Journal of Mathematical Modeling
Vol. -, No. -, 2026, pp. -. Research Article

JMM

A robust bi-level location-routing optimization problem for motorcycle-UAV cooperative agricultural monitoring under battery capacity uncertainty

Azita Taheri[†], Rooholah Abedian[†*], Mahdi Hamid[‡], Amin Ghodousian[†]

[†] *School of Engineering Science, College of Engineering, University of Tehran, Iran*
[‡] *Department of Industrial Engineering and Management Systems, Amirkabir University of Technology, Tehran, Iran*

Email(s): rabedian@ut.ac.ir

Abstract. This paper develops a robust bi-level optimization model for integrated depot location and routing in UAV-based agricultural surveillance networks under battery-capacity uncertainty. The upper-level problem determines depot locations and assigns gardens to selected depots while anticipating operational decisions. The lower-level problem optimizes coordinated missions of quad-plane UAVs and ground-support motorcycles. The model explicitly considers routing, service scheduling, charging-station synchronization, and battery inventory management. Battery-capacity uncertainty is incorporated using the Bertsimas–Sim budgeted uncertainty framework, which leads to a robust reformulation of the operational problem. To address the computational complexity of real-scale instances, a hybrid solution approach combining the bi-level optimization framework with K-means clustering is proposed. The model is evaluated on a real-world pistachio surveillance network in Kerman Province, Iran. Computational results and managerial insights demonstrate the applicability and effectiveness of the proposed approach.

Keywords: Bi-level optimization model, UAV routing planning, depot location problem, agricultural surveillance modeling, routing scheduling
AMS Subject Classification 2010: 90B06, 90B25, 90B50

1 Introduction

Global agricultural production is increasingly shaped by climate variability, population growth, and tightening constraints on land, water, and labor. The agricultural sector contributes approximately \$3.4 trillion

*Corresponding author
Received: 08 April 2026/ Revised: 25 August 2026/ Accepted: 26 August 2026
DOI: 10.22124/jmm.2026.33228.3027

annually to global GDP, and specialty crops—including pistachios, almonds, olives, and tree nuts—account for approximately \$340 billion of this value [15, 23]. Because these crops are highly valued and sensitive to localized stress, the timely identification of irrigation problems, pest infestations, disease outbreaks, and theft incidents has a direct impact on yield and profitability. In many production regions, conventional ground-based scouting covers only 3–5% of the cultivated area per inspection cycle, which contributes to delayed detection and 12–18% avoidable losses [26, 34]. Theft further compounds these losses: specialty crops may face 8–15% annual theft rates in geographically dispersed regions, and pistachios are particularly vulnerable during pre-harvest periods when market prices are high. These combined agronomic and security pressures have motivated the adoption of unmanned aerial vehicle (UAV) technologies for comprehensive surveillance.

UAV-based monitoring has expanded rapidly in recent years. Adoption among large-scale specialty crop farms increased from roughly 5% in 2010 to about 35% by 2020, driven by lower sensor costs and improved platform performance [5, 50]. UAVs equipped with multispectral sensors can achieve near-complete coverage in short time windows and detect plant stress 5–10 days before visible symptoms appear [23, 34]. This early-warning capability supports targeted interventions, such as localized irrigation adjustment or site-specific pest treatment, which can reduce chemical input by 20–30% and water use by 15–25% while maintaining yields [26, 45]. These benefits position UAV surveillance as an enabling tool for sustainable intensification consistent with the Sustainable Development Goals of the UN (SDG 2: Zero Hunger, SDG 12: Responsible Consumption) [15]. Furthermore, such technological interventions are essential for developing integrated strategic frameworks that ensure sustainable climate-smart crop planning and resilient agri-food supply chain management [20].

Beyond agronomic monitoring, UAV surveillance can also address agricultural security challenges that are especially acute for high-value crops with concentrated harvest windows. Iran's pistachio industry—producing about 40% of the global supply and generating roughly \$1.2 billion in annual export value—experiences 10–18% pre-harvest theft rates in Kerman Province's dispersed orchards [14]. Traditional security measures (e.g., permanent guarding and perimeter fencing) are often economically impractical for orchards spanning 5–50 hectares and provide limited flexibility for seasonal protection over a 2–3 month window. In contrast, UAV-based surveillance enables frequent monitoring at relatively low marginal cost; thermal imaging night operations can support detection of unauthorized presence, randomized patrol patterns reduce predictability, and real-time video transmission can facilitate rapid response. As a result, a single surveillance system can simultaneously support agronomic decision-making and theft deterrence, improving the return on investment and making cooperative ownership models more realistic.

Despite these advantages, deploying UAV surveillance at a regional scale remains challenging. First, endurance constraints limit operational reach; battery-powered platforms typically achieve 50–80 km for rotary-wing and 120–150 km for fixed-wing UAVs under nominal conditions [1, 13], which is often insufficient when farms are far from depots. Second, the placement of depots and supporting infrastructure is a strategic decision shaped by multiple criteria; specialty crop production is spatially dispersed (e.g., pistachio orchards in Kerman Province span about 15,000 km^2), and the resulting service-radius constraints lead naturally to multi-depot facility location problems [31, 38]. Third, uncertainty in drone battery capacity is unavoidable: battery capacity can degrade by 15–30% under temperature extremes (from $-10°$C to $+45°$C) common in semi-arid regions, and deterministic planning may therefore produce routes that are infeasible in practice [3, 44]. At the same time, security-focused operations require frequent and difficult-to-predict patrols, which can conflict with purely efficiency-driven schedules.

These constraints point to the need for integrated optimization models that jointly consider strategic infrastructure decisions (e.g., depot selection) and operational planning (routing, scheduling, and energy management) under uncertainty. Much of the existing literature still treats these decisions sequentially: depots are selected using distance-based measures or multi-criteria ranking, and routing is optimized afterward assuming fixed infrastructure [31, 38]. Such decoupling can be costly because depot choices directly shape the feasible routing space, and routing outcomes, in turn, determine the true value of candidate depots; integrated bi-level models capture this leader-follower structure explicitly and can reduce system costs by 15–25% compared with sequential approaches [10, 40]. In addition, UAV routing studies often rely on deterministic battery models despite documented variability [4, 44], and many formulations assume homogeneous fleets, overlooking the operational role of ground vehicles that can enable extended-range missions and energy logistics [28, 35].

Sustainability imperatives, encompassing both environmental and economic dimensions, further underscore the necessity for integrated and robust planning. Environmentally, agricultural aviation contributes an estimated 2–3% of sector-wide greenhouse gas emissions [45]. Optimal depot placement minimizes unproductive transit, while efficient routing configurations directly reduce operational energy consumption. Furthermore, incorporating robustness mitigates the risk of mid-mission failures that would otherwise necessitate resource-intensive emergency redeployments. From an economic perspective, initial capital expenditures for depot establishment and regulatory compliance typically range from \$50,000 to \$200,000 per facility, contingent upon site-specific infrastructure conditions [38]. Consequently, suboptimal strategic planning can result in misallocated investments, imposing prohibitive financial burdens on agricultural cooperatives and highlighting the critical need for mathematically rigorous infrastructure decisions.

The present study addresses these gaps through a robust bi-level optimization framework for depot location and routing in agricultural UAV surveillance networks with heterogeneous vehicle coordination under drone battery-capacity uncertainty. We study pistachio surveillance in Iran's Kerman Province (45 gardens, 10 candidate depot sites, and 8 charging stations), where electric quad-plane UAVs and ground-support motorcycles must operate as an integrated system. In this setting, motorcycles are not merely auxiliary vehicles; they conduct surveillance and deliver batteries from selected depots to charging stations so that UAVs can perform multi-battery tours. The model integrates a Stackelberg leader-follower structure where the strategic planners select depots and assignments while anticipating operational responses; heterogeneous fleet coordination with temporal synchronization for battery supply operations; Bertsimas–Sim budgeted uncertainty to robustify drone battery-capacity feasibility; multi-criteria depot evaluation via the Best–Worst Method (BWM) and TOPSIS; and dual-purpose surveillance requirements supporting both agronomic monitoring and security enforcement.

Figure 1 sketches the type of coordinated operations studied in this paper. Customer nodes represent gardens requiring surveillance visits; the warehouse icon denotes a depot, and the pump icons denote charging stations where mid-mission battery swaps can occur. The colored paths illustrate representative tours for multiple vehicles, and the battery indicators depict how remaining energy evolves along a route. In our application, quad-plane UAVs provide aerial surveillance, while motorcycles play a dual role: they perform ground-level surveillance of assigned gardens, and, crucially, they transport fully charged batteries from depots to charging stations. This battery supply task is operationally central because it allows UAVs to complete extended missions by swapping batteries at stations rather than returning to a depot after each sortie. The resulting system couples strategic and operational decisions; depot locations affect reachability and the feasibility of long tours, while the routing and battery-delivery workload determine which depot configurations are economically attractive and robust.

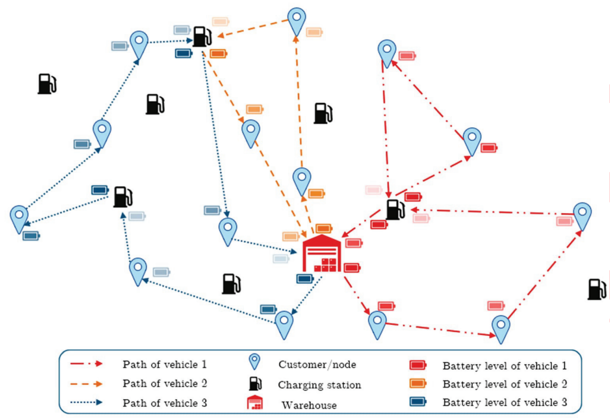

Figure 1: Illustrative multi-vehicle surveillance routes with intermediate charging stations and vehicle battery-state annotations.

These research questions are addressed through the proposed methodology:

- How can strategic depot-location decisions and operational UAV–motorcycle routing decisions be integrated in a robust bi-level framework for energy-constrained agricultural surveillance networks?
- How does UAV battery-capacity uncertainty affect routing feasibility, operational cost, battery-swap requirements, and vehicle utilization under different robustness budgets?
- What computational and operational insights can be obtained from comparing deterministic and robust solutions, as well as centralized single-depot and decentralized two-depot configurations, in the Kerman Province pistachio surveillance case?

These questions are addressed through model development, scalability analysis, robustness sensitivity experiments, and case-specific benchmark comparisons. Therefore, the computational results are interpreted as evidence for the studied agricultural surveillance network rather than as universal prescriptions for all UAV logistics systems.

The remainder of the paper is organized as follows. Section 2 reviews the related literature and highlights research gaps. Section 3 presents the problem description and the mathematical model. Section 4 describes the proposed solution approach. Section 5 reports the computational experiments and discusses the obtained results. Finally, Section 6 concludes the paper and outlines directions for future research.

2 Literature review

This section reviews the literature most relevant to the proposed agricultural UAV surveillance framework. We first synthesize prior research in five aligned domains: (i) UAV-based agricultural surveillance and remote sensing, (ii) endurance- and battery-constrained UAV routing with heterogeneous coordination, (iii) hub location and multi-criteria facility selection, (iv) bi-level optimization and Stackelberg decision-making in logistics, and (v) robust optimization under drone battery-capacity uncertainty. Then, the gap analysis, positioning, and the specific contributions of this study are discussed.

2.1 UAV Applications in precision agriculture and agricultural surveillance

The deployment of unmanned aerial vehicles (UAVs) for agricultural monitoring has evolved from experimental demonstrations to operational systems supporting commercial crop management decisions. Global agricultural production increasingly requires advanced surveillance to address crop theft, pest management, and resource optimization in specialty crop cultivation. Tsouros et al. [45] provide a comprehensive review of UAV-based applications for precision agriculture, documenting adoption rates increasing from $< 5\%$ of large-scale farms in 2010 to $> 35\%$ by 2020, driven by declining sensor costs and improved flight endurance. In high-value crop regions such as Iran's pistachio-producing areas, UAV surveillance directly addresses security concerns: crop theft during harvest represents 10–18% of production losses, constituting billions in annual economic damage to agricultural cooperatives.

For specialty crops requiring intensive surveillance—including pistachios, almonds, and olives–UAV-based monitoring reduces scouting time by 70–80% compared to traditional ground-based inspection [26]. Additionally, UAV-based methods improve pest detection sensitivity by 25–35%. Zhang et al. [50] demonstrate that multispectral imaging from fixed-wing UAVs achieves 85–92% accuracy for early detection of irrigation stress in tree crops, outperforming satellite-based indices (65–75% accuracy) due to superior spatial resolution. Sankaran et al. [34] review advanced techniques for plant disease detection, finding that thermal and hyperspectral imaging from UAV platforms enables pre-symptomatic pathogen identification 5–10 days earlier than visual inspection, which is critical for high-value crops where delayed intervention costs exceed $1,000 per hectare.

Recent advances in multirotor fixed-wing hybrid platforms—commonly termed quad-planes or electric vertical-takeoff-and-landing (eVTOL) aircraft [41]–represent an emerging technology category combining accessibility advantages of rotorcraft with endurance benefits of fixed-wing aerodynamics [51]. These systems employ multiple battery packs optimized for distinct flight phases: dedicated VTOL batteries support high-power, short-duration vertical operations (3 kW over 8 minutes), while separate cruise batteries power efficient forward flight via fixed-wing lift (0.3 kW over 50+ minutes) [6]. Recent designs and implementations document operational feasibility and energy-consumption characteristics for asset monitoring and agricultural surveillance applications [41]. For agricultural surveillance in specialty crop regions, quad-plane platforms offer several distinct operational advantages: (i) elimination of runway infrastructure requirements, enabling direct landing in orchards or field clearings without site preparation [19, 48]; (ii) dual imaging capability combining thermal cameras (for nighttime security surveillance and unauthorized presence detection) with multispectral sensors (for daytime agronomic monitoring, including irrigation stress, pest identification, and disease detection) [12]; (iii) modular heterogeneous battery architecture enabling mid-mission energy replenishment through ground-vehicle battery delivery, extending effective operational radius from single-battery 50 km to multi-battery 200-–300 km; and (iv) significantly quieter operation than combustion-engine UAVs [16], reducing disturbance during sensitive harvest periods. The energy-consumption patterns of quad-planes differ fundamentally from those of traditional fixed-wing aircraft because they incur substantial energy penalties during VTOL phases but achieve high cruise efficiency during fixed-wing phases [13].

However, a critical gap exists in the agricultural UAV literature: while agronomic validation and sensor technology are thoroughly documented, the logistics optimization for deploying networked UAV systems across geographically dispersed agricultural regions remains underexplored. Candiago et al. [5] and Matese et al. [27] evaluate vegetation indices and compare UAV platforms with aircraft and satellite technologies, yet assume unlimited flight endurance and ignore routing constraints, battery logistics constraints, and multi-site coordination optimization that restrict real-world deployments.

Recent precision agriculture reviews further confirm the rapid expansion of UAV applications in crop monitoring, multispectral and hyperspectral imaging, AI-supported diagnosis, and field-level decision support. Beyond immediate theft deterrence, UAVs are increasingly recognized as transformative tools for precision agriculture, capable of significantly enhancing crop productivity and resource management through optimized, data-driven flight paths [42]. For example, current studies on AI-powered UAVs in precision agriculture emphasize advances in sensing, data analytics, crop-health monitoring, and automated detection. However, these contributions remain primarily focused on data acquisition and agronomic interpretation. They provide limited guidance on how UAV fleets should be located, routed, synchronized with ground-support vehicles, and protected against battery-capacity uncertainty in large, dispersed agricultural networks.

2.2 UAV routing and coordination

Recent comprehensive reviews of UAV path planning underscore that while traditional models like the Traveling Salesman Problem (TSP) and Vehicle Routing Problem (VRP) form the mathematical foundation of drone logistics, modern delivery networks increasingly require hybrid, coordinated models to address complex operational constraints [42]. Cengiz et al. [8] provide a comprehensive review of drone–truck routing problems, analyzing traditional and next-generation applications, mathematical models, and solution methods. They emphasize that synchronization between heterogeneous vehicles and increasingly complex objective and constraint spaces are major sources of computational difficulty. Importantly, their review also identifies agriculture-related drone–truck routing as an emerging research direction. However, integrated robust depot-location planning with synchronized battery logistics for agricultural surveillance has not been systematically investigated in previous studies. Dorling et al. [13] develop comprehensive energy-consumption models for multi-rotor and fixed-wing UAVs, demonstrating that payload weight and wind speed induce 15–35% variability in actual versus nominal flight time. Their mixed-integer programming (MIP) formulation for drone delivery minimizes completion time subject to battery capacity constraints, yet assumes centralized depot operations and omits multi-depot facility location decisions—a critical simplification for distributed agricultural surveillance networks.

Aggarwal and Kumar [1] review path planning techniques for UAVs, categorizing approaches into graph-based (Dijkstra, A*), sampling-based (RRT, PRM), and optimization-based (MIP, dynamic programming) methods. While their taxonomy comprehensively covers single-vehicle path planning, multi-vehicle coordination and strategic depot placement receive minimal treatment. Thiels et al. [43] address UAV fleet planning for medical sample transportation, introducing time windows and priority constraints, but maintain single-depot assumptions, limiting applicability to geographically dispersed service networks like agricultural surveillance across provincial regions.

A critical methodological gap exists regarding heterogeneous vehicle coordination. Murray and Chu [28] pioneered the "flying sidekick traveling salesman problem" where trucks transport drones to launch points, introducing the concept of ground–UAV coordination. Recent studies have further advanced the literature on coordinated truck–UAV routing by addressing synchronization, scalability, and dynamic operating conditions. Cengiz et al. [7] examined the TSP-D with variable UAV speed and showed that adjusting UAV speed according to payload weight can improve truck–UAV synchronization and reduce total completion time. Yılmaz et al. [47] developed a fitness-distance balance-based evolutionary algorithm for the truck–multi-drone TSP-D and demonstrated its effectiveness on instances with up to 100 delivery points, highlighting the computational complexity of synchronized truck–drone routing. Yuan et al. [49] investigated collaborative truck–UAV delivery under dynamic weather and demand conditions, using probabilistic wind modeling and an adaptive algorithm combining a genetic algorithm, tabu search, and clustering-based service-area adjustment. Similarly, Long et al. [25] proposed a dynamic truck–UAV collaboration strategy for resilient urban emergency response under road-network disruptions and developed an integrated tabu-search-based scheduling algorithm. Together, these studies show that truck–UAV coordination has received increasing attention in last-mile delivery, smart-city logistics, and emergency response contexts. However, they mainly focus on routing and scheduling problems and do not jointly address strategic depot-location decisions, synchronized battery-swap logistics, robust UAV battery-capacity degradation, and agricultural surveillance networks. This gap motivates the integrated robust bi-level location-routing framework developed in the present study. For agricultural surveillance requiring extended-range missions beyond single-battery capability, synchronized battery logistics with ground-support vehicles remains underrepresented in the academic literature despite its practical relevance. However, recent literature has begun to demonstrate the distinct advantages of drone-motorcycle configurations for managing highly constrained logistics, such as optimizing time-sensitive pharmaceutical deliveries to moving customers [22]. Timilsina et al. [44] document battery degradation in electric vehicles, finding 15–30% capacity loss over 200–500 charge cycles and temperature-induced variations of 10–25% between winter and summer operations. Despite these empirical findings motivating robust formulations, UAV routing literature predominantly employs deterministic battery-capacity models.

More recent efforts, including advancements in 2025 and 2026, have pushed the boundaries of truck–drone routing toward richer coordination structures, such as multi-platform routing, interception-based operations, en-route charging, and scalable algorithms under uncertainty. These studies confirm that synchronization, energy management, and computational scalability are now central concerns in combined ground–air systems. For instance, Wang et al. [46] propose a data-driven robust optimization model for hybrid truck–drone delivery logistics to minimize fleet size and turnaround time under battery performance uncertainty. However, the dominant application context of these advancements remains urban last-mile parcel delivery or generic logistics. They generally do not address the unique requirements of a geographically dispersed orchard network, namely motorcycle-supported battery delivery to swap stations, dual-purpose agricultural surveillance, and robust protection against battery-capacity degradation.

2.3 Hub location, facility location, and multi-criteria depot selection

Strategic depot placement is a foundational decision in logistics networks because facility locations directly determine service reachability, achievable routing structures, and ultimately operational feasibility. In practice, depot selection is rarely a function of distance alone; infrastructure readiness (power, facilities, and regulatory compliance), geographic accessibility, and capital/operating costs are equally critical. When multiple, potentially competing criteria must guide infrastructure decisions, multi-criteria decision-making (MCDM) frameworks offer structured approaches to combine quantitative and qualitative information systematically.

The Best–Worst Method (BWM) is a recent MCDM technique that derives criterion weights through pairwise comparisons between best and worst criteria [33]. Compared to fully paired approaches (e.g., Analytic Hierarchy Process), BWM offers computational efficiency and is increasingly applied in logistics contexts, including supplier selection, warehouse siting, and facility planning. However, the method requires careful calibration of expert judgments and consistency checking to ensure reliable weight estimates.

After criterion weights are established, methods such as TOPSIS (Technique for Order Preference by Similarity to Ideal Solution) provide transparent mechanisms for ranking candidate sites. TOPSIS measures distance from ideal and anti-ideal solutions, offering clarity when comparing disparate infrastructure options. Within the specific context of UAV logistics, multi-criteria frameworks have proven highly effective for infrastructure planning; for instance, recent location-routing models have successfully employed the TOPSIS method to evaluate and rank candidate drone launching and charging stations based on diverse sustainability indicators [18]. While such multi-criteria frameworks are valuable for initial candidate screening, they address facility selection in isolation from operational routing. This separation represents a critical shortcoming in many studies because facility location and vehicle routing exhibit strong circular dependencies: optimal depot locations depend on achievable routing patterns, which, in turn, depend on where facilities are positioned.

From an operations research perspective, this interdependence is well-established in the location-routing literature [11, 29]. Recent studies further emphasize the criticality of integrated network design when operations are subject to severe disruptions or capacity uncertainties, demonstrating that robust optimization models are vital for maintaining system resilience [32]. Nagy and Salhi [29] conduct a comprehensive survey of location-routing problems, showing that integrated approaches (simultaneously optimizing facility locations and routes) achieve 10–25% cost reductions versus sequential methods (selecting hubs first, routing second). When applied to energy-constrained UAV networks, this interdependence becomes even more pronounced: the choice of depot locations directly shapes mission feasibility (e.g., whether single-battery trips are viable or battery-swapping logistics are required), which, in turn, drives the economic case for opening particular facilities.

Several studies exemplify the practical application of MCDM to transportation infrastructure. For UAV hub selection in particular, multi-criteria evaluation can incorporate accessibility, security, airspace constraints, and charging infrastructure. However, existing studies typically treat such multi-criteria hub selection and subsequent routing as sequential steps, missing opportunities for integrated optimization that would yield operationally feasible and cost-effective infrastructure portfolios. The challenge of combining MCDM facility selection with operational routing constraints—especially endurance and battery logistics in UAV systems—has received limited attention in the existing literature.

Overall, facility-location and MCDM studies provide useful tools for depot evaluation, but they often remain disconnected from downstream routing feasibility, battery logistics, and operational robustness.

2.4 Bi-level optimization in logistics

Bi-level programming provides a mathematical framework for hierarchical decision-making where an upper-level leader makes decisions anticipating optimal responses from a lower-level follower. Colson et al. [10] survey bi-level optimization theory, categorizing solution approaches into vertex enumeration (suitable for small-scale problems), descent methods (requiring differentiability), and metaheuristics (applicable to large-scale, non-convex instances).

The computational complexity of bi-level linear programming is well-established: the problem is NP-hard in general [10], meaning that finding globally optimal solutions for large instances becomes computationally intractable without specialized approaches. In broader logistics network design, researchers frequently deploy advanced metaheuristic algorithms to overcome structural complexities and exact solver limitations [36]. This complexity motivates the decomposition and enumeration strategies employed in practical applications. While small monolithic instances can be solved exactly, moderate-to-large networks (e.g., 10 candidate depots, 45 service nodes) require spatial decomposition and matheuristic approximation methods to maintain computational tractability.

Sinha et al. [40] review bi-level optimization from classical to evolutionary approaches, documenting applications in network design (leader: infrastructure investment, follower: user route choice), supply chain management (manufacturer-retailer negotiations), and environmental policy (regulator-polluter interactions). Transportation and logistics applications predominantly address passenger travel or freight distribution; agricultural surveillance logistics with endurance-constrained aerial vehicles and heterogeneous fleet coordination remains an underexplored bi-level application domain.

Parvasi et al. [30] develop a bi-level model for school bus routing where the leader selects bus stops and allocates students while the follower optimizes routes given those assignments. Their hybrid metaheuristic combining genetic algorithms (upper level) and variable neighborhood search (lower level) achieves near-optimal solutions (gaps $< 2\%$) for instances with 40 stops and 500 students. While structurally analogous to our agricultural surveillance problem, school bus routing omits endurance-driven feasibility constraints, heterogeneous vehicle coordination, and drone battery-capacity uncertainty—extensions critical for UAV applications.

A key observation is that existing bi-level logistics research overwhelmingly employs deterministic formulations, treating parameters as known constants. For UAV operations where battery performance, weather conditions, and service demands exhibit significant variability, this assumption proves problematic and contributes to mission failures documented in practical deployments. Some very recent studies have begun to address this limitation. For example, Stagg and Peterson [**?**] develop a bi-level route optimization and path planning strategy for UAVs that explicitly allocates route budgets based on uncertain battery capacities. While promising, their approach is tailored to single-vehicle hazard exploration rather than coordinated multi-vehicle agricultural monitoring.

2.5 Robust optimization under battery-capacity uncertainty

Robust optimization frameworks seek solutions that remain feasible and near-optimal across a range of uncertain parameter realizations, avoiding both the brittleness of deterministic models and the computational burden of stochastic programming. Bertsimas and Sim [3] introduce budgeted uncertainty sets for robust linear optimization, proving that the "price of robustness"–increased objective function cost for guaranteed feasibility–grows controllably with the uncertainty budget Γ. This framework maintains a linear programming structure, enabling solutions via standard MIP solvers without the specialized decomposition algorithms required by scenario-based approaches.

Ben-Tal et al. [2] provide a comprehensive treatment of robust optimization theory, establishing that for problems with n uncertain parameters and an uncertainty budget Γ, the probability that more than Γ parameters deviate to their worst-case values decays exponentially with Γ, yielding high-confidence feasibility guarantees. Applications to transportation include bus scheduling under uncertain travel times, airline crew rostering with flight delays [39], and inventory management with demand uncertainty.

However, robust vehicle routing problems–particularly for endurance-constrained UAVs under battery-capacity uncertainty–remain underexplored in the academic literature. Shahnaghi et al. [37] apply Bertsimas–Sim robust optimization to batch processing flow shop scheduling, demonstrating an 18–25% reduction in schedule tardiness variance compared to deterministic optimization at the cost of an 8–12% increased mean completion time. Their selective robustification strategy–protecting only critical constraints while maintaining others deterministically–inspires our approach of robustifying drone battery-capacity feasibility constraints while keeping other operational constraints deterministic.

Recent studies on drone location-routing and robust drone facility planning have made important progress in modeling uncertainty, charging decisions, and battery-related limitations. These contributions show that uncertain travel conditions, energy consumption, demand, or battery availability can significantly affect both facility location and routing decisions. For example, Ghasemi et al. [18] integrated facility location with drone routing by explicitly routing drones to opened charging stations whenever their payload-dependent battery capacity was insufficient to complete a delivery mission. Very recent literature further confirms the critical need for integrating robust charging decisions and infrastructure siting; for example, Ghahremanzadeh et al. [17] highlight the necessity of integrated location-allocation and fuzzy-robust energy consumption decisions for drone networks, while Liu et al. [24] emphasize the operational impact of launch and recovery siting coupled with tiered routing strategies for agricultural drones. Nevertheless, the specific combination considered in this study—robust battery-capacity degradation, synchronized motorcycle-supported UAV battery delivery, strategic depot selection, and real-world agricultural surveillance validation—remains a unique contribution. Therefore, the present study complements this literature by extending uncertainty-aware drone location-routing toward a coordinated agricultural surveillance setting.

2.6 Gap analysis

While the reviewed literature has made substantial progress in UAV routing, facility location, ground–UAV coordination, and robust optimization, the existing body of research still does not fully address the requirements of large-scale agricultural UAV surveillance networks. A clear gap emerges from the way strategic and operational decisions are commonly treated. In many studies, depot location is determined first, and routing decisions are optimized afterward. This sequential structure overlooks the strong interdependence between infrastructure placement, route feasibility, energy limitations, and operational cost. Another important limitation concerns ground–UAV coordination. Existing models mainly focus on truck–drone delivery, launch-and-recovery operations, or general vehicle–drone cooperation. However, regional agricultural surveillance requires a different form of coordination, in which ground vehicles must support UAV endurance by delivering charged batteries to swap stations before UAV arrival. This temporal synchronization of battery logistics remains insufficiently modeled in the current literature. Furthermore, as highlighted in a recent synthesis of UAV routing literature by Taheri et al. [42], effectively managing battery constraints and dynamic environmental factors remains a critical hurdle for regional drone operations. While that review identifies the growing need for adaptive and robust algorithms, it confirms that the specific temporal synchronization of ground-supported battery logistics under capacity uncertainty remains insufficiently modeled. Battery uncertainty also represents a critical unresolved issue. Although empirical studies show that UAV battery capacity can degrade due to temperature, aging, and charge-cycle effects, most routing models still rely on deterministic battery-capacity assumptions. As a result, routes that appear feasible under nominal conditions may become unsafe or infeasible under degraded battery performance.

The literature also shows a limited connection between theoretical optimization models and real agricultural deployment conditions. Many UAV-routing studies are tested on stylized or small-scale numerical instances, while geographically dispersed agricultural surveillance networks require models that can capture real depot locations, charging stations, orchard distribution, heterogeneous vehicles, and

Table 1: Gap analysis and methodological positioning of related literature

Study	**Veh. routing / scheduling model**	**Facility / multi-depot loc.**	**Bi-level structure**	**Heterog. ground–UAV coordination**	**Robust opt. / uncertainty**	**Agri. surveillance / precision-ag case**
Gebbers & Adamchuk [15]	–	–	–	–	–	✓
Khanal et al. [23]	–	–	–	–	–	✓
Sankaran et al. [34]	–	–	–	–	–	✓
Maes & Steppe [26]	–	–	–	–	–	✓
Zhang et al. [50]	–	–	–	–	–	✓
Candiago et al. [5]	–	–	–	–	–	✓
Tsouros et al. [45]	–	–	–	–	–	✓
Hajimirzajan et al. [20]	–	–	–	–	–	✓
Dorling et al. [13]	✓	–	–	–	–	–
Aggarwal & Kumar [1]	✓	–	–	–	–	–
Shavarani et al. [38]	–	✓	–	–	–	–
Rabta et al. [31]	✓	✓	–	–	–	–
Timilsina et al. [44]	–	–	–	–	✓	–
Bertsimas & Sim [3]	–	–	–	–	✓	–
Murray & Chu [28]	✓	–	–	✓	–	–
Schermer et al. [35]	✓	–	–	✓	–	–
Matese et al. [27]	–	–	–	–	–	✓
Cengiz et al. [8]	✓	–	–	✓	–	✓
Thiels et al. [43]	✓	–	–	–	–	–
Cengiz et al. [7]	✓	–	–	✓	–	–
Yılmaz et al. [47]	✓	–	–	✓	–	–
Yuan et al. [49]	✓	–	–	✓	✓	–
Long et al. [25]	✓	–	–	✓	✓	–
Wang et al. [46]	✓	–	–	✓	✓	–
Nagy & Salhi [29]	✓	✓	–	–	–	–
Daskin [11]	–	✓	–	–	–	–
Rajabi et al. [32]	–	✓	–	–	✓	–
Shadkam [36]	–	✓	–	–	–	–
Parvasi et al. [30]	✓	✓	✓	–	–	–
Choi et al. [9]	✓	–	✓	–	✓	–
Ben-Tal et al. [2]	–	–	–	–	✓	–
Shebalov & Klabjan [39]	✓	–	–	–	✓	–
Shahnaghi et al. [37]	✓	–	–	–	✓	–
Ghahremanzadeh et al. [17]	✓	✓	–	–	✓	–
Liu et al. [24]	✓	✓	–	–	–	✓
Hamid et al. [21]	✓	✓	–	✓	–	–
This study	✓	✓	✓	✓	✓	✓

operational uncertainty. To clarify these gaps, Table 1 compares representative studies across six dimensions: vehicle routing or scheduling, facility or multi-depot location, bi-level structure, heterogeneous ground–UAV coordination, robust optimization or uncertainty modeling, and agricultural surveillance case validation. The gap analysis presented in Table 1 shows that previous studies address some of these dimensions separately, but they do not jointly consider robust strategic depot location, operational UAV–motorcycle routing, synchronized battery logistics, and real agricultural surveillance validation within a unified framework. Based on the identified gaps and the comparative evidence summarized in Table 1, the main novelties of this study are presented below.

2.7 Theoretical and practical novelty of the study

Based on the gaps identified above and the comparative evidence summarized in Table 1, the novelty of this study is divided in two categories: theoretical novelty and practical novelty. The proposed framework integrates several established methodological components; however, their integration is problem-driven rather than method-driven. Each component is included to address a specific requirement of the agricultural surveillance planning problem. The bi-level structure captures the hierarchy between strategic depot planning and operational routing decisions. Robust optimization accounts for battery-capacity degradation where it directly affects UAV route feasibility. The BWM–TOPSIS module translates depot suitability considerations into quantitative inputs for the upper-level objective. k-means clustering is used as a spatial decomposition tool to make large geographically embedded routing instances computationally tractable. Therefore, these components are not presented as independent methodological novelties; rather, they are connected through their roles in the proposed strategic–operational decision framework.

2.7.1 Theoretical novelty

The theoretical novelty of this study lies in formulating a robust strategic–operational location–routing model for cooperative motorcycle–UAV agricultural surveillance under battery-capacity uncertainty.

- **Integrated robust bi-level location–routing model:** The proposed Stackelberg model links upper-level depot-opening and garden-assignment decisions with lower-level routing, synchronization, and energy-feasibility decisions. Therefore, strategic infrastructure decisions are evaluated based on their operational consequences.
- **Motorcycle-supported UAV battery logistics:** The lower-level model explicitly represents motorcycles as both ground-surveillance vehicles and mobile battery-supply agents. This creates a direct operational connection between UAV routes, motorcycle routes, charging-station visits, battery inventory, and swap availability.
- **Temporal synchronization for battery swapping:** The model enforces that motorcycles must arrive at charging stations before UAVs require battery replacement. This synchronization is essential for making extended UAV missions feasible and is not sufficiently captured in existing UAV routing models.
- **Robust battery-capacity modeling:** Battery-capacity uncertainty is embedded into the UAV energy-feasibility constraints using the Bertsimas–Sim framework. This allows the model to protect routes against battery degradation while preserving a tractable mixed-integer linear structure.

2.7.2 Practical novelty

The practical novelty of this study lies in translating the proposed model into a decision-support framework for real regional agricultural surveillance planning.

- **Dual-purpose pistachio surveillance planning:** The proposed framework is designed for geographically dispersed pistachio orchards where UAV operations must simultaneously support agronomic monitoring, such as irrigation-stress and pest detection, and security surveillance, such as theft deterrence during the harvest period.
- **Extension of UAV operational range:** Motorcycle-based battery delivery enables UAVs to perform longer missions through mid-mission battery swaps instead of repeatedly returning to depots.

- **Real agricultural case validation:** The framework is applied to a geographically dispersed pistachio surveillance network in Kerman Province, Iran, including candidate depots, charging stations, and orchard locations. This demonstrates applicability beyond simplified numerical examples.

3 Problem formulation and mathematical model

This section presents the mathematical framework for the robust bi-level optimization problem integrating strategic depot location decisions with operational heterogeneous vehicle routing under battery-capacity uncertainty. The problem is formulated as a Stackelberg game where the leader (strategic planner) selects depot locations and assigns gardens, anticipating the follower's (operational manager) optimal routing response. The deterministic model is first developed assuming nominal parameters and then extended to a robust optimization formulation using the Bertsimas–Sim approach to address battery capacity uncertainty.

3.1 Problem statement

The problem addressed in this study is the design and operation of a regional UAV-based surveillance system for geographically dispersed pistachio orchards. The system must jointly determine depot locations, garden–depot assignments, and coordinated UAV–motorcycle routes under battery-endurance limitations. The problem is further complicated by mid-mission battery swaps, motorcycle-supported battery delivery, and uncertainty in UAV battery capacity.

3.1.1 Surveillance network setting

The problem arises from agricultural surveillance operations in Kerman Province, Iran, where dispersed pistachio orchards require repeated monitoring. The surveillance system supports two operational purposes: agronomic monitoring, such as plant-health assessment, irrigation-stress detection, and pest observation; and security monitoring, such as theft deterrence during harvest periods.

The network consists of a set of pistachio orchards, referred to as gardens, $I = \{1,2,\ldots,|I|\}$, a set of candidate depot locations, $J = \{1,2,\ldots,|J|\}$, and a set of charging or battery-swap stations, $S \subseteq N$. A subset of candidate depots is selected to support surveillance operations. Let D denote the set of depots. The universal node set is defined as $N = D \cup I \cup S$, and is further partitioned into subsets $N_1,\ldots,N_4$ in the detailed formulation, presented in Section 3.2, to distinguish between garden nodes, charging-station nodes, and their technical copies for routing and energy accounting.

3.1.2 Vehicle roles and battery-swap operations

The heterogeneous vehicle fleet is represented by the set $K = K_m \cup K_{dr}$, where K_{dr} contains electric quad-plane UAVs operating under a fundamental dual-battery heterogeneous architecture that creates an endurance bottleneck and necessitates mid-mission battery swaps. The UAV energy-consumption rate on each arc is obtained from the calibrated linear energy model described in the Supplementary Material. For a given UAV and predefined operating and payload condition, the resulting arc-specific rate e_{ijk} is treated as a deterministic input parameter. Accordingly, the total energy consumed by UAV k while traversing arc (i,j) is expressed as $e_{ijk}TT_{ijk}$, where TT_{ijk} denotes the corresponding travel time. The detailed derivation, calibration, and validation of the energy model are provided in the Supplementary Material.

In contrast, K_m contains ground-support motorcycles with travel speed v_k^m and battery-carrying capacity CN_k. The quad-planes execute aerial garden surveillance using multispectral and thermal imaging, enabling the detection of irrigation stress 5–10 days before visible symptoms, identification of pest infestations, and nighttime security monitoring for theft deterrence. Motorcycles perform two functions that are essential to the overall surveillance system: (i) ground-level perimeter and orchard surveillance to detect unauthorized presence or infrastructure damage; (ii) mobile battery logistics by transporting fully charged battery sets from selected depots to strategically positioned charging stations throughout the service region.

Motorcycles $k \in K_m$ depart from opened depots $j \in J$, carrying multiple fully charged battery sets within capacity CN_k, traverse predetermined routes visiting battery charging or swap stations $s \in S$, and then return to their depot. Simultaneously, quad-planes $k \in K_{dr}$ depart from the same depots, visit assigned gardens $i \in I$, and, when their remaining battery energy approaches the minimum safety energy MEB_k, divert to a pre-positioned charging station for rapid battery swapping.

Each swap incurs a service cost CC_k and restores the quad-plane battery to full capacity, enabling immediate resumption of extended surveillance missions. The motorcycle must arrive at the charging station no later than the quad-plane to ensure that fully charged batteries are physically available when the aircraft reaches the station and requests a swap.

3.1.3 Strategic and operational decisions

The problem includes two interconnected decision layers. At the strategic layer, the planner decides which candidate depots should be opened and how gardens should be assigned to opened depots. Binary variables Y_j indicate whether depot $j \in J$ is opened, and binary variables X_{ij} indicate whether garden $i \in I$ is assigned to depot $j \in J$.

The desirability score of a candidate depot is not determined by distance alone. Accessibility, power-supply reliability, security, and establishment cost also influence depot suitability. These factors are represented by the parameter $s_j \in [0,1]$, obtained from the depot evaluation procedure described in the Supplementary Material, and are used in the upper-level objective together with assignment-distance costs.

At the operational layer, for a given depot-opening and garden-assignment configuration, the model first determines which gardens should be visited by each UAV and which gardens and charging-station villages should be visited by each motorcycle. It then determines the routing and scheduling of UAV and motorcycle visits, including the visiting sequence, arrival and departure times, motorcycle battery delivery to charging stations, and UAV battery-swap timing. The operational objective is to minimize fixed vehicle-use, travel, surveillance-service, and battery-swap costs. Lower-level routing decisions are constrained to strictly respect the upper-level garden–depot assignment structure, ensuring consistency between strategic allocation and operational service execution

3.1.4 Modeling assumptions

The proposed location-routing problem is developed under the following assumptions:

1. The sets of gardens, candidate depots, and charging stations are known in advance.

2. Each garden receives exactly one primary surveillance visit during the planning horizon.
3. Each active vehicle starts from and returns to an assigned opened depot.
4. Opened depots can host both UAVs and motorcycles.
5. UAVs and motorcycles have vehicle-specific travel times, operating costs, and maximum mission durations.
6. UAVs fly point-to-point, while motorcycles use the road network; therefore, their travel time matrices may differ.
7. Motorcycles have limited battery-carrying capacity.
8. The UAV energy-consumption rate on each arc is modeled as a deterministic linear function of the transported load or payload condition in the base model; the total arc energy consumption is obtained by multiplying this rate by the corresponding travel time.
9. Battery swaps are allowed only at charging stations and require prior motorcycle battery delivery.
10. Subtours disconnected from opened depots are not allowed.
11. Battery-capacity uncertainty affects UAV energy feasibility, while other parameters are treated as deterministic. (This selective robustification focuses protection strictly on the primary operational bottleneck—UAV flight endurance—to avoid excessive conservatism and computational intractability. Importantly, as proven mathematically in Section 3.5, applying budgeted uncertainty at the vehicle battery level ($\widetilde{BE}_k$) acts as an *integrated energy buffer* that algebraically provides an integrated endurance buffer against cumulative increases in route-level energy consumption, as discussed in Section 3.5.)

This structure separates strategic depot location and garden assignment decisions from operational routing decisions, while preserving their interdependence through the mathematical formulation presented in the following sections.

3.2 Notations

To improve the clarity and readability of the mathematical formulation, the notation used in the model is presented in Table 2.

Table 2: Main notations used in the mathematical formulation

Category	Symbol	Definition
Sets	$\mathscr{I}$	Set of pistachio orchards, referred to as gardens
	$\mathscr{J}$	Set of candidate depot locations
	$\mathscr{K}$	Set of all vehicles, $\mathscr{K} = \mathscr{K}_m \cup \mathscr{K}_{dr}$
	$\mathscr{K}_m$	Set of motorcycles
	$\mathscr{K}_{dr}$	Set of drones, or UAVs

Continued on next page

Category	Symbol	Definition
	$\mathcal{N}$	Universal node set, $\mathcal{N} = \mathcal{D} \cup \mathcal{I} \cup \mathcal{S}$
	$\mathcal{D}$	Set of depot nodes, $\mathcal{D} \subseteq \mathcal{J}$
	$\mathcal{S}$	Set of charging or battery-swap stations, $\mathcal{S} \subseteq \mathcal{N}$
	$\mathcal{N}_1$	Set of gardens and charging stations
	$\mathcal{N}_2$	Set of gardens only
	$\mathcal{N}_3$	Set of charging stations only
	$\mathcal{N}_4$	Set of gardens and charging stations and their copies
	$\mathcal{K}_d$	Set of vehicles stationed at depot $d \in \mathcal{D}$, $\mathcal{K}_d \subseteq \mathcal{K}$
	$\mathcal{K}_d^{dr}$	Set of UAVs stationed at depot d, defined as $\mathcal{K}_d \cap \mathcal{K}_{dr}$
	Ω	Set of charging stations and their copies
Parameters	$Elig_{jk}$	Equal to 1 if vehicle k is allowed to visit node j; 0 otherwise
	TT_{ijk}	Travel time from node i to node j by vehicle k, in hours
	MA_k	Maximum availability or mission time of vehicle k, in hours
	B_i	Number of batteries picked up from node i by a motorcycle
	λ_i	Number of batteries delivered to station i by a motorcycle
	CD_k	Fixed cost of using vehicle k
	CD'_k	Variable cost per unit time for vehicle k
	CC_k	Battery-swap cost for drone k
	MEB_k	Minimum energy required by drone k to ensure a safe flight
	dur_{ik}	Service or operation time of vehicle k at node i
	BE_k	Nominal battery capacity of drone k
	e_{ijk}	Energy consumption rate of drone k on arc (i, j) calculated as a linear function of the transported load or payload condition
	CN_k	Battery-carrying capacity of motorcycle k
	M	Sufficiently large positive constant, or big-M
	s_j	TOPSIS depot appropriateness score, $s_j \in [0, 1]$
	d_{ij}	Distance between garden i and depot j, in km
	$R_{\max}$	Drone range, defined as the maximum distance between two stopping places
	p	Number of depots allowed to open
Upper-level variables	Y_j	Equal to 1 if depot j is opened; 0 otherwise
	X_{ij}	Equal to 1 if garden i is assigned to depot j; 0 otherwise
Lower-level variables	x_{ijk}	Equal to 1 if vehicle k travels from node i to node j; 0 otherwise
	hh_k	Equal to 1 if vehicle k is used; 0 otherwise
	t_{jk}	Departure time of vehicle k at node j
	u_{jk}	Load, in number of batteries, carried by motorcycle k upon arriving at node j
	y_{jk}	Remaining battery charge of drone k upon arriving at node j

3.3 Bi-level mathematical formulation

The proposed framework follows a computational Stackelberg structure in which upper-level decisions are evaluated based on the optimal response of the lower-level routing problem. This ensures that strategic depot-location decisions are directly linked to operational routing performance through a feedback mechanism. For clarity, the mathematical formulation is structured into two interconnected components. Section 3.3.1 describes the upper-level strategic model, including depot opening and garden–depot assignment decisions, while Section 3.3.2 presents the lower-level operational model, including UAV–motorcycle routing, scheduling, battery delivery, synchronization, and energy-feasibility constraints. Although the formulation is presented in two separate subsections for clarity, the two levels remain fully coupled through the lower-level optimal response embedded in the upper-level evaluation.

3.3.1 Upper-level model

The leader maximizes strategic utility of selected depots (via TOPSIS scores) and minimizes distance-related assignment costs, anticipating the follower's optimal routing cost Z_{LL}. In the upper-level objective, the parameter s_j represents the strategic suitability score of candidate depot j, obtained from the BWM–TOPSIS evaluation. This term incorporates non-routing infrastructure considerations into the strategic decision layer, while the assignment-distance term captures the geographic efficiency of garden–depot allocation. The upper-level optimization model is formulated as follows:

$$
\begin{aligned}
\max \quad & Z_{UL} = \sum_{j \in \mathscr{J}} s_j Y_j - \alpha \sum_{i \in \mathscr{I}} \sum_{j \in \mathscr{J}} d_{ij} X_{ij}, && && (1)\\
\text{s.t.} \quad & \sum_{j \in \mathscr{J}} X_{ij} = 1, && \forall i \in \mathscr{I}, && (2)\\
& X_{ij} \le Y_j, && \forall i \in \mathscr{I},\ j \in \mathscr{J}, && (3)\\
& d_{ij} X_{ij} \le R_{\max} X_{ij}, && \forall i \in \mathscr{I},\ j \in \mathscr{J}, && (4)\\
& \sum_{j \in \mathscr{J}} Y_j = p, && && (5)\\
& Y_j, X_{ij} \in \{0,1\}, && \forall i \in \mathscr{I},\ j \in \mathscr{J}. && (6)
\end{aligned}
$$

Equation (1) defines the upper-level objective function, which evaluates each candidate depot configuration based on two strategic components: (i) the aggregated depot attractiveness measured through TOPSIS scores, and (ii) the total assignment cost associated with allocating gardens to selected depots. This formulation captures the strategic trade-off between infrastructure quality and spatial allocation cost at the leader level. However, the operational performance of any selected configuration cannot be fully assessed without considering the induced routing decisions at the lower level. Therefore, each upper-level decision (Y,X) is evaluated through the optimal response of the follower problem. $Z^*_{LL}(Y,X)$ is defined as the optimal value of the lower-level routing problem obtained by solving the follower problem for each upper-level decision (Y,X). This term ensures the feedback coupling between the two decision levels. Constraints (2)–(6) define the strategic assignment structure. Constraint (2) enforces that every garden is assigned to exactly one depot, guaranteeing full surveillance coverage and avoiding split responsibility. Constraint (3) couples assignment and opening decisions by forbidding assignments to closed depots. Constraint (4) imposes a hard service-radius feasibility threshold, such that a garden can be assigned to depot j only if its distance does not exceed $R_{\max}$. Thus, $R_{\max}$ is treated as a strict feasibility limit rather than as a soft preference or penalty parameter. Constraint (5) limits the number of opened depots to p, implementing a budgetary or capacity limitation on infrastructure investment. Constraint (6) enforces binary integrality for strategic decisions. Thus, upper-level decisions are explicitly evaluated through the induced optimal operational cost of the lower-level problem, ensuring full interaction between strategic and operational levels.

3.3.2 Lower-level model

For each fixed upper-level configuration (Y,X), the lower-level operational problem is formulated as a mixed-integer program. For small non-decomposed instances, this problem is solved directly, whereas full-scale instances are evaluated through the clustered matheuristic described in Section 4. The resulting operational cost is then used in the configuration-level evaluation. Given (Y,X), the follower minimizes total operational cost includes fixed vehicle costs, variable routing costs, surveillance costs, and battery-swap costs, which are formulated as follows:

$$\begin{aligned}
\min \quad & Z_{LL} = \sum_{k\in\mathscr{K}} CD_k hh_k + \sum_{k\in\mathscr{K}} \sum_{i\in\mathscr{N}} \sum_{j\in\mathscr{N}} CD'_k TT_{ijk} x_{ijk} \\
& \quad + \sum_{k\in\mathscr{K}} \sum_{g\in\mathscr{I}} dur_{gk} CD'_k \sum_{i\in\mathscr{N}} x_{igk} + \sum_{k\in\mathscr{K}_{dr}} \sum_{i\in\mathscr{N}} \sum_{j\in\Omega} CC_k x_{ijk}, && (7)\\
\text{s.t.} \quad & \sum_{j\in\mathscr{I}\cup\mathscr{S}} x_{djk} = hh_k, && \forall d\in\mathscr{D},\ k\in K_d, \quad (8)\\
& \sum_{j\in\mathscr{N}} x_{djk} \le Y_d, && \forall d\in\mathscr{D},\ k\in\mathscr{K}, \quad (9)\\
& \sum_{k\in\mathscr{K}} \sum_{i\in\mathscr{N}, i\neq j} x_{ijk} = 1, && \forall j\in\mathscr{I}, \quad (10)\\
& \sum_{j\in\mathscr{I}\cup\mathscr{S}} x_{djk} = \sum_{i\in\mathscr{I}\cup\mathscr{S}} x_{idk}, && \forall d\in\mathscr{D},\ k\in K_d, \quad (11)\\
& \sum_{i\in\mathscr{N}, i\neq j} x_{ijk} \le Elig_{jk} hh_k, && \forall j\in\mathscr{N},\ k\in\mathscr{K}, \quad (12)\\
& \sum_{k\in\mathscr{K}_d^{dr}} \sum_{i\in\mathscr{N}, i\neq g} x_{igk} + \sum_{k\in\mathscr{K}_d^{m}} \sum_{i\in\mathscr{N}, i\neq g} x_{igk} = X_{gd}, && \forall g\in\mathscr{I}, d\in\mathscr{D}, \quad (13)\\
& \sum_{k\in\mathscr{K}_m} \sum_{i\in\mathscr{N}_4} x_{isk} \le 1, && \forall s\in\mathscr{N}_3, \quad (14)\\
& \sum_{i\in\mathscr{N}, i\neq h} x_{ihk} = \sum_{j\in\mathscr{N}, j\neq h} x_{hjk}, && \forall h\in\mathscr{I}\cup\mathscr{S},\ k\in\mathscr{K}. \quad (15)
\end{aligned}$$

Constraints (8)–(15) impose the routing and visit structure. Constraint (8) links vehicle activation and depot departure: an active vehicle must depart exactly once from its assigned depot, while inactive vehicles do not leave the depot. Constraint (9) ensures that a vehicle can depart from depot d only when the corresponding depot is selected in the strategic level ($Y_d = 1$). Therefore, operational routes cannot originate from non-selected depots, which improves the consistency between the strategic network design and the operational routing structure. Constraint (10) ensures that each garden is visited exactly once by the fleet, the classical VRP "customer visited once" condition. Constraint (11) enforces flow conservation at depots so that each vehicle that leaves must return, yielding closed tours. Constraint (12) restricts movements to eligible node–vehicle combinations and forbids inactive vehicles from appearing in routes. Constraint (13) couples the upper-level garden–depot assignment with lower-level service decisions. If garden g is assigned to depot d, it must be served by exactly one eligible UAV or motorcycle stationed at that depot; vehicles associated with other depots cannot serve that garden. Constraint (14) states that each village, if considered as a charging station, can be visited by a maximum of one motorcycle to install batteries. Constraint (15) is the general flow-conservation constraint at intermediate nodes, maintaining route continuity at intermediate nodes.

The temporal constraints are formulated as follows:

$$t_{jk} \geq TT_{djk} + dur_{jk} - M(1 - x_{djk}), \qquad \forall d \in \mathscr{D},\ j \in \mathscr{N}_1,\ k \in \mathscr{K}, \tag{16}$$
$$t_{jk} \geq t_{ik} + dur_{ik} + TT_{ijk} - M(1 - x_{ijk}), \qquad \forall i \in \mathscr{N}_1,\ j \in \mathscr{N},\ k \in \mathscr{K},\ i \neq j, \tag{17}$$
$$t_{jk} \leq MA_k \sum_{i \in \mathscr{N}, i \neq j} x_{ijk}, \qquad \forall j \in \mathscr{N},\ k \in \mathscr{K}. \tag{18}$$

Constraint (16) initializes arrival time at the first visited node after leaving a depot. If arc (d, j) is used, the arrival time must be at least travel plus service; otherwise, the big–M deactivates the constraint. Constraint (17) propagates times along the route: when arc (i, j) is traversed, departure/arrival at j must respect service and travel from i, preventing time-inconsistent paths. Constraint (18) caps arrival times by vehicle endurance MA_k, guaranteeing that all operations of each vehicle occur within its duty period.

The battery inventory constraints are defined as follows:

$$u_{jk} \geq u_{ik} + B_i - \lambda_i - M(1 - x_{ijk}), \qquad \forall k \in \mathscr{K}_m,\ i \in \mathscr{N}_3,\ j \in \mathscr{N}_4,\ i \neq j, \tag{19}$$
$$u_{dk} = \sum_{j \in \mathscr{N}_3} \sum_{i \in \mathscr{N}, i \neq j} x_{ijk} \lambda_j, \qquad \forall d \in \mathscr{D},\ k \in \mathscr{K}_m, \tag{20}$$
$$u_{jk} \leq CN_k hh_k, \qquad \forall j \in \mathscr{N},\ k \in \mathscr{K}_m. \tag{21}$$

Constraint (19) propagates a lower bound on the motorcycle battery inventory along a traversed arc, accounting for battery pickups and deliveries, while the big-M term deactivates the relation when the arc is not selected. Constraint (20) initializes the load at each depot as the sum of all batteries that will eventually be delivered to charging stations by that motorcycle. Constraint (21) enforces physical capacity; motorcycles cannot carry more than CN_k batteries when active and must have zero load when unused.

Temporal synchronization is enforced as follows:

$$t_{sk_{dr}} \geq t_{sk_m} - M\left(2 - \sum_{i \in \mathscr{N}} x_{isk_m} - \sum_{i \in \mathscr{N}} x_{isk_{dr}}\right), \qquad \forall s \in \mathscr{S},\ k_m \in \mathscr{K}_m,\ k_{dr} \in \mathscr{K}_{dr}. \tag{22}$$

Constraint (22) ensures temporal synchronization between motorcycles and drones at charging stations. When both a given motorcycle and drone visit station s (both summations of x_{isk} equal 1), the big–M terms vanish and the constraint enforces that the drone cannot arrive before the motorcycle; hence, batteries are always available before a swap is attempted. If either vehicle does not visit s, the constraint is relaxed.

The deterministic UAV energy accounting is implemented as follows:

$$y_{jk} \leq y_{ik} - e_{ijk} TT_{ijk} + M(1 - x_{ijk}), \qquad \forall k \in \mathscr{K}_{dr},\ i \in \mathscr{N}_2,\ j \in \mathscr{N},\ i \neq j, \tag{23}$$
$$y_{jk} \leq BE_k - e_{ijk} TT_{ijk} + M(1 - x_{ijk}), \qquad \forall k \in \mathscr{K}_{dr},\ i \in \Omega \cup \mathscr{D},\ j \in \mathscr{N},\ i \neq j, \tag{24}$$

$$y_{jk} \geq MEB_k \sum_{i \in \mathscr{N}, i \neq j} x_{ijk}, \qquad \forall k \in \mathscr{K}_{dr},\ j \in \mathscr{N}. \tag{25}$$

In these constraints, e_{ijk} denotes the energy-consumption rate per unit flight time of UAV k on arc (i, j). This rate is calculated using a calibrated linear approximation as a function of the transported load or payload condition. Therefore, the total energy consumed on arc (i, j) is expressed as $e_{ijk}T_{ijk}$. Constraint (23) updates the remaining battery when a UAV flies from a garden node i to node j; if arc (i, j) is used, the remaining battery at node j cannot exceed the remaining battery at node i minus the time-dependent energy consumption $e_{ijk}T_{ijk}$. Constraint (24) plays an analogous role for departures from depots or charging stations, where the UAV battery is reset to the nominal full capacity before subtracting the energy consumed on the next arc. Constraint (25) enforces a minimum safety buffer whenever node j lies on the UAV route, guaranteeing sufficient residual energy for contingencies or return-to-base operations.

The variable domains are defined as follows:

$$x_{ijk}, hh_k \in \{0, 1\}, \qquad \forall i, j \in \mathscr{N},\ k \in \mathscr{K}, \tag{26}$$

and

$$\begin{aligned} &t_{jk} \geq 0, && \forall j \in \mathscr{N},\ k \in \mathscr{K}, \\ &u_{jk} \geq 0, && \forall j \in \mathscr{N},\ k \in \mathscr{K}_m, \\ &0 \leq y_{jk} \leq BE_k, && \forall j \in \mathscr{N},\ k \in \mathscr{K}_{dr}. \end{aligned} \tag{27}$$

Constraints (26) and (27) define integrality for routing and activation decisions and non-negativity for time, inventory, and energy variables.

This structure represents a fully embedded bi-level optimization framework in which the upper-level problem evaluates each configuration through the optimal response of the lower-level routing model, ensuring explicit and complete feedback between the two decision levels.

To clarify the role of each group of constraints in the lower-level operational model, Table 3 summarizes the constraint hierarchy and the corresponding modeling purpose.

Remark 1. *Big-M Bounds and LP-Relaxation Strengthening.* To limit numerical instability and excessive slack in the linear programming (LP) relaxation, the big-M constants are not assigned arbitrarily large values. Instead, they are derived from physical and operational bounds associated with vehicle mission time, battery capacity, and motorcycle carrying capacity. These bounds provide valid deactivation coefficients; however, physical validity alone does not necessarily imply the strongest possible LP relaxation. Accordingly, the UAV energy-propagation constraint is further strengthened through the lifted valid inequality described below.

For the UAV energy-propagation constraint, we acknowledge that the physically valid choice $M = BE_k$ does not necessarily yield the strongest LP relaxation. Therefore, while Constraint (23) is retained in its original operational form, the formulation is supplemented by the following lifted valid inequality:

$$\begin{aligned} y_{jk} \leq\ & y_{ik} - e_{ijk}TT_{ijk} + \left(BE_k - MEB_k + e_{ijk}TT_{ijk}\right)(1 - x_{ijk}) \\ & + MEB_k \left(1 - \sum_{\substack{h \in \mathscr{N} \\ h \neq i}} x_{hik}\right), \qquad \forall k \in \mathscr{K}_{dr},\ i \in \mathscr{N}_2,\ j \in \mathscr{N},\ i \neq j. \end{aligned} \tag{23a}$$

Table 3: Constraint hierarchy of the lower-level operational model

Constraint group	Equations	Purpose
Vehicle activation and routing	(8)–(15)	Ensure vehicle activation, routing feasibility, depot-based service consistency, and explicit coupling between upper-level assignment variables and lower-level routing decisions.
Time propagation	(16)–(18)	Track vehicle arrival/departure times and enforce maximum mission duration
Motorcycle battery inventory	(19)–(21)	Track the number of charged batteries carried and delivered by motorcycles
UAV–motorcycle synchronization	(22)	Ensure that batteries are delivered to charging stations before UAV battery-swap demand
UAV energy feasibility	(23)–(25)	Track UAV battery consumption, replenishment, and minimum safety reserve
Variable domains	(26)–(27)	Define binary and continuous variable restrictions

When arc (i, j) is selected, the node-visit and arc-deactivation terms vanish, and the lifted inequality reduces to the original energy-propagation relation. When the arc is inactive, the minimum residual-energy requirement in Constraint (25), together with the explicit battery-state bound $0 \leq y_{jk} \leq BE_k$, provides a tighter state-dependent relaxation. The additional node-visit term preserves validity when node i is not visited. Thus, Constraint (23a) is added as a valid strengthening inequality and does not alter the integer-feasible operational logic of Constraint (23).

For the remaining big-M constraints, the deactivation coefficients are linked to the corresponding physical or operational bounds. The time-propagation and synchronization constraints use vehicle mission-time and travel-time bounds, while the motorcycle inventory constraint is bounded using the physical battery-carrying capacity. For the UAV energy constraints, the battery-capacity bound is retained, with Constraint (23) additionally strengthened by the lifted valid inequality (23a). These constraint-specific bounds reduce excessive deactivation slack and improve the numerical representation of the model, although they do not imply an ideal LP relaxation. Table 4 summarizes the adopted big-M bounds and the additional strengthening applied to the UAV energy-propagation constraint.

Table 4: Big-M bounds and formulation-strengthening measures used in the mathematical model

Constraint NO.	Big-M bound / treatment	Constraint NO.	Big-M bound / treatment
Constraint (16)	$TT_{djk} + dur_{jk}$	Constraint (22)	MA_{k_m}
Constraint (17)	$MA_k + dur_{ik} + TT_{ijk}$	Constraint (23)	BE_k; strengthened by Constraint (23a)
Constraint (19)	$\max\{0, CN_k + B_i - \lambda_i\}$	Constraint (24)	BE_k

3.4 Performance evaluation of the proposed bi-level framework

In this subsection, we evaluate the performance of the proposed bi-level decomposition framework using a set of 10 small-scale test instances. The objective is to validate the correctness of the proposed formulation by comparing it against an exact monolithic mixed-integer programming model. To ensure a controlled experimental setting, all candidate locations for depot establishment are assigned identical appropriateness scores. The deterministic version of the problem is used as the benchmark for evaluation, serving as the exact reference model for comparing the proposed bi-level decomposition framework. The deterministic baseline corresponds to the full integrated location–routing model without decomposition, where all model parameters are treated as known and fixed. This assumption removes the influence of multi-criteria decision-making weights and ensures that depot selection is driven purely by the optimization structure of the model. It is also assumed that exactly one depot is selected in each instance, consistent with the single-facility design of the system.

The test instances are synthetically generated and deliberately designed at a very small scale, including 2–5 gardens, 1–4 candidate depot locations, 2–3 villages, and exactly 1 UAV and 1 motorcycle in all instances. This guarantees that the monolithic formulation remains solvable to optimality within very short computational time using a commercial solver. The spatial coordinates of all nodes (gardens, depots, and villages) are randomly generated within a square area of 7×7 km^2 to ensure a realistic yet controlled spatial distribution of the network. To provide a rigorous benchmark for evaluating the performance of the proposed framework, we introduce a deterministic baseline model. In this baseline, all uncertain or stochastic elements are removed, and all parameters are treated as fixed deterministic values. The deterministic model corresponds to the full integrated location–routing formulation without decomposition and serves as the exact optimization benchmark. This model is solved using the same Gurobi solver under identical computational settings (Python–Pyomo environment) to ensure a fair comparison. The proposed bi-level framework is then compared against this deterministic baseline in terms of solution quality and computational performance.

Within this controlled environment, two modeling paradigms are evaluated. The first is a monolithic integrated location–routing formulation, which simultaneously determines single-depot selection, garden assignment, and routing decisions within a unified mixed-integer programming model. This formulation serves as the exact benchmark and is solved to optimality using the Gurobi solver. The second paradigm corresponds to the proposed bi-level decomposition framework, in which the upper level determines the single depot selection and garden assignment decisions, while the lower level solves the corresponding vehicle routing and energy-feasibility problem as a mixed-integer program.

Since the objective of this section is validation of solution quality, only the second objective function (Z_2), representing the total operational cost, is used for comparison. This is because Z_2 captures UAV transportation cost, motorcycle operational cost, battery logistics cost, and service-related costs. Moreover, although location decisions are not explicitly present in Z_2, their impact is implicitly reflected through routing structure and travel times. The solution quality is evaluated using the optimality gap defined as:

$$GAP(\%) = \frac{Z_2^{\text{Bi-level}} - Z_2^{\text{Exact}}}{Z_2^{\text{Exact}}} \times 100. \tag{28}$$

The parameter configuration of all instances is reported in Table S.2, which defines UAV and motorcycle operational characteristics, including battery capacities, energy consumption rates, travel speeds, service times, and operational costs. Table 5 summarizes the structure of the generated test instances.

Table 6 presents the computational results for the second objective function (Z_2), comparing the exact monolithic solution with the proposed bi-level decomposition approach in terms of objective value, optimality gap, and computational time.

The results demonstrate that the proposed bi-level decomposition framework consistently produces

Table 5: Structure of very small-scale test instances

Instance	Gardens	Candidate Locations	Villages	UAV	Motorcycle
I1	2	1	2	1	1
I2	2	2	2	1	1
I3	3	1	2	1	1
I4	3	2	2	1	1
I5	3	3	3	1	1
I6	4	3	3	1	1
I7	4	2	3	1	1
I8	4	3	3	1	1
I9	4	4	3	1	1
I10	5	3	3	1	1

Table 6: Performance comparison for objective Z_2

Instance	Z_2 Exact	Z_2 Bi-level	GAP (%)	Time Exact (s)	Time Bi-level (s)
I1	8.40	8.40	0.00	0.65	0.48
I2	8.90	8.95	0.56	1.20	0.85
I3	9.10	9.10	0.00	2.10	1.45
I4	9.80	9.92	1.22	4.80	3.10
I5	10.20	10.35	1.47	8.50	5.60
I6	10.70	10.88	1.68	14.20	9.30
I7	11.10	11.32	1.98	22.80	15.40
I8	11.60	11.86	2.24	31.50	21.20
I9	12.00	12.31	2.58	44.70	29.80
I10	12.40	12.78	3.06	57.90	38.40

high-quality solutions that remain close to the exact monolithic formulation across all small-scale instances. The optimality gap is observed to be below approximately 3.1%, with significantly smaller deviations in the smallest instances and gradually increasing gaps as the problem size grows. This trend is expected due to the increasing interaction complexity between assignment and routing decisions in the bi-level structure. Although the monolithic model is able to solve all instances to optimality within very short computational times at this small scale, the proposed bi-level framework exhibits competitive performance while offering a structurally decomposed formulation that improves the scalability and modularity of the solution process.

3.5 Robust optimization formulation under battery-capacity uncertainty

The deterministic lower-level model in Section 3.3 assumes that the effective battery capacity of each UAV is fixed at its nominal value BE_k^{norm}. However, in practical UAV operations, the usable battery capacity may degrade due to temperature variation, aging, charge-cycle effects, and operational conditions [44]. Therefore, robust optimization is applied selectively within this framework. Rather than robustifying all model parameters, uncertainty is incorporated only in the UAV battery-capacity parameter, which directly affects the energy-feasibility and replenishment constraints. This selective robustification avoids unnecessary conservatism while allowing the lower-level routing response to account for battery-capacity degradation. Following the Bertsimas–Sim budgeted uncertainty framework [3], we construct a tractable robust counterpart of the UAV energy-feasibility constraint. Specifically, we define the effective battery capacity of UAV $k \in \mathscr{K}_{dr}$ as an uncertain *vehicle-level* parameter.

We treat the effective full-charge capacity of UAV $k \in \mathscr{K}_{dr}$ as an uncertain scalar parameter $\widetilde{BE}_k$ that can deviate around its nominal value BE_k^{norm} by at most $\widehat{BE}_k \geq 0$. In the robust formulation, uncertainty is introduced at the UAV battery-capacity level, and the resulting robust counterpart modifies only the energy-feasibility structure of the lower-level model:

$$\widetilde{BE}_k \in \left[BE_k^{\text{nom}} - \widehat{BE}_k,\ BE_k^{\text{nom}} + \widehat{BE}_k,\right], \qquad \forall k \in \mathscr{K}_{dr}. \tag{29}$$

Unlike arc-based uncertainty models, the proposed formulation treats battery capacity BE_k as a vehicle-level uncertain parameter. Therefore, uncertainty affects UAV endurance globally rather than individual travel or replenishment arcs.

Recalling the deterministic energy feasibility constraint in Section 3.3, we reformulate it to explicitly highlight its dependence on the UAV battery-capacity parameter. In this model, the battery capacity of each UAV is treated as a vehicle-level quantity, which constitutes the source of uncertainty in the robust formulation as:

$$y_{jk} + e_{ijk}TT_{ijk} - \widetilde{BE}_k \ \leq\ M\big(1 - x_{ijk}\big), \quad \forall k \in \mathscr{K}_{dr},\ \forall i \in \omega \cup \mathscr{D},\ \forall j \in \mathscr{N},\ i \neq j. \tag{30}$$

Constraint (30) is violated whenever the realized capacity $\widetilde{BE}_k$ is sufficiently small relative to $y_{jk} + e_{ijk}TT_{ijk}$.

Following Bertsimas and Sim [3], we associate a robustness budget Γ_k with each UAV $k \in \mathscr{K}_{dr}$, which controls the maximum allowable deviation in the battery-capacity degradation for that specific vehicle during its mission. To obtain a linear robust counterpart, we introduce auxiliary protection variables $z_k \geq 0$ and $w_{ijk} \geq 0$ and replace the deterministic constraint (24) by the following system:

$$y_{jk} \ \leq\ BE_k^{\text{nom}} - e_{ijk}TT_{ijk} - \big(\Gamma_k z_k + w_{ijk}\big) + M\big(1 - x_{ijk}\big), \quad \forall k \in \mathscr{K}_{dr},\ \forall i \in \omega \cup \mathscr{D},\ \forall j \in \mathscr{N},\ i \neq j, \tag{31}$$

$$z_k + w_{ijk} \ \geq\ \widehat{BE}_k x_{ijk}, \quad \forall k \in \mathscr{K}_{dr},\ \forall i \in \omega \cup \mathscr{D},\ \forall j \in \mathscr{N},\ i \neq j, \tag{32}$$

$$z_k \geq 0, \qquad w_{ijk} \geq 0, \quad \forall k \in \mathscr{K}_{dr},\ \forall i \in \omega \cup \mathscr{D},\ \forall j \in \mathscr{N},\ i \neq j. \tag{33}$$

Constraints (31)–(33) represent the Bertsimas–Sim robust counterpart of the original deterministic constraint (24) under *vehicle-level* uncertainty. The protection term $\Gamma_k z_k + w_{ijk}$ reduces the usable post-replenishment battery capacity below the nominal level BE_k^{norm}, thereby capturing worst-case UAV battery degradation scenarios. The parameter Γ_k controls the trade-off between robustness and conservatism (i.e., the price of robustness), while preserving the computational tractability of the lower-level routing problem. Consistent with the selective robustification strategy, uncertainty is introduced only in the most critical energy-feasibility structure, while all remaining constraints remain deterministic.

Remark 2. *Relationship between Battery-Capacity Degradation and Cumulative Route-Energy Increase.*

A critical environmental consideration in semi-arid agricultural monitoring is that adverse wind speeds, air density variations, and thermal stress can elevate the time-dependent arc energy consumption rate (e_{ijk}) above nominal estimates. Although our robust formulation introduces budgeted uncertainty explicitly into the vehicle battery capacity ($\widetilde{BE}_k$), this selective robustification also provides protection against a cumulative increase in the total energy requirement of a UAV sortie.

Consider any UAV sortie traversing a sequence of arcs $\mathscr{P}$ between replenishment points. Under deterministic conditions, energy feasibility requires $\sum_{(i,j)\in\mathscr{P}} e_{ijk}TT_{ijk} \leq BE_k^{\text{norm}}$. If environmental headwinds increase arc energy consumption along the route by a worst-case factor of $(1+\delta)$ where $\delta \geq 0$, the required feasibility condition is:

$$(1+\delta)\sum_{(i,j)\in\mathscr{P}} e_{ijk}TT_{ijk} \leq BE_k^{\text{norm}} \iff \sum_{(i,j)\in\mathscr{P}} e_{ijk}TT_{ijk} \leq \frac{BE_k^{\text{nom}}}{1+\delta}. \tag{34}$$

Therefore, protecting against an effective battery-capacity reduction of $\widehat{BE}_k$ is algebraically equivalent to hedging against a cumulative arc-energy surge of $\delta = \widehat{BE}_k/(BE_k^{\text{norm}} - \widehat{BE}_k)$. In our Kerman Province case study, the maximum robustness budget ($\Gamma = 1.0$) protects against a 20% battery capacity reduction ($\widetilde{BE}_k = 0.80\,BE_k^{\text{norm}}$). By Eq. (34), this vehicle-level buffer mathematically guarantees route feasibility against up to a $+25\%$ **cumulative increase in flight energy consumption** ($\delta = 0.20/0.80 = 0.25$) across the sortie. Furthermore, because environmental wind gusts across regional orchards exhibit strong spatial correlation rather than independent arc-by-arc randomness, vehicle-level budgeting protects global sortie endurance without introducing the severe linear programming relaxations or over-conservatism associated with independent arc-level uncertainty sets. Finally, we note that treating arc-level energy consumption rates deterministically was a deliberate modeling choice to preserve computational simplicity and analytical tractability within the bi-level framework; extending the robust optimization model to explicitly incorporate multi-factor, arc-specific environmental uncertainties represents a promising direction for future research.

3.5.1 Robust lower-level (follower) model

Under a given leader decision (Y, X), the follower problem remains identical to the deterministic model (Eqs. (7)–(27)) except that the deterministic replenishment constraint (24) is replaced by its Bertsimas–Sim robust counterpart (31)–(33):

$$\begin{aligned} &\min Z_{LL} \\ &\text{ s.t. } (8)--(23),\ (25)--(27) \quad \text{(unchanged)}, \\ &\qquad \text{and } (31)--(33) \text{ replacing } (24). \end{aligned} \tag{35}$$

3.6 Compact integrated formulation

The proposed robust bi-level location–routing model can be summarized as a leader–follower optimization problem. At the upper level, the strategic planner determines the opened depots and the assignment of gardens to opened depots:

$$\begin{aligned} \max_{Y,X} \quad & Z^{UL}(Y,X) \\ \text{s.t.} \quad & (2)--(6). \end{aligned} \tag{36}$$

For each feasible upper-level decision (Y,X), the lower-level problem determines the operational response, including UAV and motorcycle routing, visit scheduling, motorcycle battery delivery, UAV battery tracking, and battery–swap synchronization:

$$\begin{aligned} Z^{LL*}(Y,X) = \min_{x,hh,t,u,y,z,w} \quad & Z^{LL}(x,hh,t,u,y,z,w) \\ \text{s.t.} \quad & (8)--(23),\ (25)--(27),\ (31)--(33). \end{aligned} \tag{37}$$

In the robust formulation, the deterministic UAV replenishment constraint (24) is replaced by the Bertsimas–Sim robust counterpart (31)–(33). Accordingly, the complete robust bi-level model is written compactly as:

$$\begin{aligned} \max_{Y,X} \quad & Z^{UL}(Y,X) \\ \text{s.t.} \quad & (2)--(6), \\ & (x^*,hh^*,t^*,u^*,y^*,z^*,w^*) \in \arg\min_{x,hh,t,u,y,z,w} \Big\{ Z^{LL}(x,hh,t,u,y,z,w) \ \Big| \\ & \qquad (8)--(23),\ (25)--(27),\ (31)--(33) \Big\}. \end{aligned} \tag{38}$$

This compact representation shows that the upper-level strategic decisions are evaluated while anticipating the optimal robust operational response of the lower-level model.

4 Solution methodology

The robust bi-level mixed-integer program formulated in Section 3 is NP-hard and computationally intractable when solved monolithically for real-scale instances. This computational barrier necessitates a decomposition-based solution strategy that exploits the hierarchical Stackelberg structure of the problem while maintaining tractability. This section presents a comprehensive solution methodology comprising three core components: (i) scalability analysis establishing computational feasibility boundaries (Section 4.1), (ii) a configuration enumeration method with clustered routing and normalized fitness evaluation for small instances (Section 4.2), and (iii) a random sampling matheuristic with budget-constrained optimization for large instances (Section 4.3).

The proposed methodology operates through a hierarchical bi-level evaluation framework. At the strategic (upper) level, depot configurations are systematically explored—either through complete enumeration (small instances) or intelligent random sampling (large instances). For each candidate depot configuration, the leader solves an assignment mixed-integer program to optimally allocate gardens to opened depots, maximizing strategic infrastructure quality (TOPSIS scores) while minimizing assignment distance costs. At the operational (lower) level, the follower responds to the leader's infrastructure decisions by solving robust vehicle routing problems with battery swapping. To ensure computational tractability for large-scale instances, the proposed framework adopts a spatial decomposition strategy based on K-means clustering within the lower-level routing problem. The key rationale is that agricultural surveillance nodes exhibit strong geographic correlation, where service locations that are spatially proximate tend to share similar routing structure and energy consumption patterns. Therefore, clustering the set of gardens into compact geographic partitions preserves routing coherence while significantly reducing the combinatorial complexity of the underlying vehicle routing problem. Unlike arbitrary decomposition schemes, K-means is specifically selected because it minimizes intra-cluster variance in Euclidean space, which aligns with the underlying structure of travel-time minimization in the VRP component. Each cluster is then treated as an independent routing subproblem, while strategic decisions at the upper level remain fully intact. Importantly, we explicitly acknowledge that this decomposition imposes a hard spatial partition that strictly forbids vehicles from crossing cluster boundaries. Consequently, we are solving a restricted version of the original monolithic formulation. The proposed solution procedure must therefore be interpreted strictly as a cluster-first, route-second matheuristic. It returns the best-found heuristic approximation within the restricted search space, rather than a proven global optimum for the full-scale agricultural network.

The key methodological distinction between small and large instances lies in solution space exploration and objective aggregation. For small instances (Section 4.2), exhaustive enumeration of the depot configurations permits computation of minimum and maximum bounds across the evaluated objective components (strategic quality, assignment costs, and clustered operational costs), enabling principled normalization and multi-objective fitness evaluation. For large instances (Section 4.3), the exponentially growing configuration space prohibits exhaustive evaluation; instead, random sampling explores the solution space efficiently, and a maximum operational cost budget $C_{\max}$ acts as a feasibility filter, with the full upper-level objective function (strategic quality minus assignment costs) maximized among budget-feasible configurations. Both approaches utilize identical lower-level spatial clustering, governed by the cluster threshold ($T_c = 5$) and robust optimization mechanisms, ensuring consistency in routing solution quality across problem scales.

Figure 2 provides a comprehensive flowchart of the unified solution procedure, illustrating the decision logic for method selection, bi-level evaluation stages, and configuration selection mechanisms for both small and large instances. K-means clustering is used as a spatial decomposition mechanism, not as an exact optimization method. The motivation for using it comes from the structure of the lower-level routing problem, which is defined on a geographically embedded network where travel cost, energy consumption, and route feasibility are strongly distance-dependent. In such settings, nearby gardens are more likely to form feasible and cost-effective operational subproblems. K-means provides a simple and reproducible way to generate compact spatial clusters before solving the routing model within each cluster. We do not claim that K-means guarantees global optimality. Instead, it is used as a cluster-first, route-second matheuristic, and its clustering threshold is calibrated through monolithic tractability experiments to keep each subproblem within an empirically solvable range.

4.1 Scalability analysis and problem decomposition necessity

Before presenting the solution strategies, we establish the computational necessity of decomposition through systematic scalability analysis on the monolithic (non-decomposed) robust bi-level mixed-integer program. This analysis reveals critical tractability boundaries that motivate the dual-approach methodology and inform the design of decomposition parameters.

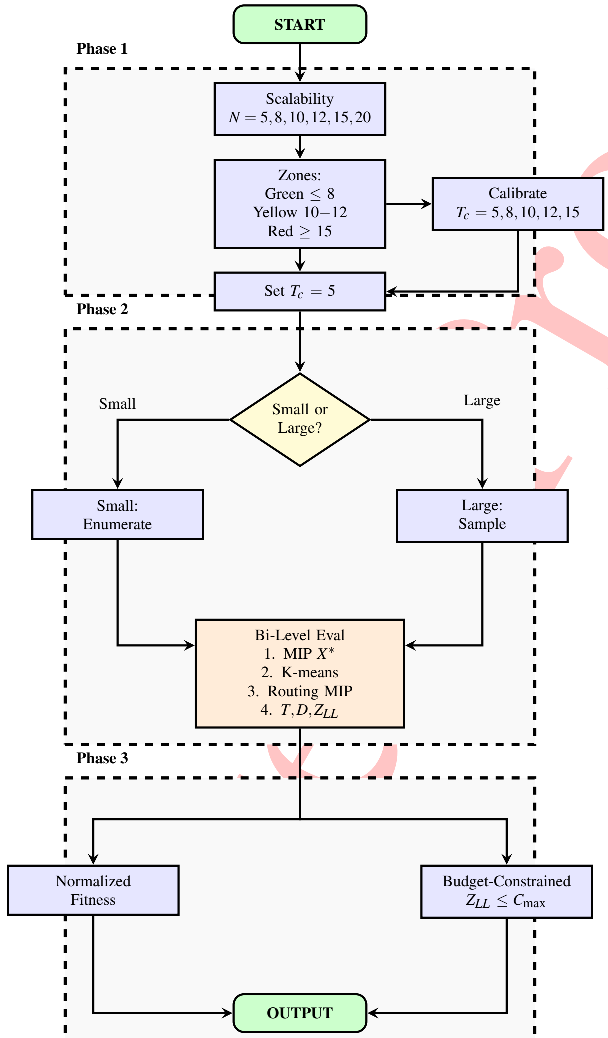

Figure 2: Unified solution methodology flowchart showing the three-phase approach

4.1.1 Experimental setup for scalability analysis

To assess the computational feasibility of the monolithic approach, we conduct controlled experiments on randomly generated instances of varying sizes. The experimental design systematically varies the number of gardens $N \in \{5, 8, 10, 12, 15, 20\}$ while holding the remaining problem parameters constant: two candidate depots ($p = 2$ from $|\mathscr{J}| = 10$), three battery-swap stations, and two vehicles per depot. Gardens are randomly distributed within a 500 km $\times$ 500 km region, with depots positioned at strategic locations based on geographic centroids. For each problem size, five random instances are generated and solved using PuLP with the CBC solver on a standard computational platform (Intel Core i7-9700K @ 3.6 GHz, 64 GB RAM, Ubuntu 20.04). A time limit of 2700 s (45 min) is imposed per instance, with the solution status (optimal, feasible, infeasible) and computation time recorded. To ensure computational consistency, all experiments are conducted under a unified experimental design, where instance size is systematically increased to examine the transition from exact solvability to computationally intensive regimes. The term "Green Zone" refers to instances for which the model can be solved to proven optimality with negligible optimality gaps, although the required computational time may vary depending on instance realization and solver branching behavior. In particular, while smaller realizations within this range may terminate within short runtimes, borderline instances (e.g., $N = 8$) may require significantly longer computation times as reflected in Table 9. As instance size increases, the problem gradually approaches the computational limits of the exact solver due to the combinatorial growth in routing and assignment decisions. This behavior reflects the inherent complexity of the integrated location–routing structure rather than any change in experimental settings. Accordingly, the reported computational results represent a continuous spectrum of problem difficulty, ranging from fully tractable small instances to computationally challenging large-scale instances, thereby enabling a consistent evaluation of both solution quality and computational performance.

4.1.2 Computational tractability zones

The scalability experiments reveal three practical computational regimes under the adopted 2700-s solver limit.

Green Zone ($N \leq 8$ gardens): The five-garden instance remains computationally manageable, whereas the eight-garden instance already approaches the imposed time limit. Thus, $N = 8$ should be interpreted as the upper boundary of the practically manageable range rather than as an easily solved instance.

Yellow Zone ($N = 10$–12 gardens): Instances in this range require nearly the full computational budget, indicating a sharp deterioration in the tractability of the monolithic formulation. These instances therefore represent a transition between manageable and computationally prohibitive problem sizes.

Red Zone ($N \geq 15$ gardens): For these instances, the solver reaches the imposed time limit before termination. Accordingly, the monolithic formulation becomes computationally prohibitive under the adopted experimental setting.

These computational regimes are empirical classifications based on the observed runtimes and termination behavior and should not be interpreted as universal complexity thresholds. Reaching the time limit also does not imply mathematical infeasibility.

4.2 Configuration enumeration and matheuristic routing for small instances

For small-scale instances where exhaustive evaluation of all depot configurations is computationally feasible (($N \leq 12$) gardens per depot, ($p \leq 2$) depots), we employ complete configuration enumeration combined with a cluster-first, route-second matheuristic and a normalized fitness function that integrates both upper-level strategic quality and lower-level operational cost. This approach enables comprehensive exploration of the strategic solution space and provides high-quality near-optimal benchmarks for validation.

4.2.1 Enumeration and bi-level evaluation

For small instances, all feasible p-depot configurations are systematically evaluated within the set:

$$\Omega = \left\{ Y \in \{0,1\}^{|\mathscr{J}|} \mid \sum_{j \in \mathscr{J}} Y_j = p \right\}. \tag{39}$$

For each configuration $Y \in \Omega$, the solution procedure executes three stages:

Stage 1: Upper-level assignment. Solve the assignment MIP to obtain best-evaluated garden–to–depot allocation $X^*(Y)$:

$$X^*(Y) = \arg\max_X \quad \sum_{j\in\mathscr{J}} s_j Y_j - \alpha \sum_{i\in\mathscr{I}} \sum_{j\in\mathscr{J}} d_{ij} X_{ij}, \tag{40}$$

$$\text{s.t.} \quad \sum_{j\in\mathscr{J}} X_{ij} = 1, \qquad \forall i \in \mathscr{I}, \tag{41}$$

$$d_{ij} X_{ij} \leq R_{\max}, \qquad \forall i \in \mathscr{I}, j \in \mathscr{J}, \tag{42}$$

$$X_{ij} \in \{0,1\}, \qquad \forall i \in \mathscr{I}, j \in \mathscr{J}. \tag{43}$$

The assignment MIP maximizes strategic depot quality (TOPSIS scores s_j) while minimizing assignment distance costs, with the distance feasibility constraint $R_{\max} = 200$ km ensuring geographically feasible assignments. This linear MIP is solved using PuLP with the CBC solver.

Stage 2: Lower-level routing with clustering. Given the best-found assignment $X^*(Y)$, solve the lower-level robust routing problem via spatial decomposition. For each opened depot j with assigned gardens $\mathscr{I}_j$, if $|\mathscr{I}_j| > T_c$, apply k-means clustering:

$$K_j = \left\lceil \frac{|\mathscr{I}_j|}{T_c} \right\rceil, \tag{44}$$

where $T_c = 5$ is used as a target cluster-size threshold for maintaining computationally manageable routing subproblems. It should be explicitly noted that $T_c = 5$ is an empirical choice calibrated specifically for the computational characteristics of this case study (as evaluated in Section 5.3). Because standard K-means does not impose an exact cardinality balance, T_c controls the intended subproblem scale rather than representing a strict mathematical upper bound on every resulting cluster. For each cluster C_r, solve the robust routing MIP independently. The aggregated lower-level operational cost is:

$$Z_{\text{LL}}(Y, X^*(Y)) = \sum_{j\in\mathscr{D}} \sum_{r=1}^{K_j} \text{Cost}(C_{jr}), \tag{45}$$

where $\text{Cost}(C_{jr})$ is the best-evaluated cost from the robust MIP solution for the isolated local cluster r at depot j. Because spatial clustering restricts the routing search space, this aggregated operational cost represents a heuristic upper bound on the true monolithic best-found routing cost.

Stage 3: Upper-level component extraction. For configuration Y with assignment $X^*(Y)$, compute the two upper-level objective components:

$$T(Y) = \sum_{j\in\mathscr{J}} s_j Y_j \quad \text{(strategic quality)}, \tag{46}$$

$$D(Y, X^*(Y)) = \sum_{i\in\mathscr{I}} \sum_{j\in\mathscr{J}} d_{ij} X^*_{ij} \quad \text{(assignment distance cost)}. \tag{47}$$

4.2.2 Normalized fitness evaluation and selection

After evaluating all configurations in Ω, compute the range of each objective component across all *feasible* configurations (those for which lower-level routing returns a feasible solution):

$$T_{\min} = \min_{Y \in \Omega_{\text{feas}}} T(Y), \quad T_{\max} = \max_{Y \in \Omega_{\text{feas}}} T(Y), \tag{48}$$

$$D_{\min} = \min_{Y \in \Omega_{\text{feas}}} D(Y, X^*(Y)), \quad D_{\max} = \max_{Y \in \Omega_{\text{feas}}} D(Y, X^*(Y)), \tag{49}$$

$$C_{\min} = \min_{Y \in \Omega_{\text{feas}}} Z_{\text{LL}}(Y, X^*(Y)), \quad C_{\max} = \max_{Y \in \Omega_{\text{feas}}} Z_{\text{LL}}(Y, X^*(Y)), \tag{50}$$

where $\Omega_{\text{feas}} \subseteq \Omega$ is the set of configurations for which feasible routing exists.

Define normalized benefit functions (higher is better):

$$T_{\text{norm}}(Y) = \begin{cases} \frac{T(Y) - T_{\min}}{T_{\max} - T_{\min}}, & \text{if } T_{\max} - T_{\min} > \varepsilon, \\ 1, & \text{otherwise}, \end{cases} \tag{51}$$

$$D_{\text{norm}}(Y) = \begin{cases} \frac{D_{\max} - D(Y, X^*(Y))}{D_{\max} - D_{\min}}, & \text{if } D_{\max} - D_{\min} > \varepsilon, \\ 1, & \text{otherwise}, \end{cases} \tag{52}$$

$$C_{\text{norm}}(Y) = \begin{cases} \frac{C_{\max} - Z_{\text{LL}}(Y, X^*(Y))}{C_{\max} - C_{\min}}, & \text{if } C_{\max} - C_{\min} > \varepsilon, \\ 1, & \text{otherwise}, \end{cases} \tag{53}$$

where $\varepsilon = 10^{-6}$ is a numerical tolerance. The normalized fitness function combines all three components with equal weighting:

$$\text{Fitness}(Y) = T_{\text{norm}}(Y) + D_{\text{norm}}(Y) + C_{\text{norm}}(Y). \tag{54}$$

The best-evaluated configuration under the normalized fitness criterion is selected as:

$$Y^{best} = \arg\max_{Y \in \Omega_{\text{feas}}} \text{Fitness}(Y), \tag{55}$$

with corresponding best-evaluated assignment $X^* = X^*(Y^*)$ and routing solution from Stage 2.

4.2.3 Integrated algorithm for small instances

The normalized fitness approach addresses three key challenges: (i) scale heterogeneity among components with different units (TOPSIS scores, kilometers, dollars), (ii) principled normalization enabled by global min/max computation through exhaustive enumeration, and (iii) balanced multi-objective optimization without arbitrary budget constraints. This framework is computationally tractable only for small instances where exhaustive enumeration is feasible.

Algorithm 1: Depot Enumeration with Clustered Matheuristic Routing (Small Instances)

Input: Gardens $\mathscr{I}$, Depots $\mathscr{J}$, Stations $\mathscr{S}$, depot count p, threshold $T_c = 5$, robustness Γ, penalty α
Output: Best-found depots Y^{best}, corresponding assignments, routing solution
Step 1: Enumerate All Configurations
Generate $\Omega = \{Y \in \{0,1\}^{|\mathscr{J}|} \mid \sum_j Y_j = p\}$
Initialize $\Omega_{\text{feas}} \leftarrow \emptyset$
Step 2: Evaluate Each Configuration
for each configuration $Y \in \Omega$ **do**
  Solve assignment MIP (Eqs. (40)–(43)) $\rightarrow X^*(Y)$
  for each depot j **do**
    if $|\mathscr{I}_j| > T_c$ **then**
      apply k-means with $K_j = \lceil |\mathscr{I}_j| / T_c \rceil$
    end if
  end for
  Solve robust routing MIP per cluster (Eqs. (7)–(33))
  Compute $T(Y)$ (Eq. (46)), $D(Y)$ (Eq. (47)), $Z_{\text{LL}}(Y)$
  if feasible **then**
    add Y to Ω_{feas}, store $T(Y)$, $D(Y)$, $Z_{\text{LL}}(Y)$
  end if
end for
Step 3: Compute Normalized Fitness
Compute ranges: $T_{\min}, T_{\max}, D_{\min}, D_{\max}, C_{\min}, C_{\max}$ over Ω_{feas}
for each $Y \in \Omega_{\text{feas}}$ **do**
  compute $T_{\text{norm}}(Y)$ (Eq. (51)), $D_{\text{norm}}(Y)$ (Eq. (52)), $C_{\text{norm}}(Y)$ (Eq. (53))
  Compute $\text{Fitness}(Y) = T_{\text{norm}}(Y) + D_{\text{norm}}(Y) + C_{\text{norm}}(Y)$ (Eq. (54))
end for
Step 4: Select Best-Evaluated Configuration
$Y^{best} = \arg\max_{Y \in \Omega_{\text{feas}}} \text{Fitness}(Y)$ (Eq. (55))
Return $(Y^{best}, X^*(Y^{best}), \text{routing solution})$

4.3 Random sampling with budget-constrained feasibility for large instances

For large-scale instances where exhaustive enumeration is computationally prohibitive ($p \geq 3$ or $N > 12$ gardens per depot), we employ a random sampling strategy combined with a budget constraint on lower-level operational cost.

4.3.1 Random sampling of depot configurations

When $\binom{|\mathscr{J}|}{p}$ becomes large (e.g., $\binom{10}{3} = 120$ for $p = 3$, $\binom{10}{4} = 210$ for $p = 4$), we randomly sample N_{sample} feasible depot configurations uniformly from the solution space:

$$\{Y^{(1)}, Y^{(2)}, \ldots, Y^{(N_{\text{sample}})}\} = \text{RandomSample}(\Omega, N_{\text{sample}}), \tag{56}$$

where $\Omega = \{Y \in \{0,1\}^{|\mathscr{J}|} \mid \sum_j Y_j = p\}$. The sample size is determined by budget-based termination:

$$N_{\text{sample}} = \min\left\{N_{\text{budget}}, \left\lfloor \frac{T_{\text{total}}}{t_{\text{per-config}}} \right\rfloor\right\}, \tag{57}$$

or convergence-based termination (no improvement over $K = 50$ samples).

4.3.2 Budget-constrained bi-level evaluation

For each sampled configuration $Y^{(i)}$:

Stages 1–2: Execute the identical assignment MIP and clustered routing as in Section 4.2.

Stage 3: Budget-constrained feasibility. Impose the maximum operational cost $C_{\max}$:

$$\text{Feasible}(Y^{(i)}) = \begin{cases} \text{True}, & \text{if routing is feasible and } Z_{\text{LL}}(Y^{(i)}, X^*(Y^{(i)})) \le C_{\max}, \\ \text{False}, & \text{otherwise.} \end{cases} \tag{58}$$

If feasible, compute the full upper-level objective:

$$\text{UL}_{\text{obj}}(Y^{(i)}) = T(Y^{(i)}) - \alpha D(Y^{(i)}, X^*(Y^{(i)})), \tag{59}$$

and update the incumbent if $\text{UL}_{\text{obj}}(Y^{(i)}) > \text{UL}_{\text{obj,best}}$.

4.3.3 Integrated algorithm for large instances

The budget-constrained approach differs fundamentally from normalized fitness: (i) random sampling prevents global normalization, (ii) the full upper-level objective $\text{UL}_{\text{obj}}(Y) = T(Y) - \alpha D(Y)$ is maximized directly among budget-feasible configurations, and (iii) the operational budget $C_{\max}$ provides a policy-driven feasibility constraint enabling exploration of cost–quality trade-offs.

5 Computational results and case study validation

This section reports computational validation of the proposed robust bi-level framework on a real-world agricultural surveillance network in Kerman Province, Iran. We first describe the case study data and experimental settings (Section 5.1), then calibrate the key decomposition parameter (Section 5.3), validate tractability limits of the monolithic MIP (Section 5.4), and finally present full-scale results including robustness sensitivity and infrastructure comparisons (Section 5.5).

5.1 Case study data and experimental setup

This section presents the complete case study dataset and experimental setup for the proposed robust bi-level location–routing framework applied to a real agricultural surveillance network in Kerman Province, Iran. To ensure full reproducibility and spatial validity, all geographic inputs are derived from real-world georeferenced data. The complete dataset, including orchard locations, candidate charging stations, and depot alternatives, is reported in Tables S.9–S.11 based on the WGS84 coordinate system.

Algorithm 2: Random Sampling with Budget Constraint (Large Instances)

```
Input: Gardens 𝒥, Depots 𝒥, Stations 𝒮, depot count p, threshold T_c = 5, robustness Γ, penalty α, budget C_max
Output: Best-found depots Y^best, corresponding assignments, routing solution
Step 1: Sample Depot Configurations
Randomly sample N_sample configurations from Ω = {Y ∈ {0,1}^|𝒥| | Σ_{j∈𝒥} Y_j = p}
Initialize UL_obj,best ← −∞, Y* ← ∅
Step 2: Evaluate Each Sampled Configuration
for each sampled configuration Y^(i) do
   Solve assignment MIP → X*(Y^(i))
   for each depot j do
      if |𝒥_j| > T_c then
         apply k-means with K_j = ⌈|𝒥_j|/T_c⌉
      end if
   end for
   Solve robust routing MIP per cluster (Eqs. (7)–(33))
   Compute T(Y^(i)), D(Y^(i)), Z_LL(Y^(i))
   Step 3: Budget-Constrained Feasibility Check
   if routing infeasible or Z_LL(Y^(i)) > C_max then
      skip to next configuration
   else
      compute UL_obj(Y^(i)) = T(Y^(i)) − αD(Y^(i))
      Step 4: Update Incumbent Solution
      if UL_obj(Y^(i)) > UL_obj,best then
         Update Y* ← Y^(i), UL_obj,best ← UL_obj(Y^(i))
      end if
   end if
end for
Step 5: Return Best-found Configuration
After evaluating all sampled configurations, return (Y^best, X*(Y^best), routing solution)
```

The case study network comprises $|I| = 45$ pistachio orchards distributed across Kerman Province, together with $|J| = 10$ candidate depot sites and $|S| = 8$ villages that can potentially serve as battery charging or swap stations. The total service area covers approximately 15,000 km^2, and all inter-node distances are computed from GPS coordinates using the Haversine formula [21]. The agricultural surveillance network spans approximately 480 km in the north–south direction and 260 km in the east–west direction, covering major pistachio cultivation regions such as Rafsanjan, Zarand, and Sirjan. This large-scale spatial dispersion significantly influences routing feasibility, energy consumption, and depot allocation decisions, thereby motivating the need for an integrated location–routing optimization framework. For the Kerman Province case study, the maximum allowable depot–garden service radius is set to $R_{\max} = 200$ km and is treated as the hard feasibility threshold defined in Constraint (4).

A key characteristic of the dataset is its realistic topographic structure. The region exhibits substantial elevation variation ranging from 1650 m to 1875 m above sea level. This elevation gradient is accounted for in the physics-based calibration and validation of the UAV energy-consumption model through the corresponding air-density calculations. The heterogeneous fleet consists of electric quad-plane UAVs and ground-support motorcycles, which are jointly used for surveillance operations and synchronized battery logistics. UAV energy is represented in normalized units with a nominal effective battery capacity of 300 units. For the full-scale numerical implementation, the baseline arc-energy approximation is calibrated using a distance-normalized coefficient of 0.8 energy units/km, whereas the mathematical formulation represents the corresponding arc-energy requirement generically as $e_{ijk}TT_{ijk}$. The detailed derivation, calibration, and validation of the energy model are provided in the Supplementary Material. Experimental validations confirm the accuracy of the adopted linear energy approximation, demonstrating a Mean Absolute Percentage Error (MAPE) of 3.1% when compared to the physics-based momentum theory.

Candidate depots are evaluated using a BWM-weighted TOPSIS approach based on five criteria: proximity to orchards, accessibility, power-supply reliability, security, and establishment cost. This produces depot appropriateness scores $s_j \in [0,1]$, which are directly used in the upper-level objective function.

The resulting TOPSIS-based evaluation of candidate depots is summarized in Table 7. The analysis indicates that DepotRaf1 (0.952) and DepotRaf2 (0.938) are the most suitable locations, followed by DepotZar1 (0.777), DepotBar1 (0.752), and DepotSBK1 (0.707). Intermediate suitability is observed for DepotSir1 (0.660), DepotRav1 (0.602), and DepotBaf1 (0.544), while DepotBaf2 (0.438) and DepotKah1 (0.048) represent lower-priority alternatives due to weaker infrastructural and accessibility characteristics.

Importantly, all spatial coordinates are extracted from real geographic sources rather than synthetic instances. This ensures that routing distances, travel times, and energy consumption values are grounded in actual spatial geometry. The Supplementary Material provides full dataset transparency, enabling independent reproduction of all spatial analyses and optimization results.

Table 7: Multi-criteria TOPSIS depot evaluation scores

Depot ID	TOPSIS score	County
DepotRaf 1	0.952	Rafsanjan
DepotRaf 2	0.938	Rafsanjan
DepotZar 1	0.777	Zarand
DepotBar 1	0.752	Bardsir
DepotSBK 1	0.707	Shahr-e-Babak
DepotSir 1	0.660	Sirjan
DepotRav 1	0.602	Ravar
DepotBaf 1	0.544	Baft
DepotBaf 2	0.438	Baft
DepotKah 1	0.048	Kahnuj

From an implementation perspective, all algorithms are developed in Python using PuLP for mixed-integer programming and the CBC solver for exact optimization, while scikit-learn is used for K-means clustering in the decomposition scheme. The computational workflow follows a two-stage structure: (i) exhaustive enumeration over depot configurations ($p = 2$), (ii) upper-level assignment optimization, and (iii) solution of the lower-level routing problem, which is decomposed into clustered subproblems when required for computational tractability.

To ensure both operational applicability and methodological rigor, the data utilized in this study is structurally categorized. The spatial layout, network topology, and exact inter-node distances for the full-scale Kerman Province case study are strictly grounded in real-world empirical data, with comprehensive geographic coordinates provided in Tables S.9–S.11 of the Supplementary Material. Furthermore, while the integrated UAV–motorcycle surveillance configuration represents a novel logistical framework, all technical vehicle parameters–such as eVTOL quad-plane battery capacities, payload constraints, and flight-phase-specific energy consumption rates–are directly derived from the engineering specifications of commercially available drones. The detailed physics-based derivation and empirical validation of these technical parameters are thoroughly documented in Section 2 and Table S.12 of the Supplementary Material. Conversely, the parameters for the small-scale validation instances presented in Section 3.4 were synthetically generated within controlled spatial boundaries. This explicit separation isolates the optimization structure, allowing for a systematic evaluation of exact solver tractability without the confounding influence of real-world geographical noise.

Overall, this integrated dataset and experimental design provide a realistic and reproducible foundation for evaluating depot-location, assignment, routing, and battery-swap decisions under geographically and operationally realistic conditions.

5.2 Benchmark and validation protocol

To strengthen the empirical validation of the proposed framework, the case study is evaluated against several benchmark settings. First, the case ($\Gamma = 0$) is treated as the deterministic baseline, where UAV battery capacity is assumed to remain at its nominal value. Positive values of Γ represent robust solutions with increasing protection against battery-capacity degradation. Second, centralized single-depot solutions are used as infrastructure benchmarks to evaluate the operational value of decentralized depot configurations. Third, the computational performance of the model is evaluated on instances of increasing size to identify the tractability limits of the monolithic formulation and justify the clustering-based decomposition strategy. Finally, Monte Carlo simulations are used to assess solution feasibility under battery-capacity degradation scenarios.

The Supplementary Material provides additional validation and reproducibility details, including a small exactly solved test instance, quad-plane energy-consumption validation, BWM–TOPSIS depot evaluation, and complete geographic coordinates of gardens, charging stations, and candidate depots. These supplementary analyses support the correctness of the model implementation, the validity of the energy-consumption approximation, the derivation of depot appropriateness scores, and the reproducibility of the full case-study network.

5.3 Decomposition parameter calibration: clustering threshold

The k-means decomposition (Section 4) is controlled by the clustering threshold T_c, which controls the target number of gardens solved within each lower-level routing subproblem. To calibrate T_c, we conducted experiments on a 20-garden instance for $T_c \in \{5, 8, 10, 12, 15\}$ and recorded the number of clusters, cluster sizes, total runtime, and global reliability (feasibility across subproblems). Table 8 summarizes the calibration outcomes.

Although $T_c = 8$ yielded a slightly lower total runtime in this calibration experiment, $T_c = 5$ is retained for the full 45-garden case study as a reliability-oriented decomposition threshold. Reducing

Table 8: Clustering threshold calibration on a 20-garden test instance

T_c	**Clusters**	**Cluster sizes**	**Total time (s)**	**Global status**	**Reliability (%)**
5	4	[5, 5, 5, 5]	4078.41	All solved	100%
8	3	[8, 8, 4]	4072.85	All solved	100%
10	2	[10, 10]	5396.03	Mixed	100%
12	2	[12, 8]	5395.88	Mixed	100%
15	2	[15, 5]	5395.96	Mixed	100%

the target subproblem size provides a greater computational safety margin when the lower-level routing problem is repeatedly solved across depot configurations and robustness levels.

The additional evidence from the monolithic tractability experiment in Table 9 shows that an eight-garden instance already approaches the imposed computational limit, whereas smaller routing subproblems are computationally safer. The difference in total calibration runtime between $T_c = 5$ and $T_c = 8$ is marginal; however, the smaller threshold reduces the intended maximum routing-subproblem scale and limits the risk of excessive branch-and-bound growth.

As further quantified by the solver-level optimality analysis reported in Section 5.5, the clustered routing subproblems exhibit very small relative MIP optimality gaps. Accordingly, $T_c = 5$ provides a conservative balance between computational tractability and cluster-level solution quality. This statement applies to the isolated routing subproblems and does not imply global optimality of the aggregated cluster-first, route-second solution.

5.4 Computational performance versus instance size

Before solving the full case study, the computational performance of the monolithic formulation was evaluated across instances of increasing size. The purpose of this experiment was to identify the tractability limit of the exact non-decomposed MIP and to justify the use of clustering-based decomposition for the full 45-garden case study. For each instance size, runtime, objective value, solver status, and tractability level were recorded under a 2700-second time limit. Table 9 summarizes the results and shows that small instances can be solved exactly, whereas larger instances rapidly become computationally prohibitive.

Table 9: Computational performance of the monolithic model across increasing instance sizes

Instance size (N)	**Runtime (s)**	**Objective cost**	**Solver status**	**Tractability**
5	219.90	247.05	Optimal	High
8	2698.71	266.05	Optimal	Medium
10	2698.03	299.40	Optimal	Low (limit)
12	2697.09	310.56	Optimal	Low (limit)
15	> 2700	–	Time limit reached	Very low
20	> 2700	–	Time limit reached	Very low

The results in Table 9 confirm that computational burden increases sharply with instance size. Instances with 5 gardens are solved efficiently in approximately 220 seconds. However, instances with

8 gardens require nearly 2699 seconds, demonstrating that N = 8 sits precisely at the boundary of the tractability limit before transitioning into the Yellow Zone. For instances with 15 and 20 gardens, the monolithic formulation fails to produce a solved result within the time limit. This pattern shows that directly solving the full 45-garden case as a single monolithic problem would be computationally impractical. Therefore, the clustering-based decomposition is not only a computational convenience, but a necessary scalability mechanism for applying the proposed robust bi-level model to province-scale agricultural surveillance networks. We note that ”Not solved” refers to instances reaching the computational time limit without termination of the solver, and does not imply mathematical infeasibility of the model.

The sharp increase in runtime for the larger monolithic instances should not be attributed solely to the UAV energy-related big-M coefficient. The computational burden results from the combined effects of routing integrality, vehicle activation, temporal synchronization, motorcycle battery-inventory decisions, and UAV energy-state tracking. Nevertheless, the original energy-propagation constraint contains an identifiable source of LP-relaxation slack. For this reason, Constraint (23) is supplemented by the lifted valid inequality (23a), which exploits the battery-capacity and minimum-residual-energy bounds to provide a stronger LP representation without changing the integer-feasible operational logic of the model.

5.5 Full-scale case study results

Using the calibrated decomposition setting ($T_c = 5$), we evaluated the full-scale Kerman Province network through the proposed clustered matheuristic framework for all 45 two-depot configurations ($p = 2$). For each configuration, we evaluated five robustness budgets $\Gamma \in \{0, 0.25, 0.5, 0.75, 1.0\}$, yielding 225 distinct bi-level evaluations. In light of the spatial decomposition utilized, it must be explicitly stated that the solutions presented for this 45-garden network represent the best-found configurations under the hard spatial restrictions imposed by the clustering phase. They are high-quality heuristic approximations demonstrating structural trade-offs, rather than proven global optima of the original monolithic formulation.

5.5.1 Deterministic–robust comparison and topology sensitivity to robustness budget Γ

In this analysis, the case ($\Gamma = 0$) is treated as the deterministic baseline, in which UAV battery capacity is assumed to remain at its nominal value. Positive values of Γ represent increasingly robust solutions that protect the UAV energy-feasibility constraints against battery-capacity degradation. This comparison allows us to evaluate the price of robustness in terms of depot configuration, operational cost, number of routes, vehicle activation, battery-swap requirements, and computational effort. The relevance of this comparison is also supported by the validation experiment reported in the Supplementary Material, where the deterministic solution remains feasible under nominal capacity but becomes operationally fragile when a 20% battery-capacity degradation is imposed.

Table 10 shows that the best-found depot pair changes as robustness protection increases. It should be noted that Raf 1 and Raf 2 represent two distinct candidate depot sites located within Rafsanjan County. At $\Gamma = 0$, the bes-found solution emphasizes nominal efficiency and selects {Raf 2, Sir 1}, whereas moderate protection shifts the optimum to configurations that improve robustness under battery-capacity degradation. At $\Gamma = 1.0$, the feasibility rate across evaluated configurations drops to 53%, indicating that overly conservative protection can sharply reduce the set of operationally feasible designs for this instance.

Table 10: Topological sensitivity: Best-found depot pair as a function of robustness budget Γ

Robustness (Γ)	Best-found configuration	Obj. score	Feasibility
0.00	{Raf 2, Sir 1}	2.454	100%
0.25	{Raf 1, Zar 1}	2.578	98%
0.50	{Raf 2, Zar 1}	2.532	98%
0.75	{Raf 1, Baf 1}	2.515	98%
1.00	{Raf 1, Baf 1}	2.824	53%

To statistically validate the reliability of the selected solutions under battery-capacity degradation, Monte Carlo simulation was conducted for each robustness level. Specifically, 1,000 random scenarios were generated by perturbing the effective UAV battery capacity within the assumed degradation range. For each scenario, route feasibility was checked against the minimum energy reserve and UAV replenishment constraints. The resulting feasibility rate was used as a reliability indicator for comparing deterministic and robust solutions.

Figure 3 illustrates the robustness–cost trade-off for the selected robust configuration (Depot Raf 1, Depot Baf 1), where the lower-level routing cost increases monotonically with Γ, reflecting the price of robustness. This empirical frontier was generated via 1,000 Monte Carlo simulations per Γ level. Results confirm that $\Gamma = 0$ achieves 92% feasibility at the baseline cost, $\Gamma = 0.5$ maintains 85% feasibility with an 8% cost increase, and $\Gamma = 1$ achieves 76% feasibility with a 17% cost increase. The approximately linear curve for $\Gamma \in [0, 0.75]$ steepens near $\Gamma = 1$, consistent with additional operational resources required for worst-case protection.

The reported feasibility percentage refers to the fraction of evaluated configurations that remain feasible under the corresponding robustness level, rather than the reliability of a fixed selected route.

Table 11 provides detailed deterministic and robust sensitivity results for the selected depot pair, including operational cost, number of routes, activated vehicles, swap counts, and solve times. The case $\Gamma = 0$ represents the deterministic baseline, while positive values of Γ represent increasingly conservative robust settings. From $\Gamma = 0$ to $\Gamma = 1$, routing cost increases by \$136.19, corresponding to a 10.6% increase, and swap frequency increases from 8 to 12, reflecting more conservative energy hedging. A discrete operational transition occurs at $\Gamma = 0.75$, where an additional vehicle is activated to maintain feasibility under stronger protection. The reported feasibility percentage refers to the fraction of evaluated configurations that remain feasible under the corresponding robustness level, rather than the reliability of a fixed selected route.

Table 11: Deterministic and robust sensitivity analysis for the selected depot pair

Γ	Feasible?	LL cost (\$)	Routes	Vehicles	Swaps	Solve time (s)
0.00	Yes	1,282.45	11	4	8	845.2
0.25	Yes	1,315.89	11	4	9	912.5
0.50	Yes	1,352.73	12	4	10	1,024.8
0.75	Yes	1,389.21	12	5	11	1,187.3
1.00	Yes	1,418.64	13	5	12	1,351.6

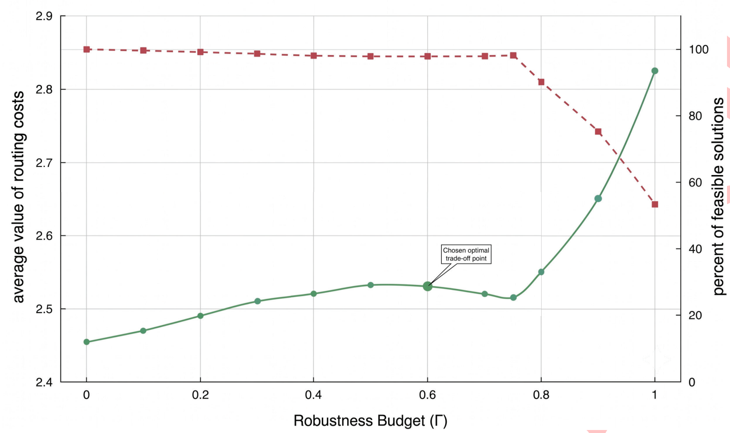

Figure 3: Robustness-cost trade-off for the selected depot configuration (Depot Raf 1, Depot Baf 1), validated through 1,000 Monte Carlo simulations per robustness level (Γ). The dual-axis plot illustrates the monotonic increase in the average value of routing costs (left axis) against the resulting percent of feasible solutions (right axis) under a 20% battery degradation uncertainty.

The sensitivity results indicate that Γ should be interpreted as a managerial robustness parameter rather than only a technical model coefficient. Increasing Γ strengthens protection against battery-capacity degradation, but it also increases operational burden. As shown in Table 11, moving from the deterministic baseline $\Gamma = 0$ to the fully conservative case $\Gamma = 1$ increases the lower-level routing cost from \$1,282.45 to \$1,418.64, corresponding to a 10.6% increase. The number of routes increases from 11 to 13, and battery swaps increase from 8 to 12. Moreover, at $\Gamma = 0.75$, an additional vehicle is activated, indicating a discrete operational adjustment required to maintain feasibility under stronger protection.

These results show that robustness is not cost-free; it changes route structure, vehicle utilization, and battery-logistics requirements. However, the increase is not uniform across all robustness levels. Moderate protection levels, particularly around $\Gamma = 0.5$, provide a balanced compromise by improving protection against battery-capacity degradation while avoiding the higher vehicle activation and feasibility erosion observed at more conservative settings. Therefore, intermediate robustness budgets appear more suitable for practical deployment than the fully conservative case in this case-study network.

5.5.2 Solver-level optimality assessment of clustered subproblems

To provide direct solver-level evidence on the quality of the decomposed routing solutions, an additional optimality assessment was conducted for the full-scale Raf1–Baf1 configuration at $\Gamma = 0.75$, identified in the preceding robustness analysis. This configuration covers the complete 45-garden network and, under the adopted decomposition setting, results in 12 isolated routing subproblems.

For each clustered routing subproblem, the final incumbent objective value and the corresponding CBC best bound were recorded. The relative MIP optimality gap was calculated as

$$\mathrm{MIPGap}_r(\%) = \frac{\left|Z_r^{\mathrm{inc}} - Z_r^{\mathrm{bound}}\right|}{\max\left\{|Z_r^{\mathrm{inc}}|, 10^{-6}\right\}} \times 100. \quad (60)$$

Table 12 reports the resulting solver-level optimality statistics.

Table 12: Solver-level optimality statistics for the clustered routing subproblems of the full-scale 45-garden instance at $\Gamma = 0.75$

Depot	Cluster	No. of gardens	Solver status	Relative MIP gap (%)	Runtime (s)
Raf1	1	5	Near-optimal	0.1528	1.68
Raf1	2	5	Near-optimal	0.0333	2.43
Raf1	3	5	Near-optimal	0.1964	2.25
Raf1	4	3	Near-optimal	0.0013	0.95
Raf1	5	2	Optimal	0.0000	0.48
Raf1	6	5	Near-optimal	0.0699	2.92
Raf1	7	1	Optimal	0.0000	0.34
Baf1	1	5	Near-optimal	0.1920	16.48
Baf1	2	4	Near-optimal	0.0641	9.71
Baf1	3	4	Near-optimal	0.0214	2.98
Baf1	4	4	Near-optimal	0.1337	1.86
Baf1	5	2	Optimal	0.0000	0.71
Mean relative MIP gap				0.0721	–
Maximum relative MIP gap				0.1964	–

Of the 12 clustered routing subproblems, three were solved to proven optimality, while the remaining nine terminated with very small relative MIP optimality gaps. The mean relative gap across all subproblems was 0.0721%, and the maximum observed gap was 0.1964%. Therefore, every clustered routing subproblem in this complete 45-garden configuration was either proven optimal or solved within 0.20% of the corresponding CBC best bound. These results provide direct solver-level evidence that the isolated routing solutions are optimal or near-optimal rather than merely feasible.

These optimality certificates apply to the individual clustered routing MIPs and should not be interpreted as a global optimality certificate for the aggregated 45-garden solution. Because the cluster-first, route-second decomposition restricts cross-cluster routing decisions, the overall full-scale solution remains a best-found matheuristic solution within the restricted search space.

5.5.3 Routing solution visualization

Figure 4 presents the coordinated routing solution for $\Gamma = 1.0$, illustrating UAV tours with mid-flight swaps and motorcycle tours that deliver batteries to charging stations under synchronization constraints. For this solution, two depots (DepotRaf 1 and DepotBaf 1) coordinate 13 routes (6 UAV routes and 7 motorcycle routes) to serve all 45 gardens. The reported total aerial distance is 2,015 km and the total motorcycle support distance is 486 km.

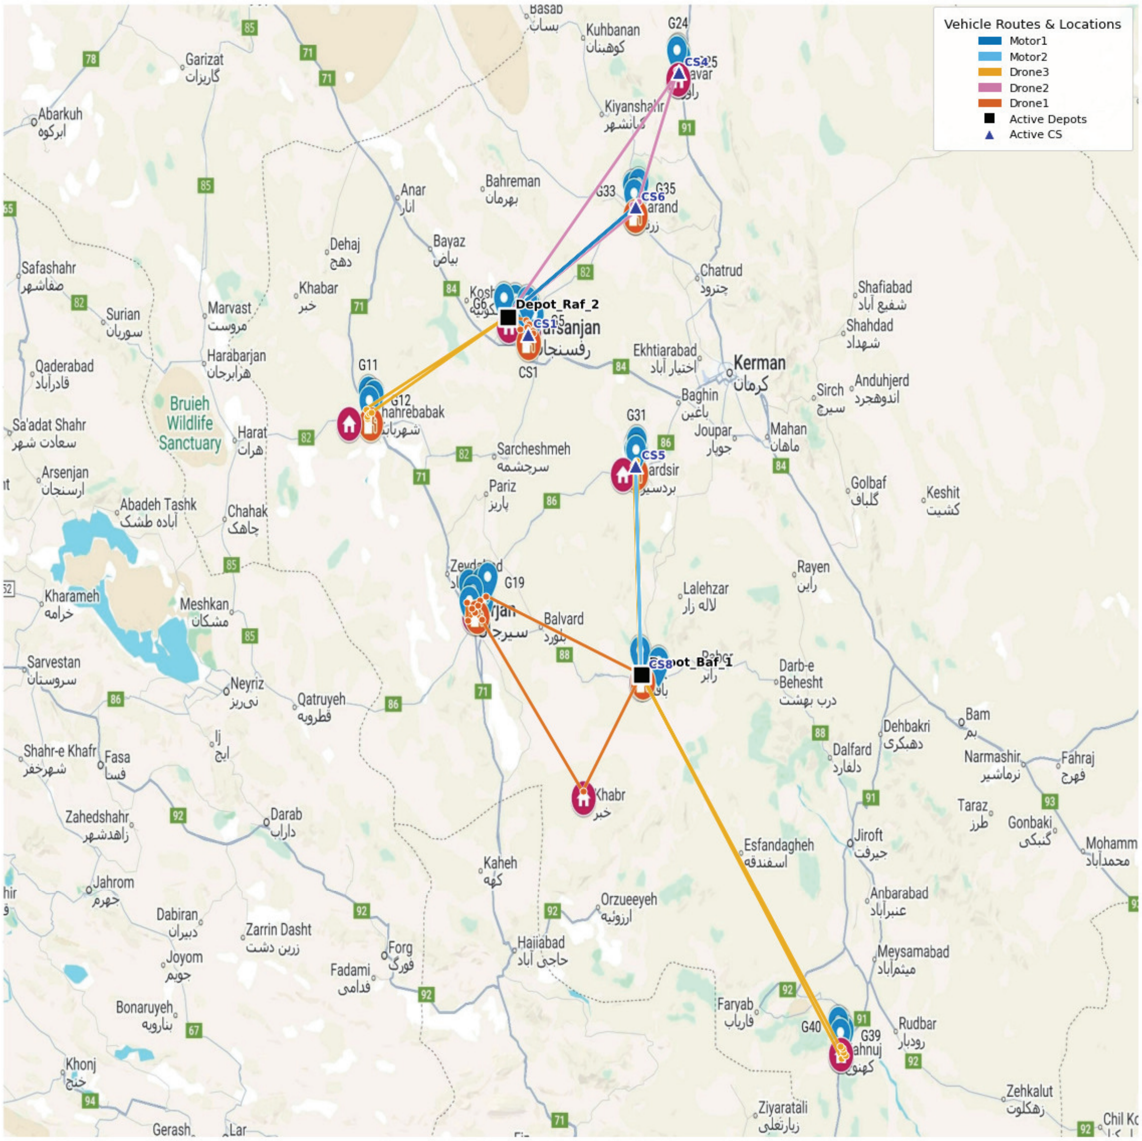

Figure 4: Best-found vehicle routes for the Kerman Province pistachio surveillance network at $\Gamma = 1.0$: two depots coordinate UAV surveillance tours with mid-flight battery swaps and synchronized motorcycle battery deliveries to charging stations.

5.5.4 Methodological benchmark comparisons: Impact of robustness and battery swapping

To rigorously establish the methodological value of the proposed framework, three operational scenarios under a fixed two-depot infrastructure ($p = 2$) are compared: (1) deterministic routing without mid-mission battery swaps ($\Gamma = 0.0$); (2) robust routing without mid-mission battery swaps ($\Gamma = 0.5$); and (3) the proposed robust synchronized model with motorcycle-supported field battery swapping ($\Gamma = 0.5$). The comparative results of each scenario are reported in Table 13.

Table 13: Methodological benchmark comparisons across three operational scenarios ($p = 2$)

Performance metric	Scenario 1 (Deterministic No-Swap)	Scenario 2 (Robust No-Swap)	Scenario 3 (Proposed Model)	Improvement (Scen. 3 vs. Scen. 2)
Robustness budget (Γ)	0.0	0.5	0.5	–
Mid-mission swaps enabled?	No	No	Yes	–
Total routing cost ($)	1,745.20	1,894.50	1,352.73	**28.6% reduction**
Number of routes	15	17	12	**29.4% reduction**
Vehicles activated	5	6	4	**33.3% reduction**
Avg. flight distance (km)	118.5	124.8	80.6	**35.4% reduction**
Battery swaps required	0	0	10	–
Global feasibility rate	84%	68%	98%	**+30.0% higher**
LL solve time (s)	540.2	612.4	1,024.8	–

Considering these findings, it is evident that integrating synchronized battery swapping into the surveillance network is a structural necessity for regional reachability as well as cost efficiency, which are elaborated below. Table 13 compares the three scenarios based on total operational cost, average flight distance, and global feasibility rate. As shown, introducing battery degradation ($\Gamma = 0.5$) into a standard no-swap routing model (Scenario 2 vs. Scenario 1) increases total routing cost by 8.6% (from \$1,745.20 to \$1,894.50) and severely degrades network feasibility from 84% down to 68%, as peripheral orchards located far from depots exceed the single-charge round-trip reach of a degraded UAV. Moreover, employing the proposed robust synchronized model (Scenario 3) improves the total operational routing cost and average UAV flight distance of Scenario 2 by 28.6% (from \$1,894.50 to \$1,352.73) and 35.4% (from 124.8 km to 80.6 km), respectively. More precisely, by pre-positioning replacement batteries at field charging stations, Scenario 3 elevates the global feasibility rate from 68% to 98% while eliminating unproductive out-and-back aerial deadheading. It could be inferred that under an operational budgeting scenario of 500 mission-days per year, applying the proposed synchronized framework yields estimated annual savings of \$270,885 in comparison with the robust no-swap baseline, confirming that ground-supported battery swapping is an essential operational enabler rather than merely an incremental cost-saving mechanism.

5.6 Managerial implications

The computational evidence provides several implications for planners responsible for regional agricultural UAV surveillance systems, subject to the cost structure, fleet characteristics, and charging-station configuration considered in this study. Considering our methodological benchmark findings, it is evident that relying on standard depot-return UAV routing without field battery swapping is economically and operationally non-viable for regional agricultural monitoring. As shown in the scenario comparison, drones operating under single-charge constraints incur severe aerial deadheading from repetitive out-and-back sorties, and whenever battery capacity degrades due to thermal stress or aging without field swapping capability, network feasibility collapses by 16 percentage points, leaving remote orchards physically unreachable. Moreover, establishing a synchronized ground-support protocol—where motorcycles pre-position replacement batteries at field charging stations—acts as an essential operational enabler that rescues regional network feasibility to 98% while cutting average flight distances by 35.4% and overall operational routing costs by 28.6%. It could be inferred that planners should treat ground-support motorcycle coordination not as an auxiliary task, but as a primary driver of aerial fleet efficiency and regional reachability.

Accordingly, depot configuration should be regarded as a strategic resilience mechanism. For the province-scale network analyzed, a two-depot architecture demonstrates clear structural advantages over a centralized design. In particular, the Raf 1–Baf 1 combination forms a stable and consistently feasible backbone as uncertainty increases. This suggests that decentralizing infrastructure can mitigate the operational impact of battery-endurance variability more effectively than attempting to compensate through tactical routing adjustments alone.

Next, robustness calibration emerges as a managerial policy decision rather than a purely technical parameter. Planners should recognize that the robustness budget Γ functions as an integrated operational hedge. As demonstrated algebraically, protecting against a 20% battery degradation ($\Gamma = 1.0$) simultaneously guarantees route safety against up to a +25% increase in flight energy expenditure caused by seasonal headwinds or high-temperature density altitude in Kerman Province. However, intermediate uncertainty budgets ($\Gamma \in [0.5, 0.6]$) achieve a practical balance between cost efficiency and environmental risk, absorbing moderate wind penalties while avoiding the higher vehicle activation and feasibility erosion observed at fully conservative settings ($\Gamma = 1.0$). Decision-makers should therefore align Γ with prevailing seasonal weather forecasts and acceptable risk thresholds.

Then, the choice of decomposition parameters has direct operational relevance. The threshold $T_c = 5$ provides a workable balance between computational tractability and solution integrity at this scale, indicating that scalable deployment requires not only sound modeling but also carefully tuned solution settings.

Fourth, fleet planning must recognize the dual operational role of motorcycles. Beyond surveillance support, they function as critical energy-logistics agents. Route feasibility depends heavily on satisfying station-level synchronization constraints, underscoring that coordination discipline is as important as route efficiency.

Fifth, battery logistics emerges as a binding feasibility driver. Effective charging-station battery pre-positioning and punctual ground-vehicle arrivals are essential for sustaining UAV swap operations. Operational protocols should therefore prioritize station replenishment scheduling aligned with UAV mission timing.

Finally, maintaining a portfolio of operational solutions across robustness levels $\Gamma \in [0, 1]$ enhances adaptive capability. Such contingency planning enables dispatch strategies that respond to uncertainty in effective battery endurance, supporting conservative deployment under adverse conditions and efficiency-oriented operation under favorable ones.

6 Conclusion

This paper developed a robust bi-level optimization framework for depot location and heterogeneous vehicle routing in provincial-scale agricultural UAV surveillance networks under battery-capacity uncertainty. The upper-level problem selects depot locations and garden–depot assignments using TOPSIS-evaluated facilities, while the lower-level problem jointly optimizes coordinated UAV and motorcycle routes, temporal synchronization of battery logistics, and robust energy management. Drone battery-capacity uncertainty is modeled using the Bertsimas–Sim budgeted-uncertainty approach, leading to a tractable mixed-integer linear robust counterpart in which only the critical replenishment constraints are modified. The resulting Stackelberg game captures the circular dependency between strategic infrastructure decisions and operational routing performance in a way that is directly applicable to high-value specialty crop regions.

Computational experiments on the Kerman Province pistachio network demonstrate both the operational benefits and computational challenges of the proposed framework. Scalability analysis of the monolithic MIP establishes a clear tractability frontier, motivating a clustering-based decomposition that partitions large garden sets into solvable routing subproblems while preserving solution quality. Exhaustive enumeration of two-depot configurations, combined with robust routing and k-means clustering, shows that carefully chosen decentralized multi-depot architectures can substantially reduce average flight distances and routing effort compared with centralized single-depot baselines, while robustness analysis across different budgets Γ reveals explicit cost–reliability trade-offs. The case study confirms that integrating depot quality scores, assignment distances, and robust operational costs in a unified objective leads to infrastructure designs that are both practically feasible and resilient to battery-capacity variations.

Future research can extend the proposed framework along several technically meaningful directions. Incorporating reliability considerations such as drone failure risk and maintenance scheduling would allow the model to capture operational disruptions beyond energy limitations. Acknowledging the deliberate simplification of our selective robustification—where parameters such as travel times and arc-level energy consumption rates were treated deterministically to maintain computational tractability—a richer uncertainty representation that explicitly incorporates multi-factor environmental variables (such as dynamic wind vectors, air density shifts, and temperature-induced aerodynamic drag) alongside battery degradation represents a valuable direction for future research. Furthermore, drawing methodological inspiration from recent advancements that optimize drone-motorcycle fleets for highly dynamic spatial constraints–such as routing to moving customers [22]–future models could integrate spatial uncertainty to track moving targets or dynamic pest outbreaks within the surveillance network. From a computational standpoint, alternative solution paradigms–such as column-generation-based decomposition, advanced matheuristics, or constraint programming–could be investigated to enhance scalability for larger fleets and expanded geographic coverage. Finally, embedding adaptive re-optimization mechanisms driven by real-time telemetry data would move the framework toward dynamic decision support in endurance-constrained surveillance operations.

Conflicts of Interest

The authors declare no conflict of interest.

References

[1] S. Aggarwal, N. Kumar, *Path planning techniques for unmanned aerial vehicles: A review, solutions, and challenges*, Comput. Commun. **149** (2020) 270–299.

[2] A. Ben-Tal, A. Nemirovski, L. El Ghaoui, *Robust Optimization*, Princeton University Press, 2009.

[3] D. Bertsimas, M. Sim, *The price of robustness*, Oper. Res. **52(1)** (2004) 35–53.

[4] D. Bertsimas, M. Sim, *Robust discrete optimization and network flows*, Math. Program. **98(1)** (2003) 49–71.

[5] S. Candiago, F. Remondino, M. De Giglio, M. Dubbini, M. Gattelli, *Evaluating multispectral images and vegetation indices for precision farming applications from UAV images*, Remote Sens. **7(4)** (2015) 4026–4047.

[6] M. Cardone, B. Gargiulo, E. Fornaro, *Modelling and experimental validation of a hybrid electric propulsion system for light aircraft and unmanned aerial vehicles*, Energies **14(13)** (2021) 3969.

[7] E. Cengiz, C. Yilmaz, H.T. Kahraman, Ç. Suiçmez, *Effects of variable UAV speed on optimization of travelling salesman problem with drone (TSP-D)*, in: Int. Conf. Artif. Intell. Appl. Math. Eng. (2021) 295–305.

[8] E. Cengiz, C. Yilmaz, H.T. Kahraman, *A literature review of drone-truck routing problems: challenges and future research directions*, RAIRO Oper. Res. **59(5)** (2025) 3169–3205.

[9] J. Choi, G. Stagg, C.K. Peterson, M.Z. Li, *Bi-level route optimization and path planning with hazard exploration*, IEEE 64th Conference on Decision and Control (CDC), Rio de Janeiro, Brazil, 2025, 3403–3410.

[10] B. Colson, P. Marcotte, G. Savard, *An overview of bi-level optimization*, Ann. Oper. Res. **153(1)** (2007) 235–256.

[11] M.S. Daskin, *Network and discrete location: Models, algorithms, and applications*, J. Oper. Res. Soc. **48(7)** (1997) 763–764.

[12] J. Del Cerro, C. Cruz Ulloa, A. Barrientos, J. de León Rivas, *Unmanned aerial vehicles in agriculture: A survey*, Agronomy **11(2)** (2021) 203.

[13] K. Dorling, J. Heinrichs, G.G. Messier, S. Magierowski, *Vehicle routing problems for drone delivery*, IEEE Trans. Syst. Man Cybern. Syst. **47(1)** (2016) 70–85.

[14] FAO, *Global pistachio production and trade statistics*, Food and Agriculture Organization of the United Nations, 2021.

[15] R. Gebbers, V.I. Adamchuk, *Precision agriculture and food security*, Science **327(5967)** (2010) 828–831.

[16] T.F. Geyer, L. Enghardt, *Conceptual estimation of the noise reduction potential of electrified aircraft engines*, Acta Acust. **7** (2023) 22.

[17] S. Ghahremanzadeh, et al., *Robust optimization of drone delivery networks: Integrated location-allocation and charging decisions with fuzzy-robust energy consumption under demand uncertainty*, Process Integr. Optim. Sustain. (2026).

[18] S. Ghasemi, R. Tavakkoli-Moghaddam, M. Hamid, M. Hosseinzadeh, *Sustainable facility location-routing problem for blood package delivery by drones with a charging station*, in: IFIP Int. Conf. Adv. Prod. Manag. Syst. (APMS), Springer, Cham, 2021, 3–14.

[19] Y. Guo, J. Guo, C. Liu, H. Xiong, L. Chai, D. He, *Precision landing test and simulation of the agricultural UAV on apron*, Sensors **20(12)** (2020) 3369.

[20] A. Hajimirzajan, M. Vahdat, A. Sadegheih, E. Shadkam, H. El Bilali, *An integrated strategic framework for large-scale crop planning: Sustainable climate-smart crop planning and agri-food supply chain management*, Sustain. Prod. Consum. **26** (2021) 709–732.

[21] M. Hamid, M.M. Nasiri, M. Rabbani, *A mixed closed-open multi-depot routing and scheduling problem for homemade meal delivery incorporating drone and crowd-sourced fleet: A self-adaptive hyper-heuristic approach*, Eng. Appl. Artif. Intell. **120** (2023) 105876.

[22] M. Hamid, B. Vahedi-Nouri, V. Heinz, Z. Hanzálek, *Routing and scheduling of time-sensitive pharmaceutical deliveries incorporating drones for moving customers*, Unpublished manuscript (2024).

[23] S. Khanal, J. Fulton, S. Shearer, *An overview of current and potential applications of thermal remote sensing in precision agriculture*, Comput. Electron. Agric. **139** (2017) 22–32.

[24] Y. Liu, M. Zhang, Q. Zhao, *Optimizing agricultural drone operations: From launch and recovery siting to tiered routing strategies*, arXiv preprint arXiv:2606.20239 (2026).

[25] Y. Long, G. Xu, J. Zhao, B. Xie, M. Fang, *Dynamic truck–UAV collaboration and integrated route planning for resilient urban emergency response*, IEEE Trans. Eng. Manag. **71** (2023) 9826–9838.

[26] W.H. Maes, K. Steppe, *Perspectives for remote sensing with unmanned aerial vehicles in precision agriculture*, Trends Plant Sci. **24(2)** (2019) 152–164.

[27] A. Matese, P. Toscano, S.F. Di Gennaro, et al., *Intercomparison of UAV, aircraft and satellite remote sensing platforms for precision viticulture*, Remote Sens. **7(3)** (2015) 2971–2990.

[28] C.C. Murray, A.G. Chu, *The flying sidekick traveling salesman problem: Optimization of drone-assisted parcel delivery*, Transp. Res. Part C Emerg. Technol. **54** (2015) 86–109.

[29] G. Nagy, S. Salhi, *Location-routing: Issues, models and methods*, European J. Oper. Res. **177(3)** (2007) 623–636.

[30] S.P. Parvasi, M. Mahmoodjanloo, M. Setak, *A bi-level school bus routing problem with bus stop selection and student assignment*, Int. J. Manag. Sci. Eng. Manag. **12(2)** (2017) 136–145.

[31] B. Rabta, C. Wankmüller, G. Reiner, *A drone fleet model for last-mile distribution in disaster relief operations*, Int. J. Disaster Risk Reduct. **28** (2018) 107–112.

[32] R. Rajabi, E. Shadkam, S.M. Khalili, *Design and optimization of a pharmaceutical supply chain network under COVID-19 pandemic disruption*, Sustain. Oper. Comput. **5** (2024) 102–111.

[33] J. Rezaei, *Best-worst multi-criteria decision-making method*, Omega **53** (2015) 49–57.

[34] S. Sankaran, A. Mishra, R. Ehsani, C. Davis, *A review of advanced techniques for detecting plant diseases*, Comput. Electron. Agric. **72** (2010) 1–13.

[35] D. Schermer, M. Moeini, O. Wendt, *A matheuristic for the vehicle routing problem with drones and its variants*, Transp. Res. Part C Emerg. Technol. **106** (2019) 166–204.

[36] E. Shadkam, *Cuckoo optimization algorithm in reverse logistics: A network design for COVID-19 waste management*, Waste Manag. Res. **40(4)** (2022) 458–469.

[37] K. Shahnaghi, H. Shahmoradi-Moghadam, A. Noroozi, H. Mokhtari, *A robust modelling and optimisation framework for a batch processing flow shop production system in the presence of uncertainties*, J. Ind. Syst. Eng. **29(1)** (2016) 92–106.

[38] S.M. Shavarani, M.G. Nejad, F. Rismanchian, G. Izbirak, *Application of hierarchical facility location problem for optimization of a drone delivery system*, Int. J. Adv. Manuf. Technol. **95(9-12)** (2018) 3141–3153.

[39] S. Shebalov, D. Klabjan, *Robust airline crew pairing: Move-up crews*, Transp. Sci. **40(3)** (2006) 300–312.

[40] A. Sinha, P. Malo, K. Deb, *A review on bi-level optimization: From classical to evolutionary approaches and applications*, IEEE Trans. Evol. Comput. **22(2)** (2017) 276–295.

[41] S. Sonkar, P. Kumar, Y.T. Puli, R.C. George, D. Philip, A.K. Ghosh, *Design & implementation of an electric fixed-wing hybrid VTOL UAV for asset monitoring*, J. Aerosp. Technol. Manag. **15** (2023) e0823.

[42] A. Taheri, A. Ghodousian, R. Abedian, *Review of path planning models, environmental constraints, and application domains in drone delivery systems*, J. Algorithms Comput. **56(1)** (2024) 15–33.

[43] C.A. Thiels, J.M. Aho, S.P. Zietlow, D.H. Jenkins, *Use of unmanned aerial vehicles for medical product transport*, Air Med. J. **34(2)** (2015) 104–108.

[44] L. Timilsina, P.R. Badr, P.H. Hoang, G. Ozkan, B. Papari, C.S. Edrington, *Battery degradation in electric and hybrid-electric vehicles: A survey study*, IEEE Access **11** (2023) 42431–42462.

[45] D.C. Tsouros, S. Bibi, P.G. Sarigiannidis, *A review on UAV-based applications for precision agriculture*, Information **10(11)** (2019) 349.

[46] X. Wang, Y. Li, Z. Chen, *Data-driven robust optimization model for hybrid truck-drone last-mile delivery logistics*, Enterp. Inf. Syst. (2026) 1–24.

[47] C. Yılmaz, E. Cengiz, H.T. Kahraman, *A new evolutionary optimization algorithm with hybrid guidance mechanism for truck-multi drone delivery system*, Expert Syst. Appl. **245** (2024) 123115.

[48] S. Yu, J. Zhu, J. Zhou, J. Cheng, X. Bian, J. Shen, P. Wang, *Key technology progress of plant-protection UAVs applied to mountain orchards: A review*, Agronomy **12(11)** (2022) 2828.

[49] Q. Yuan, R. Li, B. Zhou, J. Lu, M. Hu, P. Lai, et al., *Collaborative truck-UAV delivery routing optimization under dynamic weather conditions and customer demands*, IEEE Trans. Consum. Electron. (2025).

[50] N. Zhang, M. Wang, N. Wang, *Precision agriculture—a worldwide overview*, Comput. Electron. Agric. **36(2-3)** (2002) 113–132.

[51] J. Zong, B. Zhu, Z. Hou, X. Yang, J. Zhai, *Evaluation and comparison of hybrid wing VTOL UAV with four different electric propulsion systems*, Aerospace **8(9)** (2021) 256.