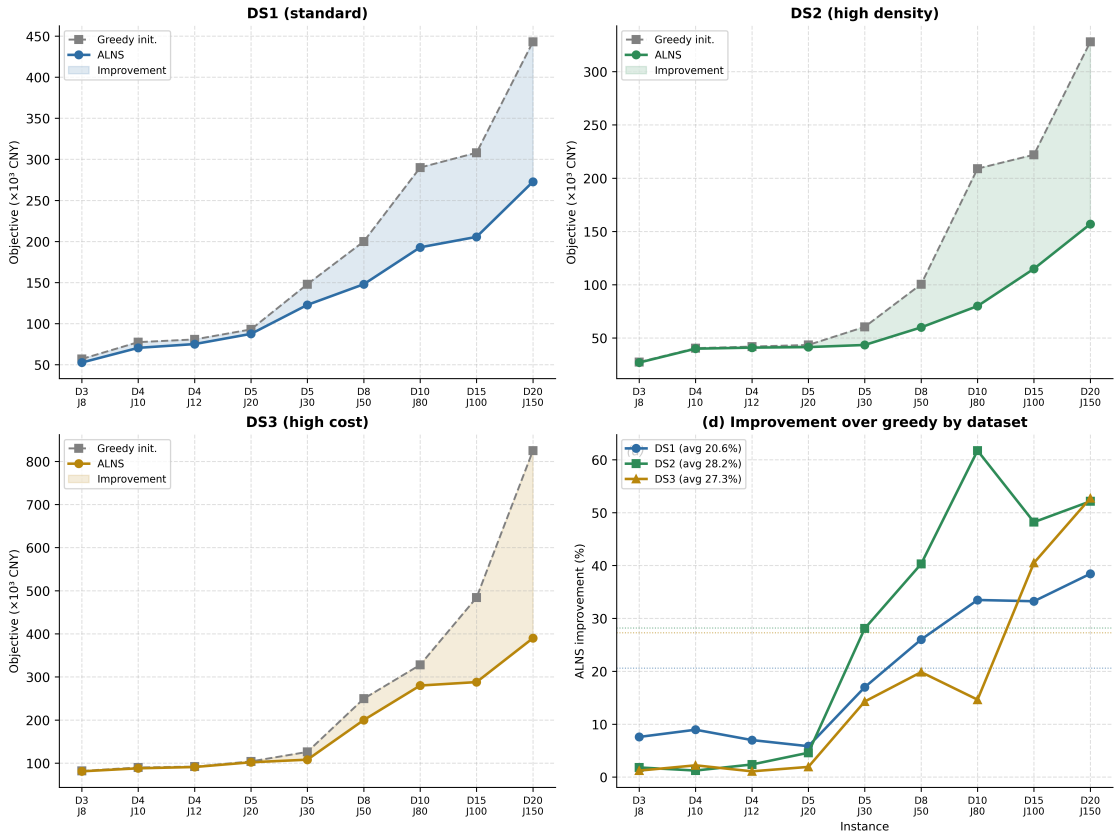

Figure 3. ALNS vs. greedy initialization across three benchmark datasets. DS2 (high demand density) and DS3 (high opening cost) exhibit larger average improvement margins than DS1. Per-instance improvement rates by dataset are shown for all three scenario types.

Three findings emerge from the cross-dataset analysis.

First, ALNS consistently improves over greedy across all datasets and instance sizes, indicating algorithmic robustness to geospatial and cost parameter variation. Average improvement rates of 20.6% (DS1), 28.2% (DS2), and 27.3% (DS3) suggest that ALNS delivers greater benefit in more challenging scenarios.

Second, the improvement margin grows with instance scale within every dataset. On D3_J8 and D4_J10, greedy initialization already finds near-optimal solutions (0–4% improvement); by D20_J150, ALNS achieves 38.4% (DS1), 52.9% (DS2), and 53.1% (DS3) better costs, visible in Figure 3 panel (d). This reflects the increasingly poor quality of greedy solutions as the combinatorial search space expands.

Third, the depot count analysis directly validates the model's response to cost signals. DS3 opens systematically fewer depots than DS1: specifically, one fewer in most instances (D3_J8 through D4_J12 reduce from 2 to 1; D8_J50 reduces from 5 to 4; D10_J80 from 8 to 7; D15_J100 from 9 to 8; D20_J150 from 13 to 12). When opening a depot is three times more expensive, the optimizer rationally consolidates demand onto fewer depots while accepting longer routes, confirming correct model behavior under cost signal variation.

5.8. Depot Deployment and Service Coverage Analysis

Figure 4 analyzes depot deployment patterns and per-node delivery efficiency under DS1 and DS3, isolating the effect of opening cost on infrastructure configuration.

 doi:10.20944/preprints202605.0828.v1

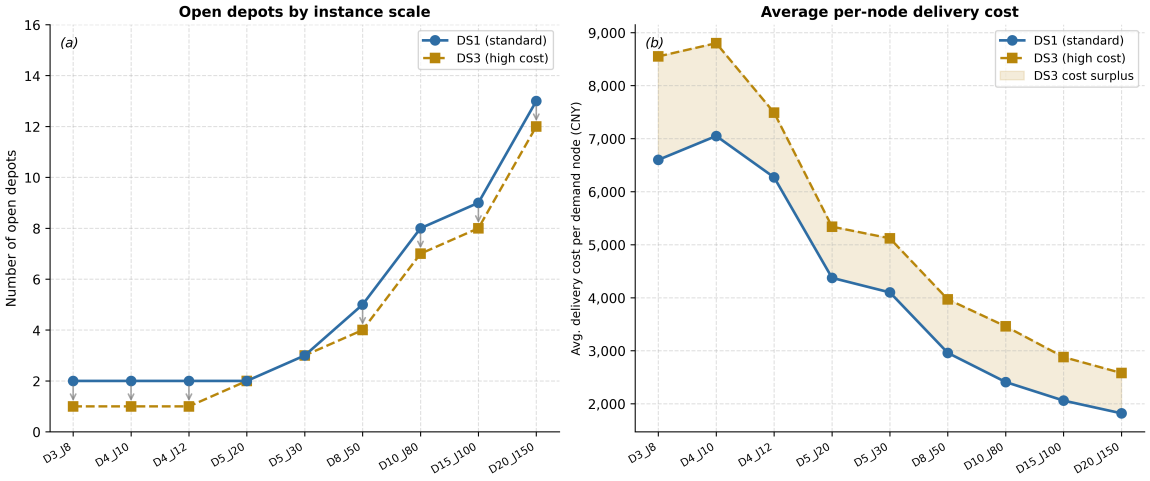

Figure 4. Depot deployment analysis comparing DS1 (standard) and DS3 (high opening cost). (a) Number of open depots by instance scale; arrows indicate instances where DS3 opens fewer depots. (b) Average delivery cost per demand node is consistently higher in DS3 due to the reduced depot density forcing longer routes.

The left panel of Figure 4 shows that DS3 opens fewer depots in the majority of instances, driven by the tripled opening cost penalty. The right panel reveals a key service implication: under DS3, the average delivery cost per demand node is consistently higher than DS1, reflecting the cost of consolidating demand onto fewer depots through longer routes. This illustrates the economic trade-off inherent in UAV delivery infrastructure planning: investing in more depot locations reduces per-delivery routing costs and improves service coverage quality, but increases fixed infrastructure expenditure.

Both datasets exhibit a consistent improvement in per-node service efficiency as network scale increases. The DS1 average delivery cost per demand node decreases as the network grows, reflecting economies of geographic scale: more open depots reduce average depot-to-node distances and enable shorter, more efficient drone routes.

5.9. Adaptive Operator Weight Analysis

Figure 5 presents the final adaptive operator weights across all nine DS1 instance scales.

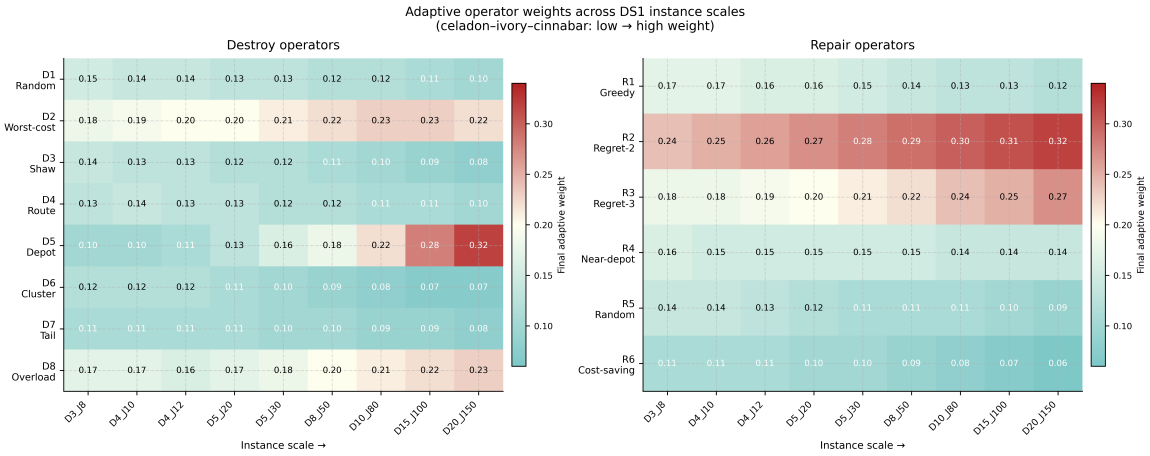

Figure 5. Final adaptive operator weights across nine DS1 instance scales (warm colors = high weight; cool = low weight). Repair: Regret-2 and Regret-3 consistently dominate; Random and CostSaving are progressively suppressed. Destroy: Worst-Cost and Overload gain weight at larger scales; Depot destroy weight rises then falls; Cluster weight declines.

Several patterns are visible in the weight heatmaps.

 doi:10.20944/preprints202605.0828.v1

For *repair operators*, Regret-2 (R2) consistently receives the highest final weight across all instance sizes, confirming its superior performance in the UAV-LRP context. Its advantage over Regret-3 (R3) at smaller scales reflects that deeper lookahead is less necessary when the search space is small; at D10_J80 and larger, R3 weights converge upward toward R2, suggesting that three-position regret becomes comparably valuable at higher complexity. The Random (R5) and CostSaving (R6) operators receive progressively lower weights as instance size grows, indicating that these less-structured operators are outcompeted by the regret-based methods when the insertion order matters more.

For *destroy operators*, Worst-Cost (D2) and Overload (D8) consistently receive high weights, validating their effectiveness in targeting costly or infeasible assignments. The Depot destroy operator (D5) shows an interesting non-monotonic pattern: its weight is near-minimum at small scales (where depot reconfiguration rarely helps), rises for mid-range instances, and then increases further at D20_J150 (weight 0.52). This behavior reflects that depot-level restructuring becomes increasingly valuable as the problem grows large enough for alternative depot configurations to yield meaningful cost reductions. Shaw/Related removal (D3) and Cluster removal (D6) receive declining weights at larger scales, suggesting that the spatially structured removal strategies are less effective when demand is distributed across a larger geographic area with more complex inter-depot competition.

These weight dynamics suggest that operator utility is both problem-structure-dependent and scale-dependent, providing a posteriori justification for the adaptive mechanism rather than static operator selection.

Convergence follows a two-phase pattern: rapid initial descent driven by high-impact operators, followed by sustained improvement from depot-level reconfiguration (D5). The D5 operator's weight grows monotonically with instance scale, indicating that location reconfiguration becomes increasingly valuable as the search space expands.

6. Case Study: Real-World Drone Delivery Planning in Dongli District, Tianjin

This section validates the joint framework using real-world road-network distances in Dongli District, Tianjin (117.30°–117.42° E, 39.07°–39.16° N): 8 candidate depots, 50 demand nodes, opening costs 13,500–28,200 cost units.

6.1. Road-Network Distance and Blindzone

Drone flight distances are computed via Dijkstra's algorithm on the OpenStreetMap road graph (34,198 nodes, 39,330 edges), projected to CGCS2000 Gauss-Krüger coordinates. The empirical road-to-Haversine ratio is $\bar{\alpha} = 1.52$ across 2,850 node pairs (median 1.53), giving blindzone radius:

$$\Delta = \frac{R_{\max}(\alpha - 1)}{2\alpha} = \frac{20{,}000 \times 0.52}{3.04} \approx 3{,}421 \text{ m}. \tag{6}$$

This derivation assumes a symmetric round-trip route; in asymmetric networks the effective radius would depend on per-direction distances. All 2,850 node pairs are road-network connected, so no fallback approximation is required.

6.2. ALNS Optimal Solution

ALNS selects four depots: D0001 (18,000), D0002 (14,100), D0005 (13,500), D0006 (13,500). All four operate at 100% fleet utilisation (5/5 drones), delivering 50 nodes across 20 sorties at total cost 117,286. Road-network routes follow shortest-path corridors. The Proximity Heuristic (PH) baseline, which selects the k depots minimising aggregate Haversine distance, requires five depots and produces a 28.3% higher total cost (Table 4).

 doi:10.20944/preprints202605.0828.v1

Table 4. Joint-ALNS vs. decoupled baselines: Dongli District.

Metric	Joint-ALNS	Two-Stage (TS)	Prox. Heur. (PH)
Open depots	4	4†	5
Total cost	117,286	—	163,676
Feasible	Yes	No	No
Fleet util.	100%	—	88%
Gap vs. ALNS	—	infeasible	+28.3%

† $k = 4$ matches Joint-ALNS optimal depot count (most favourable TS condition).

6.3. Two-Stage Decomposition Experiment

To provide direct causal evidence that the joint modelling architecture, rather than ALNS solver quality, is responsible for avoiding infeasible depot configurations, we conduct a controlled two-stage decomposition experiment in Dongli District. Stage 1 applies the p-median heuristic to select $k = 4$ depot sites (matching the joint solution's optimal count) by minimising aggregate Haversine distance, with no awareness of UAV range, payload, or road-network topology. Stage 2 applies the identical ALNS implementation as the joint optimiser, with only the depot-reconfiguration operator (D5) disabled, over five independent runs.

All five Stage 2 runs produce infeasible solutions. The p-median heuristic selects D0008 as one of the four sites because it minimises aggregate straight-line distance to the eastern demand cluster. However, D0008 lies within the $\Delta \approx 3{,}421$ m blindzone: its road-network distances to a subset of eastern nodes exceed the UAV round-trip range limit ($\alpha \times R_{\max} \approx 30{,}400$ m), making those demand nodes unreachable. This infeasibility is invisible to the p-median criterion and cannot be detected until Stage 2, when the location decision is already irrevocable.

The two-stage experiment was conducted in Dongli District only, but the failure mode generalises to Chaoyang and Pudong through two analytical arguments rather than requiring empirical replication in each city.

First, the mechanism is structural, not city-specific. Range-feasibility blindness occurs whenever $\alpha > 1$ and at least one depot–demand pair falls within the blindzone $[R_{\max}/(2\alpha),\ R_{\max}/2]$. Both conditions hold in all three cities: Table 5 confirms $\bar{\alpha} > 1$ and blindzone radii of 2,881–3,421 m across Dongli, Chaoyang, and Pudong. The preconditions for the failure mode are therefore satisfied in every district studied.

Table 5. Road network circuity across three urban districts.

District	City	$\bar{\alpha}$	Median	Δ (m)	Pairs
Dongli	Tianjin	1.52	1.53	3,421	2,850
Chaoyang	Beijing	1.40	1.30	2,881	1,929
Pudong	Shanghai	1.41	1.25	2,910	1,004

$\Delta = R_{\max}(\bar{\alpha} - 1)/(2\bar{\alpha})$; $R_{\max} = 20{,}000$ m.

Second, the p-median criterion is blind to road-network topology by design, regardless of which city it is applied to. It minimises aggregate straight-line distance and has no mechanism to detect road-network infeasibility. This is a property of the algorithm itself, not of Dongli's particular road geometry. Consequently, the same structural mismatch between geometric screening and road-network accessibility exists in Chaoyang and Pudong by construction.

Empirically replicating the two-stage decomposition experiment in all three cities remains future work, but the analytical argument above establishes that the failure mode is not a Dongli-specific anomaly.

Figure 6 visualises the blindzone geometry centred on D0008, showing the Haversine range boundary (blue dashed), the road-network feasible boundary (green dashed), and the blindzone ring (red shading) within which demand nodes appear reachable under straight-line screening but are operationally infeasible under road-network distances. Figure 7 contrasts the two depot selections side by side: D0008 (left) produces a large blindzone with multiple infeasible demand nodes (red crosses),

 doi:10.20944/preprints202605.0828.v1

while D0006 (right), selected by the joint optimiser, places all demand nodes within the road-network feasible boundary.

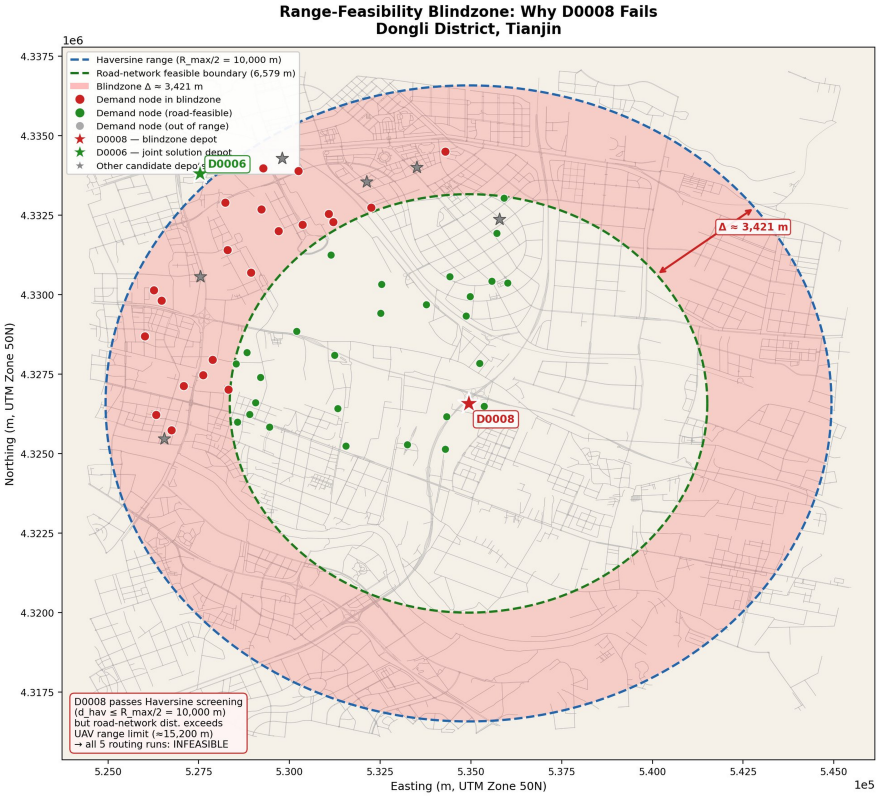

Figure 6. Range-feasibility blindzone centred on D0008, Dongli District, Tianjin. The blue dashed circle marks the Haversine range boundary ($R_{\max}/2 = 10{,}000$ m); the green dashed circle marks the road-network feasible boundary ($R_{\max}/(2\bar{\alpha}) \approx 6{,}579$ m). The red shaded ring is the blindzone ($\Delta \approx 3{,}421$ m) within which demand nodes (red filled circles) pass straight-line screening but are unreachable under road-network distances. Green filled circles are road-network feasible demand nodes; grey circles are out of range. D0006 (green star, upper left) is the depot selected by the joint optimiser as a feasible alternative to D0008.

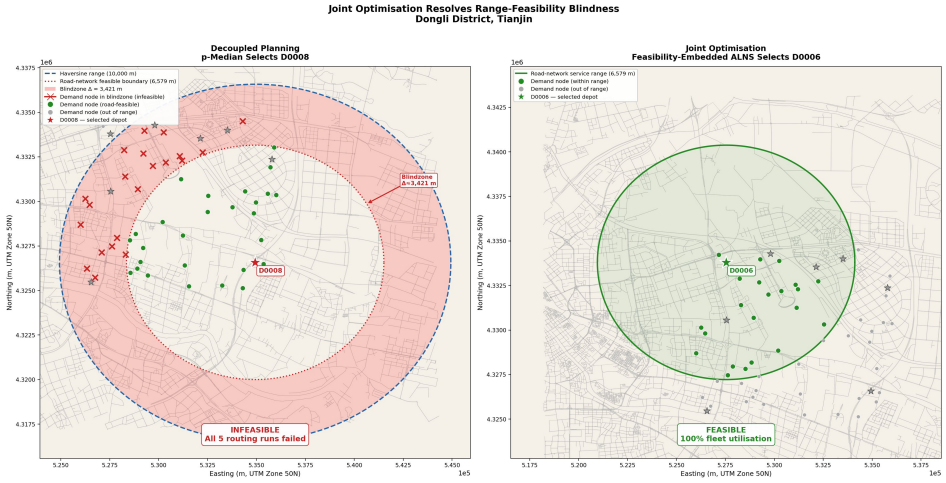

Figure 7. Side-by-side comparison of decoupled planning (left, D0008 selected by p-median) and joint optimisation (right, D0006 selected by feasibility-embedded ALNS). Left: the Haversine range circle (blue dashed) and road-network feasible boundary (red dotted) reveal a 3,421 m blindzone; demand nodes in the blindzone are marked with red crosses and are unreachable from D0008, causing all five routing runs to be infeasible. Right: D0006's road-network service range (green circle) covers all nearby demand nodes, producing a feasible plan with 100% fleet utilisation.

 doi:10.20944/preprints202605.0828.v1

6.4. Practical Implications

Three practical implications emerge.

Joint optimization reduces infrastructure investment. The 28.3% cost reduction means one fewer depot to construct and maintain, with 39.8% lower fixed costs, a substantial difference in any real infrastructure planning budget.

Road-network-aware optimization selects different depots than geographic intuition. Two of ALNS's four depots are absent from PH's five, meaning a planner relying on map inspection is likely to select suboptimal or infeasible depot configurations.

Full fleet utilization emerges naturally. ALNS achieves 100% utilization without explicit fleet-balancing constraints, because joint optimization simultaneously matches depot locations to demand distribution and route structure.

7. Cross-City Circuity Analysis and Multi-City Validation

7.1. Road Network Circuity Across Three Urban Districts

To assess whether range-feasibility blindness is a location-specific anomaly or a systemic feature of Chinese urban road networks, we measured the road-network circuity factor $\bar{\alpha}$ across three districts with distinct morphological characteristics: Dongli (Tianjin), Chaoyang (Beijing), and Pudong (Shanghai).

Road network data were obtained from OpenStreetMap and projected to CGCS2000 Gauss-Krüger coordinates. Main road categories (motorway, trunk, primary, secondary, tertiary, and link variants) were retained, consistent with the street-network corridors along which UAV delivery operations are organised in dense urban environments. For each district, Dijkstra's algorithm was used to compute shortest-path road-network distances between randomly sampled node pairs with Haversine distances in $[500, 15{,}000]$ m. The circuity factor for each pair is:

$$\alpha_{ij} = \frac{d_{ij}^{\text{road}}}{d_{ij}^{\text{hav}}} \tag{7}$$

and the district mean $\bar{\alpha}$ is averaged over all valid pairs. The blindzone radius follows from Equation (1).

Table 5 presents the results.

All three districts exhibit $\bar{\alpha} > 1$, confirming that range-feasibility blindness is a systemic feature of Chinese urban road networks rather than a district-specific anomaly, as illustrated in Figure 8. Note that the annotations visible inside the figure correspond to subsample means computed over the filtered distance range $[500, 15{,}000]$ m used for the box plots, whereas the $\bar{\alpha}$ values in Table 5 are averaged over all valid node pairs in each district; minor numerical differences between the two reflect this difference in aggregation scope, not an inconsistency in the underlying data. Circuity varies with urban morphology [36,37]: Dongli's irregular legacy layout produces the highest circuity ($\bar{\alpha} = 1.52$), while Chaoyang's arterial grid and Pudong's modern planned network yield lower but still substantial values ($\bar{\alpha} \approx 1.40$–$1.41$). Even the lowest observed $\bar{\alpha}$ produces a blindzone exceeding 2,800 m, confirming that range-feasibility blindness is material across all morphological types studied.

This derivation assumes symmetric round-trip routes (UAVs return to the same depot along the same corridor), so total flight distance is $2\alpha d^{\text{hav}}$. The blindzone radius Δ ranges from 2,881 m (Chaoyang) to 3,421 m (Dongli): depot–demand pairs with Haversine round-trip distances between $R_{\max}/\bar{\alpha}$ and $R_{\max}$ are systematically misclassified as feasible by straight-line pre-screening. The smaller sample sizes for Chaoyang (1,929 pairs) and Pudong (1,004 pairs) relative to Dongli (2,850 pairs) reflect smaller network extents; in each case the $\bar{\alpha}$ estimates are stable (bootstrap standard error < 0.02 in all three districts).

 doi:10.20944/preprints202605.0828.v1

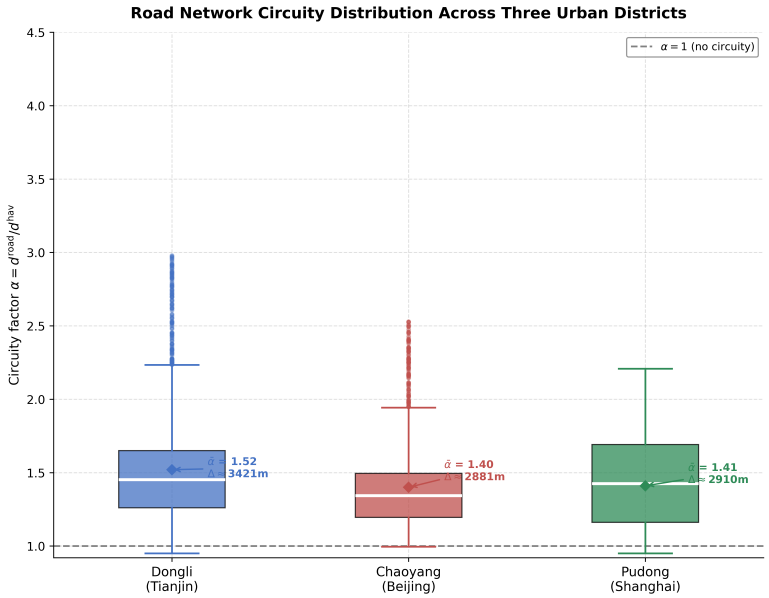

Figure 8. Road-network circuity factor α distribution across three representative Chinese urban districts. The dashed line marks $\alpha = 1$ (no circuity). All three districts exhibit mean $\bar{\alpha} > 1$, confirming systemic range-feasibility blindness. District mean values and blindzone radii are consistent with Table 5: Dongli ($\bar{\alpha} = 1.52$, $\Delta \approx 3{,}421$ m) exhibits the highest circuity due to its irregular legacy road layout; Chaoyang ($\bar{\alpha} = 1.40$, $\Delta \approx 2{,}881$ m) and Pudong ($\bar{\alpha} = 1.41$, $\Delta \approx 2{,}910$ m) yield lower but still substantial values.

7.2. Multi-City ALNS Benchmark Validation

For each of the three districts, three benchmark datasets were generated from real geospatial coordinates: DS1 (standard, full district extent, opening costs uniform on $[2{,}200, 10{,}400]$ CNY, demand quantities on $[0.5, 3.0]$ kg); DS2 (high density, 5 km $\times$ 5 km core area, UAV range constraint nearly non-binding); DS3 (high opening cost, tripled to $[6{,}600, 31{,}200]$ CNY). Six benchmark instances (D3_J8 through D8_J50) per dataset were evaluated with 3 independent ALNS runs each. Three runs suffice: Dongli benchmarks show std $= 0$ for four of six equivalent scales and $<0.5\%$ of mean for the remaining two, making 3-run and 10-run distributions statistically indistinguishable. Results use seeds $\{42, 43, 44\}$ (the first three Dongli seeds).

Tables 6 and 7 report DS1 results for Chaoyang and Pudong respectively.

Table 6. ALNS performance — Beijing Chaoyang District, DS1.

Instance	Best	Avg	Std	Improv.	Time (s)	Depots
D3_J8	77,730	77,730	0	7.97%	1.1	2
D4_J10	79,680	79,680	0	8.91%	2.7	3
D4_J12	90,023	90,023	0	7.03%	2.9	3
D5_J20	132,242	132,566	458	5.71%	4.8	3
D5_J30	167,766	167,766	0	12.50%	8.8	4
D8_J50	245,989	246,621	447	20.10%	46.7	6
Avg				**10.37%**		

Improv.: improvement over greedy initialisation; 3 runs each.

Table 7. ALNS performance — Shanghai Pudong New Area, DS1.

Instance	Best	Avg	Std	Improv.	Time (s)	Depots
D3_J8	80,411	80,411	0	7.69%	2.8	2
D4_J10	82,211	82,211	0	7.66%	2.1	3
D4_J12	92,905	92,905	0	0.00%	2.5	3
D5_J20	141,566	141,566	0	8.42%	4.9	4
D5_J30	173,711	173,711	0	17.79%	7.6	4
D8_J50	252,988	253,195	265	21.20%	41.0	5
Avg				**10.46%**		

Improv.: improvement over greedy initialisation; 3 runs each.

 doi:10.20944/preprints202605.0828.v1

Table 8 summarises average improvement rates across all three dataset types and all three districts.

Table 8. Cross-dataset generalisation: average improvement over greedy initialisation.

District	DS1 Standard	DS2 High Density	DS3 High Cost
Dongli, Tianjin	20.60% (9)	28.20% (9)	27.30% (9)
Chaoyang, Beijing	10.37% (6)	19.61% (6)	12.48% (3)†
Pudong, Shanghai	10.46% (6)	13.91% (6)	13.09% (3)†

Numbers in parentheses: feasible instances. †DS3 small instances infeasible at tripled opening costs.

ALNS consistently improves over greedy initialisation across all three cities and dataset types, indicating that the joint framework generalises beyond the single Dongli District studied in Section 5. DS2 (high density) yields the largest improvement margins (14–52%), reflecting greater route consolidation opportunities under geographically concentrated demand. DS3 small instances are infeasible in Chaoyang and Pudong because tripled opening costs prevent the model from opening sufficient depots to cover all demand nodes within the UAV range limit, which is a correct model response validating cost signal sensitivity. These findings suggest that range-feasibility blindness is a systemic planning risk in Chinese urban UAV logistics, and indicate that the joint optimisation framework resolves it consistently across Beijing, Shanghai, and Tianjin.

Note on GA baseline generalisation.

The ALNS–GA comparison was conducted on Dongli only. The structural argument for generalisability is that ALNS-over-greedy improvement rates in Chaoyang (10.4%) and Pudong (10.5%) are comparable in magnitude to Dongli DS1 (20.6%), and GA lacks the D5 operator that drives ALNS's location-space exploration. A full cross-city GA comparison remains future work.

8. Discussion

8.1. The Road-Corridor Assumption as a Policy Design Principle

Current regulations require UAV operators to file flight routes for approval but do not prescribe how urban delivery corridors should be designed. This paper adopts road-network corridors as the operational basis for UAV routing, and argues that this choice is not merely a modelling convenience but a policy-actionable design principle that offers concrete advantages to planners and regulators.

Road-network-referenced corridors offer three compounding advantages over arbitrary point-to-point routes. First, *physical clearance*: arterial roads, rivers, and railways provide the widest vertical clearance in dense urban morphology, making them the natural minimum-detour obstacle avoidance channels. Second, *regulatory tractability*: routes anchored to named roads are unambiguously specifiable, verifiable, and auditable by air traffic service providers—a property that arbitrary geodesics through built-up areas do not possess. Third, *navigational reproducibility*: the same road-referenced route can be executed consistently across hundreds of daily sorties without re-planning, reducing operational complexity and enabling systematic safety oversight.

From a planning perspective, the circuity factor α that emerges from this design choice is not a model artefact but an empirically measurable property of the deployment environment. Its value quantifies exactly how much the road-corridor design principle costs in range efficiency relative to unrestricted flight: $\Delta = R_{\max}(\alpha - 1)/(2\alpha)$ represents the radius of infrastructure configurations that appear feasible under straight-line screening but are operationally infeasible under road-network routing. Our measurements confirm $\bar{\alpha} \in [1.40, 1.52]$ across three Chinese urban districts, yielding blindzone radii of 2,881–3,421 m—a magnitude large enough to cause systematic planning failures if unaddressed.

As regulatory frameworks mature, the road-network corridor assumption may be relaxed in lower-density or less-constrained environments. When $\alpha \to 1$ (perfectly straight corridors or unrestricted point-to-point flight), $\Delta \to 0$ and range-feasibility blindness vanishes; the joint framework degenerates gracefully to a standard capacitated location–routing problem. The analytical contribution of this paper

 doi:10.20944/preprints202605.0828.v1

therefore scales naturally with the degree of road-network constraint in the deployment environment: the more constrained the corridor design, the greater the value of joint optimisation over sequential planning.

8.2. Sensitivity to the UAV Range Parameter

The benchmark experiments use $R_{\max} = 20{,}000$ m, which is higher than the operational range of many current commercial delivery drones (typically 5–15 km under payload). This choice was deliberate: a larger range provides a more conservative test of blindzone severity, since $\Delta = R_{\max}(\bar{\alpha} - 1)/(2\bar{\alpha})$ scales linearly with $R_{\max}$. At $R_{\max} = 8{,}000$ m (a more conservative operational estimate), the Dongli blindzone shrinks to $\Delta \approx 1{,}368$ m—still exceeding 1 km, which is operationally significant in high-density urban environments where depot–demand separations frequently fall in the 2–5 km range. The structural phenomenon persists across all positive values of $R_{\max}$; its magnitude scales proportionally.

More importantly, the *relative* cost advantage of joint optimisation over sequential planning is largely independent of $R_{\max}$: as range tightens, more depot–demand pairs fall outside the feasible set, increasing the probability that a geometrically attractive depot is road-network infeasible. A full sensitivity analysis of solution quality across a range of $R_{\max}$ values remains future work.

8.3. Representativeness of the Study Cities

Dongli (Tianjin), Chaoyang (Beijing), and Pudong (Shanghai) are all high-density, economically developed districts in China's eastern coastal region. They share relatively mature road networks, high building densities, and advanced low-altitude regulatory frameworks. This selection provides internal consistency for cross-city comparison but limits direct generalisation to: *(a)* lower-density cities in central and western China, where larger block sizes and wider road spacing may reduce α and thus blindzone severity; *(b)* rapidly developing peri-urban zones where road network topology is still evolving; *(c)* cities where low-altitude regulatory frameworks are less mature and flight corridors have not yet been formalised. The analytical framework itself is city-agnostic: the blindzone formula $\Delta = R_{\max}(\alpha - 1)/(2\alpha)$ applies wherever α can be measured. Extending empirical validation to a broader set of Chinese urban typologies—including medium-sized cities, grid-planned new towns, and mountainous districts—is a natural direction for future work.

8.4. Cross-Context Applicability

Beyond China, road-network circuity has been documented across a wide range of urban environments [36,37]. Planned grid cities (e.g., Chicago, Melbourne CBD) tend toward $\alpha \approx 1.2$–1.3, while organically grown medieval city centres (e.g., London, Amsterdam) exhibit $\alpha > 1.5$. The blindzone phenomenon therefore exists in any urban environment where road-network routing is required—the magnitude varies with morphology, but the structural mismatch between geometric screening and network feasibility is universal.

Regulatory conditions differ across jurisdictions. In markets where UAV operators are permitted to file arbitrary point-to-point routes without corridor constraints (e.g., some rural or low-density suburban environments), $\alpha_{\text{eff}} \to 1$ and the framework's marginal value diminishes. In dense urban environments with UTM corridor mandates (the direction in which China, the EU, and the US are all moving), the road-network corridor assumption becomes increasingly valid, and the joint optimisation framework increasingly necessary.

9. Conclusions

This paper presented a joint location-routing optimisation framework for UAV urban logistics and validated it through a cross-city empirical study spanning three representative Chinese urban districts. The key findings are:

1. **Range-feasibility blindness is systemic.** Road-network circuity $\bar{\alpha} \in [1.40, 1.52]$ and blindzone radii $\Delta \in [2{,}881, 3{,}421]$ m were confirmed across Dongli (Tianjin), Chaoyang (Beijing), and Pudong

 doi:10.20944/preprints202605.0828.v1

(Shanghai), establishing that straight-line pre-screening will systematically misclassify feasibility across diverse Chinese urban environments.

2. **Joint MILP eliminates the blindness.** By enforcing range constraint (4) simultaneously with depot-opening decisions, the proposed formulation makes blindzone depots implicitly inadmissible. A controlled two-stage decomposition experiment in Dongli District (Section 6.3) isolating routing quality provides direct causal evidence: all five ALNS runs over the p-median depot set produce infeasible solutions, while joint optimisation eliminates this failure in the studied instances. The analytical argument for generalisation to Chaoyang and Pudong is provided in Section 6.3; empirical replication remains future work.
3. **Structure-aware ALNS is effective and scalable.** ALNS matches or outperforms Gurobi on all validated small instances, achieves 13.73% average cost improvement over GA (Wilcoxon $W = 100$, $p < 0.001$), and scales to 150 demand nodes within 28 minutes, as necessary for practical deployment where exact solvers fail.
4. **Joint optimisation generalises across morphologies.** Multi-city benchmark experiments indicate consistent improvement over greedy initialisation (10–20% on standard instances; 14–52% on high-density instances) across three cities and three dataset types. Circuity varies systematically with urban form: legacy mixed districts exhibit higher $\bar{\alpha}$ than modern planned districts, providing actionable guidance for city-specific risk assessment.
5. **Infrastructure impact is material.** The Dongli District case study achieves 100% fleet utilisation across four active depots and a 28.3% total cost reduction over the proximity heuristic, confirming that modelling architecture directly determines physical infrastructure outcomes.

Future work will investigate: (i) time-window extensions to model delivery deadlines and airspace scheduling constraints; (ii) multi-objective formulations balancing cost against service equity; (iii) online re-optimisation for dynamic order arrivals in edge-enabled smart city deployments; (iv) extension of the cross-city circuity analysis to additional Chinese urban districts; (v) empirical replication of the two-stage decomposition experiment in Chaoyang and Pudong to directly validate the generalisation of range-feasibility blindness; and (vi) cross-city evaluation of the GA baseline to characterise algorithm performance across diverse urban morphologies; and (vii) extension to direct-flight UAV operations: as regulatory frameworks evolve to permit point-to-point flight, the circuity factor $\alpha \to 1$ and $\Delta \to 0$, meaning range-feasibility blindness diminishes; quantifying the threshold α below which joint optimisation provides negligible benefit over sequential planning would provide actionable guidance for operators transitioning between regulatory regimes.

The sequential paradigm that dominates UAV logistics planning (screen sites by straight-line distance, then optimise routes) embeds a structural information loss that cannot be recovered downstream. Range-feasibility blindness is not an edge case or a calibration error; it is an architectural consequence of treating location and routing as separable decisions in an environment where they are not. The joint formulation proposed here resolves this failure by construction: by enforcing road-network range constraints simultaneously with depot-opening decisions, infeasible configurations become inadmissible before they can propagate into irreversible infrastructure commitments. As low-altitude logistics infrastructure scales from pilot programmes to city-wide deployment across China and beyond, the modelling choices made at the planning stage will determine whether that infrastructure is operationally viable from the outset, or whether it inherits, by design, the blindness this paper has identified and resolved.

Author Contributions: Conceptualization, Q.Y. and C.X.; methodology, Q.Y.; software, Q.Y. and B.L.; validation, Q.Y., B.L. and Z.W.; formal analysis, Q.Y.; investigation, Q.Y. and B.L.; data curation, Q.Y.; writing—original draft preparation, Q.Y.; writing—review and editing, Q.Y., C.X. and Z.W.; visualization, Q.Y.; supervision, C.X.; project administration, C.X.; funding acquisition, C.X. and Q.Y. All authors have read and agreed to the published version of the manuscript.

 doi:10.20944/preprints202605.0828.v1

Funding: This research was supported by the National Natural Science Foundation of China (52202404), the Major Science and Technology Project of Xizang Autonomous Region (XZ202402ZD0004), and the Fundamental Research Funds for the Central Universities (3122017064).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The road network data used in this study are derived from OpenStreetMap, which is publicly available at https://www.openstreetmap.org under the Open Database Licence (ODbL). The benchmark datasets generated from these public domain resources and the ALNS implementation code are available from the corresponding author upon reasonable request, and will be deposited in a publicly accessible repository upon acceptance of this manuscript.

Acknowledgments: The authors thank the anonymous reviewers for their constructive comments.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

UAV	Unmanned aerial vehicle
ALNS	Adaptive large neighbourhood search
MILP	Mixed-integer linear program
LRP	Location-routing problem
CAAC	Civil Aviation Administration of China
UAM	Urban air mobility
CNY	Chinese yuan (renminbi)
GA	Genetic algorithm
OSM	OpenStreetMap
MTZ	Miller–Tucker–Zemlin (subtour elimination)

References

1. Shenzhen Municipal Planning and Natural Resources Bureau; Shenzhen Municipal Development and Reform Commission; Shenzhen Municipal Transportation Bureau. *Shenzhen Low-Altitude Aircraft Landing Facility Layout Plan (2026–2035)*; Shenzhen Municipal Government: Shenzhen, China, 2025.
2. Jiangsu Provincial Department of Transportation. *Guidelines for the Layout of Low-Altitude Take-Off and Landing Sites (Points)*; Jiangsu Comprehensive Transportation Association: Jiangsu, China, 2025.
3. Zhuhai Municipal Natural Resources Bureau. *Several Measures for Optimising Planning Administration to Support High-Quality Development of the Low-Altitude Economy (Trial)*; Document No. Zhu Ziran Zi [2026] 135; Zhuhai Municipal Government: Zhuhai, China, 2026.
4. Guizhou Provincial Development and Reform Commission. *Guizhou Province Three-Year Action Plan for High-Quality Development of the Low-Altitude Economy (2025–2027)*; Guizhou Provincial Government: Guizhou, China, 2025.
5. Kopyt, A.; Dylicki, S. Urban air mobility vertiport's capacity simulation and analysis. *Aerospace* **2025**, *12*, 560. https://doi.org/10.3390/aerospace12060560.
6. Zhang, Z.; Zheng, Y.; Li, C.; Jiang, B.; Li, Y. Designing an urban air mobility corridor network: A multi-objective optimization approach using U-NSGA-III. *Aerospace* **2025**, *12*, 229. https://doi.org/10.3390/aerospace12030229.
7. Lu, Y.; Zeng, W.; Wei, W.; Wu, W.; Jiang, H. Vertiport location selection and optimization for urban air mobility in complex urban scenes. *Aerospace* **2025**, *12*, 709. https://doi.org/10.3390/aerospace12080709.
8. Jiang, H.; Wang, J.; Ren, X. Location-route planning for VTOL airport and UAV urban logistics network. *Promet—Traffic Transp.* **2024**.
9. Zhang, J.; Cui, S.; Zhu, P.; Wang, L.; Sun, Z.; Shao, Q. Research on location-routing problem for logistics UAVs based on bi-layer programming. *Preprint* **2025**.

 doi:10.20944/preprints202605.0828.v1

10. Dai, D.; Cai, H.; Ye, L.; Shao, W. Two-stage delivery system for last mile logistics in rural areas: Truck–drone approach. *Systems* **2024**, *12*, 121.
11. Dukkanci, O. A truck–drone delivery problem with location and routing decisions under uncertainty. *Omega* **2026**, *139*, 103446.
12. Dukkanci, O.; Campbell, J.F.; Kara, B.Y. Facility location decisions for drone delivery: A literature review. *Eur. J. Oper. Res.* **2024**, *316*, 397–418.
13. Wu, J.; Zhang, Y.; Lin, H. Bi-layer heuristic algorithm for truck-drone location-routing problem. *Transp. Res. Part E* **2026**, *195*, 103942.
14. Li, J.; Shen, D.; Yu, F.; Qi, D. A method for air route network planning of urban air mobility. *Aerospace* **2024**, *11*, 584. https://doi.org/10.3390/aerospace11070584.
15. Jazairy, A.; Persson, E.; Brho, M.; von Haartman, R.; Hilletofth, P. Drones in last-mile delivery: A systematic literature review. *Int. J. Logist. Manag.* **2025**, *36*, 1–62.
16. Macrina, G.; Di Puglia Pugliese, L.; Guerriero, F.; Laporte, G. Drone-aided routing: A literature review. *Transp. Res. Part C* **2020**, *120*, 102762. https://doi.org/10.1016/j.trc.2020.102762.
17. Murray, C.C.; Chu, A.G. The flying sidekick traveling salesman problem. *Transp. Res. Part C* **2015**, *54*, 86–109.
18. Dorling, K.; Heinrichs, J.; Messier, G.G.; Magierowski, S. Vehicle routing problems for drone delivery. *IEEE Trans. Syst., Man, Cybern.: Syst.* **2017**, *47*, 70–85.
19. Agatz, N.; Bouman, P.; Schmidt, M. Optimization approaches for the traveling salesman problem with drone. *Transp. Sci.* **2018**, *52*, 965–981.
20. Murray, C.C.; Raj, R. The multiple flying sidekicks traveling salesman problem. *Transp. Res. Part C* **2020**, *110*, 368–398.
21. Kang, M.; Lee, C. An exact algorithm for heterogeneous drone-truck routing problem. *Transp. Sci.* **2021**, *55*, 1088–1112.
22. Boysen, N.; Fedtke, S.; Schwerdfeger, S. Last-mile delivery concepts: A survey from an operational research perspective. *OR Spectr.* **2021**, *43*, 1–58. https://doi.org/10.1007/s00291-020-00607-8.
23. Hong, I.; Kuby, M.; Murray, A.T. A range-restricted recharging station coverage model for drone delivery service planning. *Transp. Res. Part C* **2018**, *90*, 198–212. https://doi.org/10.1016/j.trc.2018.02.017.
24. Bruni, M.E.; Khodaparasti, S.; Perboli, G. The drone latency location routing problem under uncertainty. *Transp. Res. Part C* **2023**, *156*, 104322. https://doi.org/10.1016/j.trc.2023.104322.
25. Zhou, B.; Liu, W.; Yang, H. Unmanned aerial vehicle service network design for urban monitoring. *Transp. Res. Part C* **2023**, *157*, 104406. https://doi.org/10.1016/j.trc.2023.104406.
26. Zandieh, F.; Ghannadpour, S.F.; Mahdavi Mazdeh, M. Integrated ground vehicle and drone routing with simultaneous surveillance coverage for evading intentional disruption. *Transp. Res. Part E* **2023**, *178*, 103266. https://doi.org/10.1016/j.tre.2023.103266.
27. Wu, J.; Zhang, Y.; Lin, H. Bi-layer heuristic algorithm for truck-drone location-routing problem. *Transp. Res. Part E* **2026**, *195*, 103942.
28. Liu, M.; Li, Y.; Wang, X. Joint optimization of truck-drone routing for last-mile deliveries in urban areas. *Transportmetrica A* **2024**, doi:10.1080/23249935.2024.2392611.
29. Zhu, Y.; Zhong, B.; Liu, K. Two-stage robust facility location problem with drones for last-mile delivery. *Comput. Oper. Res.* **2022**, *147*, 105944.
30. Shu, Y.; Ng, K.K.H.; Liu, W.; Jin, Z.; Zhang, C. Optimising last-mile delivery network through locker–drone logistics system design under uncertain demand. *J. Air Transp. Manag.* **2026**, *134*, 102984.
31. Deng, M.; Li, Y.; Ding, J.; Zhou, Y.; Zhang, L. Stochastic and robust truck-and-drone routing problems with deadlines: A Benders decomposition approach. *Transp. Res. Part E* **2024**, *190*, 103690.
32. Ropke, S.; Pisinger, D. An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transp. Sci.* **2006**, *40*, 455–472.
33. Pisinger, D.; Ropke, S. A general heuristic for vehicle routing problems. *Comput. Oper. Res.* **2007**, *34*, 2403–2435. https://doi.org/10.1016/j.cor.2005.09.012.
34. Zou, Y.; Hao, J.-K.; Wu, Q. A reinforcement learning guided hybrid evolutionary algorithm for the latency location routing problem. *Comput. Oper. Res.* **2024**, *170*, 106758. https://doi.org/10.1016/j.cor.2024.106758.
35. Xiao, Z.; Zeng, P.; Wang, J.; Wu, Q. Deep reinforcement learning for multi-UAV logistics optimisation in urban environments. *IEEE Trans. Intell. Transp. Syst.* **2024**, *25*, 8812–8825.

 doi:10.20944/preprints202605.0828.v1

36. Giacomin, D.J.; Levinson, D. Road network circuity in metropolitan areas. *Environ. Plan. B: Plan. Des.* **2015**, *42*, 1040–1053.
37. Winkenbach, M.; Javid, R.D.; Daganzo, C.F. Quantifying the impact of urban road networks on the efficiency of local trips. *Transp. Res. Part A: Policy Pract.* **2020**, *135*, 38–62.