

Drone-enhanced last-mile medical delivery in rural areas: A weather-aware optimization approach

Pardis Bahmani^a, Shakiba Enayati^{a,*}, Sina Ansari^b, Srikanth Gururajan^c, Ziteng Wang^d

^a Supply Chain and Analytics Department, College of Business Administration, University of Missouri-St. Louis, 1 University Blvd, St. Louis, MO 63121, USA

^b Department of Management and Entrepreneurship, Driehaus College of Business, DePaul University, 1 E Jackson Blvd, Chicago, IL 60604, USA

^c Department of Aerospace and Mechanical Engineering, Saint Louis University, 3450 Lindell Blvd, St. Louis, MO 63103, USA

^d Department of Industrial and Systems Engineering, Northern Illinois University, 590 Garden Rd, DeKalb, IL 60115, USA

ARTICLE INFO

Keywords:

Scheduling

Routing

Multi-modal networks

Drone delivery

Weather-aware healthcare logistics

ABSTRACT

Rural communities face persistent challenges in accessing timely and reliable medical supply delivery due to sparse infrastructure, long travel distances, and limited healthcare resources. This study presents a weather-aware, multi-modal optimization framework for operational-level medical delivery in rural regions using drones and ground vehicles. The proposed model jointly integrates routing and scheduling decisions with transshipment, recharging operations, and spatiotemporal energy consumption derived from forecast-based weather information, enabling operationally realistic planning under environmental and infrastructure constraints. Computational experiments based on real-world data from the Shield Illinois testing program evaluate instances across four operating scenarios that vary in demand volume and spatial dispersion. Relative to decentralized ground-based delivery, centralized optimization of ground vehicle routing alone reduces total system cost by 40–83%. Incorporating drone-ground coordination yields an additional 38–62% cost reduction, while drone-enabled configurations reduce weighted average delivery time by up to 85% and CO₂ emissions by more than 99%. High-volume scenarios favor drone-only operation due to parallel point-to-point service and sufficient fleet capacity, whereas low-volume but spatially dispersed networks benefit from multi-modal coordination through reduced travel distances and improved feasibility. The results demonstrate that spatial dispersion, fleet capacity, and energy feasibility, rather than demand volume alone, determine the operational value of multi-modality in rural medical logistics. The proposed framework provides a practical decision-support tool for energy-aware, weather-resilient planning in drone-assisted healthcare delivery systems. Compared to simplified range-based models commonly used in the literature, the proposed approach more accurately captures energy demands and enables better-informed routing decisions, offering a robust platform for navigating trade-offs between cost, timeliness, and network complexity in rural healthcare logistics.

* Corresponding author.

E-mail addresses: pbhp5@umsystem.edu (P. Bahmani), senayati@umsl.edu (S. Enayati), sina.ansari@depaul.edu (S. Ansari), srikanth.gururajan@slu.edu (S. Gururajan), zwang3@niu.edu (Z. Wang).

<https://doi.org/10.1016/j.trc.2026.105920>

Received 20 July 2025; Received in revised form 29 June 2026; Accepted 7 August 2026

Available online 18 August 2026

0968-090X/© 2026 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

1. Introduction

The medical supply chain is a critical component of global health systems, ensuring that essential medicines and supplies reach healthcare facilities and patients in need (Kruk et al., 2018). Within this system, last-mile delivery represents the final and most challenging stage, connecting central distribution hubs to often remote or underserved locations. The efficiency and reliability of last-mile delivery are crucial for developing sustainable, equitable, and responsive health systems, which are vital to global health security and human well-being (Backman et al., 2008).

Despite its critical role, last-mile delivery continues to face significant challenges in rural and underserved regions. Inadequate infrastructure, difficult terrain, and high transportation costs often delay access to essential medical supplies such as vaccines and test kits (Enayati et al., 2023). These challenges are intensified in regions with provider shortages, high poverty, and aging populations (Administration, n.d.). During the COVID-19 pandemic, such delivery barriers hindered timely response efforts and contributed to worsened health outcomes (Goldschmidt and Stasko, 2022). Innovative and resilient solutions are therefore essential to overcome the persistent challenges of last-mile delivery in rural areas.

Drones, or Unmanned Aerial Vehicles (UAVs), offer a promising alternative to traditional ground-based systems for last-mile medical delivery in rural areas (Jeon et al., 2022; Macrina et al., 2020). In such settings, drones benefit from lower airspace density, reduced interaction with dense built infrastructure, and operational conditions that can facilitate long-distance aerial delivery, factors that enhance their feasibility and practical applicability. Their speed, route flexibility, and ability to bypass challenging or underdeveloped road networks make them especially effective in reaching remote areas (Lamprey and Serwaa, 2020; Griffith et al., 2023). As healthcare logistics increasingly prioritize automation, sustainability, and rapid response, drones present an efficient and scalable last-mile solution (Stierlin et al., 2024; Bosona, 2020).

However, integrating drones into last-mile healthcare delivery introduces operational complexities that couple routing and scheduling decisions with stringent, weather-dependent energy feasibility constraints. Drone performance, determined by factors such as flight range, battery capacity, and payload, is highly sensitive to environmental conditions, particularly weather (Dorling et al., 2017; Troudi et al., 2018). Time-dependent wind patterns, for instance, can substantially increase energy consumption and reduce feasible flight range, while favorable conditions such as tailwinds may improve efficiency (Thibbotuwawa et al., 2019a; Betti Sorbelli, 2024). Despite growing interest in UAV-based delivery, many existing studies rely on fixed or range-based energy representations that assume a constant effective flight range, without explicitly capturing how spatiotemporal variations in environmental conditions alter the drone's actual operational range and routing feasibility (Zhang et al., 2021; Chen et al., 2021). Addressing this gap is critical for enabling resilient and operationally feasible last-mile healthcare logistics systems.

To address the overlooked impact of environmental fluctuations on delivery system performance, this study develops a deterministic, operational-level decision-support optimization model for last-mile medical delivery over a short single-period planning horizon (e.g., one day). The model represents a multi-modal logistics network comprising drones, ground vehicles, and other potential transportation modes, and explicitly considers weather-aware drone energy consumption to improve operational feasibility, reliability, efficiency and resilience under resource constraints. The proposed model, however, does not focus on high-fidelity flight dynamics or vehicle design; instead, vehicle behavior is deliberately abstracted to balance physical realism with optimization tractability and implementation feasibility in real-world healthcare logistics settings.

The objective is to minimize total delivery costs, including transportation, recharging, and transshipment, while jointly optimizing routing decisions, final depot assignment, selection of recharging/refueling and transshipment stations from pre-specified candidate locations, and pickup and delivery scheduling within the operational period. The logistics network features designated operational bases, from which each vehicle departs and completes its route at an optimally selected final depot. Each transportation mode operates along its own set of feasible network connections. To enable coordination across modes, designated transshipment points enable the transfer of medical orders between vehicles. The model also determines when and where recharging/refueling is required to maintain minimum energy levels, thereby ensuring continuous and reliable operation across the delivery network.

To capture the localized impact of environmental conditions on drone performance, the service area is partitioned into weather-based sub-regions using a Voronoi diagram, with each region assigned to its nearest weather station. This partitioning is used solely to associate every candidate drone route segment with localized time-dependent forecast data for energy estimation, and does not impose airspace or operational boundaries. For each drone arc, energy consumption is represented as a departure-time-dependent function derived from a performance model calibrated to the drone's technical specifications and the localized weather forecasts. These arc-level time- and location-dependent energy functions are then embedded directly into the optimization framework, enabling energy-aware routing and scheduling decisions under spatiotemporal environmental conditions. While this study focuses on aerial vehicles, the proposed modeling structure is sufficiently general to accommodate departure-time-dependent energy consumption for other transportation modes.

This study offers three key contributions to the medical supply chain literature. First, it develops an operational-level optimization model that jointly determines scheduling, routing, and the identification of critical operational elements, such as transshipment points, recharging stations, and final depots, within a single-period planning horizon. Second, to the best of our knowledge, this paper is the first to introduce a time-dependent energy consumption modeling approach, where drone energy use is predicted based on weather forecasts and incorporated into the optimization model as a function of departure time, capturing spatiotemporal environmental impacts on operational performance. Third, the model is applied to a real-world case study, demonstrating its effectiveness in optimizing last-mile medical deliveries in rural areas and improving healthcare access in resource-constrained settings through weather-aware, multi-modal coordination.

The remainder of this paper is structured as follows: [Section 2](#) provides a comprehensive review of the literature on drone-integrated delivery logistics networks and weather-aware supply chain optimizations that integrate drones. [Section 3](#) outlines the problem description and model assumptions. [Section 4](#) presents the network inputs, mathematical formulation of the proposed optimization model, and the energy consumption modeling framework. [Section 5](#) presents the computational preprocessing and solver-based implementation procedures used to improve the tractability of the proposed optimization model. [Section 6](#) explains the case study that contextualizes our research. The numerical results, sensitivity analysis, and managerial insights are presented in [Section 7](#). Finally, [Section 8](#) concludes the paper and suggests areas for future research.

2. Literature review

This section reviews relevant literature on drone-integrated delivery networks and the integration of weather conditions into supply chain optimization. We first examine existing work on drone-based logistics, followed by studies on weather-integrated optimization in healthcare delivery. Finally, we identify key research gaps and position the contributions of this study within the existing body of work.

2.1. Drone-integrated delivery networks

Recent studies highlight the growing interest in drone delivery networks, particularly for reducing costs and improving delivery times in sectors like healthcare and humanitarian logistics ([Zhang et al., 2023](#)). [Betti Sorbelli \(2024\)](#) offers a broad review of drone delivery applications, while [Gunaratne et al. \(2022\)](#) demonstrate their effectiveness in healthcare delivery in low-income regions. [Stierlin et al. \(2024\)](#) discuss key advancements and challenges in using drones for transporting medical samples such as blood and organ tissue.

Optimization models for drone routing have addressed a range of operational goals and constraints. [Kim et al. \(2021\)](#) develop a routing algorithm to minimize delivery distance, while [Dorling et al. \(2017\)](#) focus on reducing cost and time in two delivery scenarios. [Schermer et al. \(2019\)](#) design a truck-drone logistics network to minimize the makespan, while [Radzki et al. \(2019\)](#) emphasize customer satisfaction, incorporating time, payload, and battery constraints. [Shi et al. \(2022\)](#) propose a bi-objective mixed-integer linear program to balance cost efficiency and delivery urgency in medical supply logistics. Building on this literature, our model similarly minimizes total system costs, including transportation and operational expenses, while accounting for time windows, payload limits, and battery capacity.

A key challenge in drone-integrated delivery networks is limited battery capacity, which constrains flight range. To address this, prior studies have proposed solutions such as strategically placed charging stations ([Pinto and Lagorio, 2022](#)), machine learning-based routing ([Alyassi et al., 2023](#)), and optimization models minimizing energy and time costs ([Dukkanci et al., 2021](#)) under limited battery resources ([Liu et al., 2025](#)). [Moadab et al. \(2022\)](#) further explore the use of public transit vehicles as mobile charging stations to improve energy efficiency. However, these approaches often overlook the impact of time-dependent weather on drone energy consumption, which is critical for determining optimal recharging times and locations. This study addresses this gap by incorporating time-dependent forecast weather data into energy estimation and routing decisions.

Several studies have also explored multi-modal logistics networks that combine drones with ground vehicles to enable transshipment and improve delivery efficiency. [Kim et al. \(2024\)](#) present a framework leveraging the complementary strengths of both modes. [Wolfinger et al. \(2019\)](#) formulate a long-haul routing model with transshipment nodes and time windows, demonstrating cost savings through the integration of both short- and long-haul routes. [Enayati et al. \(2023\)](#) propose a strategic vaccine distribution model that optimizes drone bases, relay stations, and distribution centers under cold chain and drone flight range constraints. [Ramos and Vigo \(2023\)](#) integrate cars and drones for same-day medical deliveries under lead-time limits in dynamic order settings. Building on this line of work, our study integrates drones and ground vehicles into a unified network, jointly determining optimal transshipment points and final depot locations to enhance operational efficiency.

Given the critical role of timing in healthcare delivery, several studies have incorporated time windows into their optimization models ([Enayati et al., 2023](#); [Wolfinger et al., 2019](#); [Ramos and Vigo, 2023](#)). [Lakhwani et al. \(2025\)](#) propose a metaheuristic approach inspired by natural processes to optimize blood bag delivery using trucks and drones, accounting for route degradation, vehicle range, and time-sensitive perishability. Similarly, our model includes given time windows for pickups and deliveries to reflect real-world healthcare constraints.

Despite recent advances, the weather remains a major barrier in drone delivery systems, affecting energy consumption, flight range, and operational time ([Koshta et al., 2024](#); [Beigi et al., 2022](#)). This challenge has prompted the development of optimization models that incorporate weather impacts into network design and decision-making ([Bocewicz et al., 2022](#)). The following subsection reviews studies that examine the effects of weather on drone-integrated supply chain networks.

2.2. Weather-integrated optimization for drone delivery

Weather conditions significantly influence the operational performance of drone-integrated logistics networks. Motivated by this, recent studies have increasingly incorporated weather effects, either directly or indirectly, into optimization frameworks to evaluate and improve drone performance.

Some studies consider static or simplified representations of weather conditions (e.g., scenario-based) in optimization models. For instance, [Thibbotuwawa et al. \(2020\)](#) develop a drone mission planning model that accounts for wind effects on energy consumption

by comparing constant ground-speed and constant airspeed flight strategies. [Chen et al. \(2023\)](#) examine a vehicle routing problem (VRP) under multiple wind scenarios, while [Gürel and Serdarasan \(2024\)](#) explore wind-integrated planning for sustainable last-mile delivery. These models do not explicitly incorporate forecast-driven, time-dependent environmental conditions. This study addresses this limitation by embedding time- and location-dependent weather forecasts into the optimization model, enabling more accurate energy estimation and weather-aware routing and scheduling decisions.

Other studies move toward more dynamic representations of weather. Building on soft time window VRPs ([Guerriero et al., 2014](#)) and mission planning under weather dependencies ([Thibbotuwawa et al., 2019b](#)), [Radzki et al. \(2019\)](#) develop a constraint programming model that integrates non-linear, weather-dependent energy consumption for drone routing. Their approach ensures collision-free deliveries within time windows while maximizing customer satisfaction and delivery volume. [Lakhwani et al. \(2025\)](#) incorporate wind speed as a limiting factor affecting drone range over time, while [Cheng et al. \(2024\)](#) propose a two-period data-driven scheduling model that incorporates wind information to address flight-time uncertainty and enable schedule adjustments, highlighting the importance of weather information in drone operations within a multi-period decision structure. Our approach complements this line of research by focusing on single-period operational planning and embedding forecast-based, time- and location-dependent energy consumption directly into the routing and scheduling formulation. By modeling energy use as an explicit function of departure time and arc-specific wind conditions, the proposed framework captures spatiotemporal variability in energy requirements while preserving a tractable single-stage optimization structure. Moreover, we extend weather-aware modeling to a multi-modal logistics setting involving drones and ground vehicles, candidate facilities, and transshipment decisions, thereby linking localized environmental conditions to integrated routing and mode-assignment decisions. To the best of our knowledge, no prior study has integrated forecast-based, time-dependent weather information into an operational, drone-assisted, multi-modal routing and scheduling framework in this manner.

2.3. Summary of research gaps and model contributions

[Table 1](#) highlights key characteristics of the literature most pertinent to this study. While the growing body of literature on drone-integrated medical delivery has advanced understanding of key operational factors, such as multi-modal coordination, time windows, payload limits, battery capacity, and recharging or transshipment logistics, many existing models fall short in capturing the complexity of time-dependent, system-wide operational decision-making. One critical limitation is the limited integration of time-dependent weather conditions, which directly affects drone performance by altering energy consumption, flight range, and travel time.

Although some studies consider weather through static assumptions or scenario-based models, they often treat environmental conditions as fixed or uniform, failing to capture their time-dependent and location-specific nature. As a result, energy estimation and routing feasibility are often inaccurate. While some studies have explored the effects of weather variability on drone operations through scenario analysis or flight time uncertainties, none, to the best of our knowledge, have integrated forecast-based, time-dependent prediction of energy consumption into the optimization model. Embedding this predictive capability alongside critical operational decisions, such as routing, scheduling, recharging, and transshipment, remains a key gap in the literature.

Building on the identified gaps, this study proposes a comprehensive operational-level optimization framework that jointly determines scheduling, routing, and the selection of key operational elements such as transshipment points, recharging stations, and final depots during a single-period planning horizon. The model incorporates forecast-based, time-dependent weather information to improve predictions of drone energy consumption and supports coordination across multiple vehicle types. To our knowledge, no existing study has integrated all of these components into a unified decision-support tool for medical logistics under forecast-based environmental variability. [Section 3](#) formalizes the problem and outlines the structure of the delivery network under study.

3. Problem description

This section describes the operational medical delivery problem considered in this study. It begins with an overview of the problem scope, followed by a detailed description of the delivery system and its key components. Finally, the modeling assumptions underlying this operational setting are introduced.

3.1. Problem overview

This study considers a deterministic, operational-level last-mile medical delivery problem in a rural service area over a single operating period. A set of medical orders, including laboratory specimens, vaccines, test kits, and medications, must be transported from geographically dispersed pickup locations to designated delivery destinations, such as pharmacies, clinics, laboratories, or patient homes, within specified time requirements. Deliveries are carried out using a heterogeneous, multi-modal fleet composed of drones and ground vehicles. The objective is to determine cost-efficient routing and scheduling plans that satisfy payload limits, pickup and delivery time windows, routing feasibility, and drone battery endurance, while explicitly accounting for time- and location-dependent weather impacts on drone energy consumption through forecast-based energy inputs. Within this framework, the model also determines which pre-specified candidate facilities are activated and utilized during the operating period.

The operational network consists of pickup and delivery locations, predetermined vehicle departure depots, and a set of pre-specified candidate sites that may function as final depots, transshipment points, and/or recharging locations. Strategic decisions, such as facility siting, fleet sizing, and infrastructure investment, are assumed to have been made prior to the planning horizon and

Table 1

Summary of the most relevant studies.

References	Objective Function	Problem Type	Network Structure				Transportation Mode			Weather Integration		Operational Constraints				
			Transshipment	Pickup	Recharge	Drone	Ground	Static	Time-Dependent	Weather Incorporation Approach	Battery Capacity	Time Window	Payload Capacity	Solution		
F. Guerriero et al. (2014)	Min TD, Max CS, Min NV	non-linear				*						*		heuristic		
Dorling et al. (2017)	Min TC, Min DT	linear	*			*					*	*	*	exact		
Wolffinger et al. (2019)	Min TC			*			*					*	*	heuristic		
Thibbotuwawa et al. (2019b)	Max CS	non-linear		*		*			*	Time-segmented wind updates	*	*	*	heuristic		
Schermer et al. (2019)	Min DT	linear		*		*	*			Scenario-based changing wind conditions	*	*	*	metaheuristic		
Radzki et al.(2019)	Max CS			*		*				Wind-adjusted air-speed & ground speed	*	*	*	CP		
Thibbotuwawa et al. (2020)	Max CS			*		*		*						CP		
Dukkanci et al.(2021)	Min TC	non-linear		*		*					*	*	*	exact		
Kim et al. (2021)	Min TD	linear		*		*							*	exact		
Moadab et al. (2022)	Min EC	linear		*	*	*					*	*	*	exact		
Alvassi et al. (2023)	Min TT	linear		*	*	*					*	*	*	heuristic		
Shi et al. (2022)	Min TC, Min TD	linear		*	*	*					*	*	*	heuristic		
Pinto and Lagorio (2022)	Min TC, Min TT	linear		*	*	*					*	*	*	exact		
Bocewicz et al. (2022)	Min TT	non-linear		*	*	*			*	Reactive routing with dynamic wind changes	*	*	*	heuristic		
Chen et al. (2023)	Min TC	linear		*	*	*		*		Scenario-based stationary wind	*	*	*	heuristic		
Ramos et al. (2023)	Min TC	linear		*	*	*	*				*	*	*	exact		
Enayati et al. (2023)	Min TC	linear		*	*	*	*				*	*	*	exact		
Cheng et al. (2024)	Min R	linear		*	*	*			*	Data-driven clustered wind uncertainty (two-stage scheduling)	*	*	*	exact		
Gurel et al. (2024)	Min TT	linear		*	*	*		*		Static wind-adjusted drone speed	*	*	*	exact		
Lakhwani et al. (2025)	Min TD, Min TT, Min EC	linear		*	*	*	*		*	Dynamic wind-based range reduction	*	*	*	metaheuristic		
Liu et al. (2025)	Min TC	linear		*	*	*	*		*	Forecast-based, time- and location-specific weather-driven energy prediction	*	*	*	exact		
This paper	Min TC	linear	*	*	*	*	*		*		*	*	*	exact		

Abbreviations: TC = total cost, DT = delivery time, TT = travel time, TD = traveled distance, EC = energy consumption, R = risk of late delivery, CS = customer satisfaction, NV = number of used vehicles, CP = constraint programming.




Fig. 1. Schematic illustration of a feasible solution for the multi-modal delivery system in this study. The solution includes a fleet consisting of two drones and one ground vehicle. The ground vehicle departs from an initial depot that also serves as a pickup location, transfers its load to a drone at a transshipment location, and then travels to its final depot. Each drone serves one or multiple pickup locations and undergoes one recharging operation, one of which occurs at a node that also serves as a pickup location, before completing delivery to the delivery location.

are therefore treated as exogenous to the operational planning problem considered in this study. Fig. 1 illustrates a feasible solution for this multi-modal delivery system.

3.2. Delivery system description

The delivery system modeled in this study includes medical orders, transportation modes, logistics facilities, transshipment operations, and energy replenishment mechanisms within a unified operational framework.

Each medical order is defined by a known quantity and a delivery time window, including an earliest allowable pickup time and a latest delivery deadline, reflecting the time-sensitive nature of healthcare logistics. In addition, delivery destinations may be subject to handling or receiving capacity limits during the operating period, which restrict the volume of orders processed within the planning horizon.

Transportation is performed by a heterogeneous fleet of ground vehicles and drones operating over mode-specific, capacitated networks. Ground vehicles travel along the road network, while drones traverse a set of pre-validated aerial arcs; both modes are associated with mode-specific travel times and transportation costs.

Coordination between modes is enabled through candidate transshipment nodes, where medical orders may be transferred between vehicles. These transfers incur handling time but enhance operational flexibility by leveraging the complementary capabilities of aerial and ground transportation. All vehicles depart from given initial operational bases and complete their routes at a final depot selected by the model.

Energy replenishment is a critical operational component of the system. Drones undergo battery swapping, while ground vehicles may be recharged or refueled at pre-specified candidate locations. These activities incur time and operating costs and must be jointly coordinated with routing and scheduling decisions to maintain feasibility throughout the operating period. For drones, energy consumption along each arc varies by time and location due to weather conditions, such as wind. This dependence is captured through forecast-based, time-dependent energy-consumption functions incorporated into the optimization model as fixed inputs over the planning horizon, thereby linking routing and recharging decisions to prevailing environmental conditions.

To obtain localized weather inputs across the service area, each network element is associated with its nearest monitoring station using a Voronoi partition. Because weather data are available only at discrete stations, this Voronoi-based association serves as a spatial interpolation mechanism that assigns time-dependent weather forecasts to each arc for energy-consumption estimation, without imposing routing or airspace constraints.

3.3. Modeling assumptions

The model relies on the following assumptions:

- **Operational scope:** Facility locations, vehicle types, fleet size, and the set of candidate sites are predetermined. The model optimizes routing, scheduling, order assignment, vehicle utilization, recharging/refueling, transshipment, and energy-feasibility decisions over a single operating period. Accordingly, the available fleet is treated as an input to the operational planning problem, and the model determines how the given vehicles are used rather than how many vehicles should be acquired or located. Strategic decisions such as vehicle acquisition, fleet sizing, and depot-location planning are outside the scope of this study.
- **Fleet ownership and cost structure:** Fixed vehicle acquisition costs are not modeled explicitly, as the problem focuses on short-term operational planning over a single operating period. In many real-world healthcare logistics settings, organizations responsible for medical deliveries (e.g., hospitals, laboratories, and public health agencies) outsource transportation services to third-party logistics providers. Under such agreements, fleet ownership and long-term capital costs are borne by service providers and are reflected implicitly in the transportation fees charged per route or delivery. Accordingly, the model represents variable costs through mode- and arc-specific operating parameters that aggregate operator fees, vehicle usage, maintenance, and infrastructure-related expenses, and are assumed to remain time-independent over the operating period.

- **Deterministic inputs:** Order quantities, travel times, and service requirements are assumed to be known and deterministic for the operating period.
- **Forecast-based weather:** Time-dependent weather conditions are incorporated using forecast information available prior to the operating period. The resulting energy-consumption parameters are treated as fixed inputs during optimization. Real-time updates, dynamic recourse actions, and explicit uncertainty modeling are beyond the scope of this study.
- **Energy modeling abstraction:** Drone energy consumption is represented using arc-based, departure-time-dependent parameters derived from established performance models and incorporated into the optimization model as fixed inputs over the planning horizon. Trajectory-level flight dynamics, in-flight control, and continuous vehicle dynamics are not explicitly modeled and are assumed to be handled by onboard flight-control systems operating under standard cruise conditions. This abstraction is a deliberate modeling choice: incorporating high-fidelity vehicle dynamics and platform-specific behaviors would significantly increase model complexity and shift the problem from operational optimization to simulation-based flight modeling. Given the operational planning scope of this study and the limited availability of platform-specific flight data, the adopted representation provides a pragmatic and implementable basis for network-level decision-making.
- **Energy replenishment:** At the start of the operating period, all vehicles are assumed to be fully charged or fueled. Ground vehicle recharging/refueling and drone battery swapping are permitted only at designated candidate locations and require fixed, vehicle-type-specific service times that are standardized across all locations.
- **Battery degradation:** Long-term battery degradation, aging effects, and capacity fade are not considered, as the study focuses on short-term, single-period operational planning where such effects are negligible.
- **Service region characteristics:** The model focuses on a predominantly rural service area and therefore does not explicitly account for urban-specific factors such as traffic congestion, complex airspace restrictions, or dense infrastructure. Consistent with this scope, the model emphasizes network-level operational decisions rather than platform-specific vehicle design, which would require substantially different data, modeling paradigms, and validation frameworks.
- **Feasible arc set and connectivity:** The network includes only pre-validated arcs. Drone arcs are assumed to satisfy baseline airspace feasibility and command-and-control (C2) connectivity requirements (i.e., the availability of reported communication coverage between the drone and its control infrastructure along the flight path), and are down-selected by preprocessing that excludes infeasible arcs prior to optimization.
- **Orders and handling:** Each medical order is indivisible, and split deliveries are not permitted. Transfers may occur only at selected transshipment points. Loading, unloading, and transfer times are represented at an aggregate level and incorporated as fixed parameters. No additional infrastructure or resources are explicitly modeled to facilitate transfers.

4. Optimization model

This section presents the mixed-integer optimization model for the deterministic single-period multi-modal medical delivery problem described in Section 3. After defining the sets, parameters, and decision variables, we present the objective function and constraints that capture routing, scheduling, facility selection, and drone battery feasibility, accounting for weather-dependent energy consumption. The section concludes with a description of how precomputed, time-dependent energy parameters are incorporated into the model.

4.1. Network input notations

The medical delivery network is represented as a directed network, where the set of nodes $\mathcal{N}$ represents discrete geographic locations, and the set of arcs $\mathcal{A}$ defines the connections between nodes in a rural setting. Each arc in $\mathcal{A}$ is represented by a tuple (i, j, v) , where node $i \in \mathcal{N}$ is connected to node $j \in \mathcal{N}$ via vehicle $v \in \mathcal{V}$. Set $\mathcal{V}$ indicates the fleet of available vehicles in the network, which may include various transportation modes, e.g., drones and cars. The vehicle index is used to represent an individual vehicle rather than solely a physical vehicle type. Although multiple vehicles may share the same physical characteristics, they are not necessarily interchangeable in the operational planning setting considered here because the same vehicle types may differ in their starting operational bases, feasible final-depot sets, and operational roles within the network. Thus, the tuple (i, j, v) represents both travel between two locations and the feasibility of that movement for a specific vehicle.

The network includes five types of facilities: (1) Operational bases for each vehicle $v \in \mathcal{V}$, where the vehicle must start and end its route. The starting operational base for vehicle v is denoted by parameter $O^v \in \mathcal{N}$, while the set of possible operational final depots is defined by $\mathcal{O}^v \subset \mathcal{N}$; (2) Recharging or refueling candidate locations for each vehicle type $v \in \mathcal{V}$ takes place at designated nodes $\mathcal{P}^v \subset \mathcal{N}$, where vehicles can restore energy to continue their routes; (3) Transshipment candidate locations, where orders can be transferred between vehicles, denoted by the set $\mathcal{H} \subset \mathcal{N}$; (4) Pickup locations $P_r \in \mathcal{N}$ for each order $r \in \mathcal{R}$, from which the order r must be picked up; (5) Destination locations, where the orders must be delivered, indicated by the set $\mathcal{D} \subset \mathcal{N}$. Each node $i \in \mathcal{N}$ may serve one or more facility roles. The overall set of nodes is defined as: $\mathcal{N} = (\{O^v \mid v \in \mathcal{V}\} \cup \{P^v \mid v \in \mathcal{V}\} \cup \mathcal{H} \cup \mathcal{D} \cup \{P_r \mid r \in \mathcal{R}\} \cup \{O^v \mid v \in \mathcal{V}\})$.

Each order $r \in \mathcal{R}$ originates from a pickup location $P_r \in \mathcal{N}$, with a quantity denoted by Q_r . The total amount of orders carried by any vehicle $v \in \mathcal{V}$ across its assigned arcs must not exceed its capacity U^v . Each destination node $d \in \mathcal{D}$ can receive a maximum order amount of K_d . The cost of transporting orders on arc $(i, j, v) \in \mathcal{A}$ is given by C_{ij}^v , which is distance-dependent. Recharging or refueling vehicle v incurs a cost ω^v , while transferring orders between vehicles incurs a fixed cost λ .

Table 2
Summary of the network notation for the input sets and indexes.

Notation	Description
$\mathcal{R}$	Set of the orders, $r \in \mathcal{R}$
$\mathcal{P}^v$	Set of the candidate nodes for recharging vehicle $v \in \mathcal{V}$, $\mathcal{P}^v \subset \mathcal{N}$
$\mathcal{O}'^v$	Set of the candidate operational bases at which vehicle $v \in \mathcal{V}$ ends its route, $\mathcal{O}'^v \subset \mathcal{N}$
$\mathcal{H}$	Set of the available transshipment nodes, $\mathcal{H} \subset \mathcal{N}$
$\mathcal{D}$	Set of destination nodes within the network, $d \in \mathcal{D}$
$\mathcal{N}$	Set of all nodes in the network
$\mathcal{V}$	Set of all available vehicles in the network, $v \in \mathcal{V}$
$\mathcal{A}$	Set of arcs connecting nodes $i, j \in \mathcal{N}$ via vehicle $v \in \mathcal{V}$, $(i, j, v) \in \mathcal{A}$

Table 3
Summary of the model input parameters.

Parameter	Description
C_{ij}^v	Cost of shipping orders through the arc $(i, j, v) \in \mathcal{A}$
ω^v	Cost of recharging/refueling vehicle $v \in \mathcal{V}$
λ	Cost of transshipping an order from one vehicle to another vehicle
Q_r	Quantity of order $r \in \mathcal{R}$
K_d	Maximum order amount that can be shipped into destination $d \in \mathcal{D}$
U^v	Transportation capacity of vehicle $v \in \mathcal{V}$
P_r	Pickup location associated with order $r \in \mathcal{R}$
O^v	Operational base for vehicle $v \in \mathcal{V}$ to begin its route, $O^v \in \mathcal{N}$
a_{P_r}	Earliest available time for order $r \in \mathcal{R}$ to be picked up from location P_r
b_r	Latest allowable time for order $r \in \mathcal{R}$ to be delivered to a destination
l_d	Latest allowable delivery time for any order to be delivered to destination $d \in \mathcal{D}$
θ_{ij}^v	Time required for traveling from node i to node j via arc $(i, j, v) \in \mathcal{A}$
g^v	Time required for recharging vehicle $v \in \mathcal{V}$
h^{vn}	Time required to transfer orders from vehicle $v \in \mathcal{V}$ to vehicle $n \in \mathcal{V}$
$E_{\min}^v$	Minimum energy reserve required for vehicle $v \in \mathcal{V}$ during its route to ensure safety
$E_{\max}^v$	Maximum energy capacity of vehicle $v \in \mathcal{V}$
M	A positive large number
$\mathcal{E}_{ij}^v(\mathcal{T}'_{i^v})$	Function to calculate energy consumption for vehicle $v \in \mathcal{V}$ departing from node $i \in \mathcal{N}$ to node $j \in \mathcal{N}$ upon variable departure time $\mathcal{T}'_{i^v}$

Order $r \in \mathcal{R}$ must be within a specified time window starting at a_{P_r} , and delivered to a destination $d \in \mathcal{D}$ no later than b_r . Also, l_d indicates the latest allowable delivery time for any order delivered to destination $d \in \mathcal{D}$. The parameter θ_{ij}^v represents the travel time for vehicle $v \in \mathcal{V}$ on arc $(i, j, v) \in \mathcal{A}$, while h^{vn} indicates the time required to transship an order from vehicle $v \in \mathcal{V}$ to vehicle $n \in \mathcal{V}$. Each vehicle $v \in \mathcal{V}$ has a maximum energy capacity of $E_{\max}^v$ and a minimum energy reserve requirement of $E_{\min}^v$ to ensure operational safety. The energy consumed while traversing arc $(i, j, v) \in \mathcal{A}$ is modeled as a parametric function of the vehicle's departure time $\mathcal{T}'_{i^v}$ from node $i \in \mathcal{N}$, denoted by $\mathcal{E}_{ij}^v(\mathcal{T}'_{i^v})$. The time required for recharging/refueling each vehicle $v \in \mathcal{V}$ is given by g^v .

The decision variables in the model capture key operational aspects, including arc selection for routing, assignment of final depots, order transport and transshipment, recharging/refueling operations, time scheduling, and energy management within the network. The binary variable $\mathcal{X}_{ij}^v$ indicates whether arc $(i, j, v) \in \mathcal{A}$ is included in the route of vehicle $v \in \mathcal{V}$, while $\mathcal{F}_i^v$ indicates whether node $i \in \mathcal{O}'^v$ is selected as the final depot for vehicle v . The binary variable $\mathcal{Y}_{ijr}^v$ tracks whether order $r \in \mathcal{R}$ is carried on the arc $(i, j, v) \in \mathcal{A}$. The binary variable $\mathcal{Z}_i^v$ shows if vehicle $v \in \mathcal{V}$ is recharged or refueled at node $i \in \mathcal{P}^v$. Order transfers are captured by γ_{rn}^v , which equals 1 if order $r \in \mathcal{R}$ is transshipped from vehicle $v \in \mathcal{V}$ to vehicle $n \in \mathcal{V}$ at node $i \in \mathcal{H}$. The variable δ_i indicates whether any transfer occurs at node $i \in \mathcal{H}$. Time-related decision variables $\mathcal{T}_i^v$ and $\mathcal{T}'_{i^v}$ represent the arrival and departure times of vehicles $v \in \mathcal{V}$ at node $i \in \mathcal{N}$, respectively. Similarly, energy variables E_i^v and E'_i^v monitor the remaining energy of the vehicle $v \in \mathcal{V}$ upon arrival and departure at node $i \in \mathcal{N}$. The complete notation for all sets, parameters, and decision variables of the model is provided in Tables 2–4.

4.2. Mathematical formulation

This section presents the model's objective function and associated constraints.

Objective function. As shown by Eq. (1), the objective function minimizes the total operating costs in the delivery network, including transportation (first term), vehicle recharging/refueling (second term), and order transshipment (third term) costs.

$$\min TC = \sum_{(i,j,v) \in \mathcal{A}} C_{ij}^v \mathcal{X}_{ij}^v + \sum_{v \in \mathcal{V}} \sum_{i \in \mathcal{P}^v} \omega^v \mathcal{Z}_i^v + \sum_{i \in \mathcal{H}} \lambda \delta_i \quad (1)$$

Vehicle routing constraints. Constraints (2) to (5) define feasible routes and the selection of arcs within the network. Constraints (2) ensure that if vehicle $v \in \mathcal{V}$ departs from its operational base located at O^v , it must travel to exactly one subsequent node via arc $(i, j, v) \in \mathcal{A}$. Constraints (3) require that vehicle $v \in \mathcal{V}$ completes its route at a final depot $i \in \mathcal{O}'^v$. Constraints (4) link the

Table 4
Summary of the model decision variables.

Decision Variable	Description
$\mathcal{X}_{ij}^v$	Equal to 1 if arc $(i, j, v) \in \mathcal{A}$ is selected to interconnect nodes $i \in \mathcal{N}$ and $j \in \mathcal{N}$ via vehicle $v \in \mathcal{V}$; 0 otherwise
$\mathcal{Y}_{ijr}^v$	Equal to 1 if order $r \in \mathcal{R}$ is transited from node i to node j through arc $(i, j, v) \in \mathcal{A}$; 0 otherwise
$\mathcal{Z}_i^v$	Equal to 1 if vehicle $v \in \mathcal{V}$ is recharged at node $i \in \mathcal{P}^v$; 0 otherwise
$\mathcal{F}_i^v$	Equal to 1 if node $i \in \mathcal{O}^v$ is selected to be the final operational base for vehicle $v \in \mathcal{V}$
$\mathcal{W}_{ij}^v$	Equal to 1 if node $i \in \mathcal{N}$ precedes node $j \in \mathcal{N}$ in the route of vehicle $v \in \mathcal{V}$
γ_{ir}^v	Equal to 1 if order $r \in \mathcal{R}$ is transshipped from vehicle $v \in \mathcal{V}$ to vehicle $n \in \mathcal{V}$ at node $i \in \mathcal{H}$; 0 otherwise
δ_i	Equal to 1 if at least one order is transshipped from one vehicle to another at node $i \in \mathcal{H}$; 0 otherwise
T_i^v	Arrival time of vehicle $v \in \mathcal{V}$ at node $i \in \mathcal{N}$
$T_i^{\prime v}$	Departure time of vehicle $v \in \mathcal{V}$ from node $i \in \mathcal{N}$
E_i^v	Remaining energy for vehicle $v \in \mathcal{V}$ upon arrival at node $i \in \mathcal{N}$
$E_i^{\prime v}$	Remaining energy for vehicle $v \in \mathcal{V}$ upon departure from node $i \in \mathcal{N}$

selection of final depots to the inclusion of their associated incoming arcs; if a node is selected, its incoming arc must be used, and vice versa. More details on how to implement these constraints are provided in Section 5.1. Finally, constraints (5) ensure that if a vehicle visits any intermediate node $j \in \mathcal{N}$ that is not an operational base, it must also depart from that node via a valid arc.

Constraints (6) through (8) address sub-tour elimination constraints, enforcing precedence relationships between nodes through linear ordering. To achieve this, a binary variable $\mathcal{W}_{ij}^v$ is defined, where $\mathcal{W}_{ij}^v = 1$ if node $i \in \mathcal{N}$ precedes node $j \in \mathcal{N}$ in the route of vehicle $v \in \mathcal{V}$, otherwise $\mathcal{W}_{ij}^v = 0$. Notably, candidate locations for the vehicles' final depot are not excluded from these constraints, as they have no outgoing arcs (see Section 5.1). Constraints (7) prohibit the immediate reversal of arcs, ensuring that if vehicle v traverses arc $(i, j, v) \in \mathcal{A}$, it cannot subsequently traverse the reverse arc $(j, i, v) \in \mathcal{A}$. Constraints (8) prevent the formation of sub-tours at non-depot nodes, maintaining route continuity.

$$\sum_{j: (O^v, j, v) \in \mathcal{A}} \mathcal{X}_{O^v j}^v \leq 1, \quad \forall v \in \mathcal{V} \quad (2)$$

$$\sum_{j: (O^v, j, v) \in \mathcal{A}} \mathcal{X}_{O^v j}^v = \sum_{l \in \mathcal{O}^v} \mathcal{F}_l^v, \quad \forall v \in \mathcal{V} \quad (3)$$

$$\sum_{i: (i, l, v) \in \mathcal{A}} \mathcal{X}_{il}^v = \mathcal{F}_l^v, \quad \forall v \in \mathcal{V}, l \in \mathcal{O}^v \quad (4)$$

$$\sum_{i: (i, j, v) \in \mathcal{A}} \mathcal{X}_{ij}^v = \sum_{i: (j, i, v) \in \mathcal{A}} \mathcal{X}_{ji}^v, \quad \forall v \in \mathcal{V}, j \in \mathcal{N} \setminus (\mathcal{O}^v \cup \{O^v\}) \quad (5)$$

$$\mathcal{X}_{ij}^v \leq \mathcal{W}_{ij}^v, \quad \forall (i, j, v) \in \mathcal{A}, i \neq O^v \quad (6)$$

$$\mathcal{W}_{ij}^v + \mathcal{W}_{ji}^v = 1, \quad \forall (i, j, v), (j, i, v) \in \mathcal{A}, i \neq O^v \quad (7)$$

$$\mathcal{W}_{ij}^v + \mathcal{W}_{jl}^v + \mathcal{W}_{li}^v \leq 2, \quad \forall (i, j, v), (j, l, v), (l, i, v) \in \mathcal{A}, i, j \neq O^v \quad (8)$$

Order delivery constraints. The following constraints track the movement of each order throughout the network. Constraints (9) ensure that order $r \in \mathcal{R}$ is picked up from its designated pickup location P_r by one of the available vehicles in the network, while constraints (10) ensure that the order is eventually delivered to a destination node. If an order is transferred to an intermediate (non-destination) node, constraints (11) require it to continue on to another node. Constraints (12) indicate that order r may travel from node i to node j via vehicle v only if arc $(i, j, v) \in \mathcal{A}$ is selected.

$$\sum_{(P_r, j, v) \in \mathcal{A}} \mathcal{Y}_{P_r j r}^v = 1, \quad \forall r \in \mathcal{R} \quad (9)$$

$$\sum_{(P_r, j, v) \in \mathcal{A}} \mathcal{Y}_{P_r j r}^v = \sum_{(i, d, v) \in \mathcal{A}, d \in \mathcal{D}} \mathcal{Y}_{id r}^v, \quad \forall r \in \mathcal{R} \quad (10)$$

$$\sum_{i, v: (i, j, v) \in \mathcal{A}} \mathcal{Y}_{ij r}^v = \sum_{i, v: (j, i, v) \in \mathcal{A}} \mathcal{Y}_{ji r}^v, \quad \forall r \in \mathcal{R}, j \in \mathcal{N} \setminus (\mathcal{D} \cup \{P_r\}) \quad (11)$$

$$\mathcal{Y}_{ij r}^v \leq \mathcal{X}_{ij}^v, \quad \forall (i, j, v) \in \mathcal{A}, r \in \mathcal{R} \quad (12)$$

Capacity constraints. Constraints (13) and (14) ensure capacity limitations are respected throughout the network. Constraints (13) restrict the total amount of orders transported through arc $(i, j, v) \in \mathcal{A}$ to not exceed the capacity of vehicle v . Similarly, constraints (14) limit the total amount of orders delivered to each destination $d \in \mathcal{D}$ to its maximum handling capacity.

$$\sum_{r \in \mathcal{R}} Q_r \mathcal{Y}_{ij r}^v \leq U^v \mathcal{X}_{ij}^v, \quad \forall (i, j, v) \in \mathcal{A} \quad (13)$$

$$\sum_{i, v: (i, d, v) \in \mathcal{A}} \sum_{r \in \mathcal{R}} Q_r \mathcal{Y}_{id r}^v \leq K_d, \quad \forall d \in \mathcal{D} \quad (14)$$

Transshipment constraints. Constraints (15) define a transshipment at node $i \in \mathcal{H}$ as occurring when an order $r \in \mathcal{R}$ arrives via vehicle $v \in \mathcal{V}$ through arc $(j, i, v) \in \mathcal{A}$ and is subsequently transported from node i to node $j' \in \mathcal{N}$ by vehicle $n \in \mathcal{V}$ via arc $(i, j', n) \in \mathcal{A}$.

Constraints (16) activate the binary variable δ_i , indicating whether any order is transshipped at node $i \in \mathcal{H}$. In this constraint, the Big- M constant is bounded by the total number of orders, i.e., $M = |\mathcal{R}|$, which represents the maximum number of transshipment operations that may occur at a single node. Constraints (17) enforce temporal feasibility by ensuring that vehicle $n \in \mathcal{V}$ cannot depart node $i \in \mathcal{H}$ until vehicle $v \in \mathcal{V}$ has arrived and the required order handoff has been completed. In this constraint, the Big- M constant is set to $M = |\mathcal{R}| \cdot |\mathcal{V}| \cdot h^{vn}$, providing an upper bound on the total transshipment-related time interactions at node i .

$$\sum_{j:(j,i,v) \in \mathcal{A}} \mathcal{Y}_{jir}^v + \sum_{j':(i,j',n) \in \mathcal{A}} \mathcal{Y}_{ij'r}^n \leq \gamma_{ir}^{vn} + 1, \quad \forall i \in \mathcal{H}, r \in \mathcal{R}, v, n \in \mathcal{V}, v \neq n \quad (15)$$

$$M\delta_i \geq \left(\sum_{r \in \mathcal{R}} \sum_{v \in \mathcal{V}} \sum_{n \in \mathcal{V}} \gamma_{ir}^{vn} \right), \quad \forall i \in \mathcal{H} \quad (16)$$

$$\mathcal{T}_i^v + \sum_{r \in \mathcal{R}} h^{vn} \gamma_{ir}^{vn} - \mathcal{T}'_{i^n} \leq M(1 - \delta_i), \quad \forall i \in \mathcal{H}, v, n \in \mathcal{V}, v \neq n \quad (17)$$

Time constraints. Constraints (18) establish the relationship between a vehicle's arrival and departure times at node $i \in \mathcal{N}$, accounting for operational activities performed at that location. Specifically, the time difference reflects the duration required for actions such as recharging/refueling and loading/unloading during transshipment. Constraints (19) ensure that if vehicle v travels from node $i \in \mathcal{N}$ to node $j \in \mathcal{N}$ via arc $(i, j, v) \in \mathcal{A}$, it arrives at node j after corresponding travel time θ_{ij}^v . Time window constraints are enforced by (20) and (21): the former prevents vehicle $v \in \mathcal{V}$ from picking up order $r \in \mathcal{R}$ before its earliest allowable pickup time at location P_r , and the latter ensures that the order is delivered to destination $d \in \mathcal{D}$ by both the order-specific deadline and the latest allowable time for that destination. For the time-related constraints (19) and (21), the Big- M constant is defined as the length of the planning horizon, providing an upper bound on the largest possible difference between arrival and departure times and ensuring proper relaxation when the associated binary routing variable equals zero.

$$\mathcal{T}'_i^v - \mathcal{T}_i^v \geq \begin{cases} 0, & \forall v \in \mathcal{V}, i \in \mathcal{N} \setminus (\mathcal{P}^v \cup \mathcal{H}) \\ g^v \mathcal{Z}_i^v + \sum_{r \in \mathcal{R}} h^{vn} (\gamma_{ir}^{vn} + \gamma_{ir}^{nv}), & \forall v, n \in \mathcal{V}, n \neq v, i \in (\mathcal{P}^v \cap \mathcal{H}) \\ \sum_{r \in \mathcal{R}} h^{vn} (\gamma_{ir}^{vn} + \gamma_{ir}^{nv}), & \forall v, n \in \mathcal{V}, n \neq v, i \in \mathcal{H} \setminus \mathcal{P}^v \\ g^v \mathcal{Z}_i^v, & \forall v \in \mathcal{V}, i \in \mathcal{P}^v \setminus \mathcal{H} \end{cases} \quad (18)$$

$$\mathcal{T}'_i^v + \theta_{ij}^v - \mathcal{T}_j^v \leq M(1 - \mathcal{X}_{ij}^v), \quad \forall (i, j, v) \in \mathcal{A} \quad (19)$$

$$a_{p_r} \leq \mathcal{T}'_{p_r}^v, \quad \forall r \in \mathcal{R}, v \in \mathcal{V} \quad (20)$$

$$\mathcal{T}_d^v \leq \min(b_r, l_d) \sum_{i:(i,d,v) \in \mathcal{A}} \mathcal{Y}_{idr}^v + M(1 - \sum_{i:(i,d,v) \in \mathcal{A}} \mathcal{Y}_{idr}^v), \quad \forall d \in \mathcal{D}, v \in \mathcal{V}, r \in \mathcal{R} \quad (21)$$

Energy constraints. Constraints (22) and (23) ensure energy sufficiency and operational safety for all vehicles. Constraints (22) enforce that a vehicle's energy upon departing any node does not exceed its maximum capacity, while constraints (23) ensure that its energy upon arrival remains above a predefined safety threshold. Constraints (24) initialize each vehicle with a full energy charge at its operational base. Constraints (25) allow recharging or refueling only at designated nodes $i \in \mathcal{P}^v$ if the vehicle reaches that node via an available arc. If recharging occurs, constraints (26) reset the vehicle's energy to full capacity. Conversely, constraints (27) maintain that energy remains unchanged between arrival and departure at a node if no recharging occurs. Finally, constraints (28) calculate the remaining energy upon arrival at node j as the departure energy from node i minus the energy required to traverse arc $(i, j, v) \in \mathcal{A}$. Constraints (27) and (28) employ a Big- M formulation to conditionally enforce the vehicle energy balance based on routing decisions. The parameter M is chosen as an upper bound on the maximum possible energy variation along any arc, equal to the vehicle's maximum energy capacity, which ensures the validity of the formulation without introducing unnecessary relaxation. The energy consumption function $\mathcal{E}_{ij}^v(\mathcal{T}'_i^v)$ is a parametric expression whose coefficients are precomputed to capture the impact of forecast weather conditions on the energy required for vehicle v to traverse arc (i, j, v) , when departing from location i at time $\mathcal{T}'_i^v$. Further details on energy consumption modeling are provided in Section 4.3.

$$E_i^v \leq E_{\max}^v, \quad \forall v \in \mathcal{V}, i \in \mathcal{N} \quad (22)$$

$$E_i^v \geq E_{\min}^v, \quad \forall v \in \mathcal{V}, i \in \mathcal{N} \quad (23)$$

$$E_{O^v}^v \geq E_{\max}^v \sum_{j:(O^v,j,v) \in \mathcal{A}} \mathcal{X}_{O^v j}^v, \quad \forall v \in \mathcal{V} \quad (24)$$

$$\mathcal{Z}_i^v \leq \sum_{j:(j,i,v) \in \mathcal{A}} \mathcal{X}_{ji}^v, \quad \forall v \in \mathcal{V}, i \in \mathcal{P}^v \quad (25)$$

$$E_i^v \geq E_{\max}^v \mathcal{Z}_i^v, \quad \forall v \in \mathcal{V}, i \in \mathcal{P}^v \quad (26)$$

$$E_i^v - E_i^v \leq \begin{cases} M\mathcal{Z}_i^v, & \forall v \in \mathcal{V}, i \in \mathcal{P}^v \\ 0, & \forall v \in \mathcal{V}, i \in \mathcal{N} \setminus \mathcal{P}^v \end{cases} \quad (27)$$

$$E_i^v - \mathcal{E}_{ij}^v(\mathcal{T}'_i^v) \geq E_j^v - M(1 - \mathcal{X}_{ij}^v), \quad \forall (i, j, v) \in \mathcal{A} \quad (28)$$

If the energy consumption function $\mathcal{E}_{ij}^v(\mathcal{T}_i^v)$ is a linear expression, the primary binary routing variables $\mathcal{X}_{ij}^v$ scale with $|\mathcal{A}|$, while order-flow variables $\mathcal{Y}_{ijr}^v$ scale with $|\mathcal{A}| \times |\mathcal{R}|$. Transshipment decisions introduce additional binary variables γ_{ir}^{vn} , whose cardinality grows with $|\mathcal{H}| \times |\mathcal{R}| \times |\mathcal{V}|^2$, where $\mathcal{H}$ denotes the set of candidate transshipment nodes. Time and energy tracking variables scale linearly with $|\mathcal{N}| \times |\mathcal{V}|$.

The number of constraints is dominated by routing feasibility, flow conservation, time propagation, energy balance, and transshipment consistency constraints. In particular, time and energy propagation constraints scale with $|\mathcal{A}|$, while transshipment feasibility constraints scale with $|\mathcal{H}| \times |\mathcal{R}| \times |\mathcal{V}|^2$. As a result, the overall model size grows polynomially with respect to the number of nodes, orders, and vehicles. Despite this polynomial model size, the problem remains NP-hard because it generalizes several well-known NP-hard problems, including the VRP, pickup-and-delivery with time windows, and joint routing, scheduling, and assignment problems.

Remark 1 (Symmetry handling and type-based specialization). The set $\mathcal{V}$ indexes individual vehicle resources, each defined by a base $\mathcal{O}^v$, a final-depot set $\mathcal{O}'^v$, and a feasible arc set obtained after preprocessing. When a subset of vehicles shares all of these attributes, its units are interchangeable and label symmetry may arise. Let $\mathcal{V}_h \subseteq \mathcal{V}$ be a maximal subset of mutually identical vehicles, sharing the same base, final-depot set, feasible arc set, and energy functions $\mathcal{E}_{ij}^v(\cdot)$ (which depend on the arc and platform rather than on vehicle identity), with members ordered arbitrarily as $v_1, v_2, \dots, v_{|\mathcal{V}_h|}$. Defining the activation indicator

$$u^v = \sum_{j: (\mathcal{O}^v, j, v) \in \mathcal{A}} \mathcal{X}_{\mathcal{O}^v j}^v, \quad (29)$$

symmetry can be suppressed by imposing the lexicographic activation-ordering constraints

$$u^{v_{m+1}} \leq u^{v_m}, \quad \forall \mathcal{V}_h, m = 1, \dots, |\mathcal{V}_h| - 1, \quad (30)$$

which require a lower-indexed identical vehicle to be activated before a higher-indexed one. When two identical vehicles are simultaneously active, symmetry can additionally be broken by ordering them by delivered load:

$$\sum_{\substack{(i, d, v_m) \in \mathcal{A} \\ d \in \mathcal{D}}} \sum_{r \in \mathcal{R}} Q_r \mathcal{Y}_{idr}^{v_m} \geq \sum_{\substack{(i, d, v_{m+1}) \in \mathcal{A} \\ d \in \mathcal{D}}} \sum_{r \in \mathcal{R}} Q_r \mathcal{Y}_{idr}^{v_{m+1}}, \quad \forall \mathcal{V}_h, m = 1, \dots, |\mathcal{V}_h| - 1. \quad (31)$$

Constraints (30) and (31) are valid symmetry-breaking constraints: any optimal solution can be relabeled within $\mathcal{V}_h$ to satisfy them, so at least one optimal solution is preserved. The individual-vehicle formulation also nests a type-based one: vehicles sharing a base, final-depot set, and arc set may be aggregated into a type k or type–location pair (k, s) with fleet-availability constraints, reducing symmetry in large homogeneous fleets. We retain the individual index v in this study because per-vehicle energy state (recharging and battery swapping) and pairwise transshipment do not aggregate cleanly under a count-based representation, while remaining compatible with the symmetry-reduction and aggregation strategies above.

4.3. Forecast-based spatiotemporal energy consumption modeling

Accurate modeling of drone energy consumption in operational planning requires linking routing decisions to spatially and temporally varying weather conditions, which are typically observed through forecasts available only at discrete locations and time intervals (e.g., hourly). This subsection describes how such forecast-based spatiotemporal weather data are systematically processed and embedded into the optimization model to represent time-dependent energy consumption across a geographically continuous service area.

We represent the service area as a directed graph $G(\mathcal{N}, \mathcal{A})$, where nodes $i, j \in \mathcal{N}$ are connected by multi-modal arcs $(i, j, v) \in \mathcal{A}$. To associate network elements with localized weather information, we employ Voronoi regions as a spatial mapping mechanism. Each node and arc segment is assigned to its nearest weather station, ensuring that energy consumption is estimated using the most relevant forecast data available in practice. Voronoi regions are used solely as a data-driven spatial interpolation mechanism to assign discrete meteorological observations to network elements for predictive energy modeling. They do not represent airspace boundaries, routing corridors, geofencing constraints, or operational restrictions, and therefore do not impose any constraints on routing or vehicle operations. This construction mirrors how weather information is accessed in practice and avoids introducing unnecessary modeling complexity or additional decision variables.

Formally, the service area is partitioned into Thiessen polygons (Voronoi diagrams) (Ruehle, 2020) based on the set of weather station locations $S = \{s_1, s_2, \dots, s_n\}$. Partitioning a D -dimensional space $\mathbb{R}^D$, the Voronoi region VR_i corresponding to station s_i is defined as:

$$\text{VR}_i = \{p \in \mathbb{R}^D \mid d(p, s_i) \leq d(p, s_j), \forall j \neq i\}, \quad (32)$$

where $d(p, s_i)$ denotes the Euclidean distance between location p and station s_i . The resulting regions are non-overlapping and collectively cover the entire service area.

To capture temporal variations in energy consumption due to changing weather conditions, the planning horizon T can be discretized into intervals $t = \{1, 2, \dots, |T|\}$ based on the frequency of available weather forecast data (e.g., hourly updates). This segmentation enables the integration of forecast-based, time-indexed weather data into energy consumption estimates, ensuring that each vehicle's energy usage reflects the prevailing environmental conditions in its assigned region.

For each arc $(i, j, v) \in \mathcal{A}$ and time interval t , energy consumption is computed using a mode-specific function F_{ijt}^v that accounts for vehicle characteristics, terrain, and localized weather conditions. If an arc lies entirely within a single Voronoi region, the corresponding regional forecast is used directly. If an arc crosses multiple regions, it is segmented at region boundaries and energy consumption is computed for each segment using the appropriate local forecast, with total arc energy obtained by aggregation.

The resulting arc- and time-specific energy estimates form a discrete series F_{ijt}^v for all $(i, j, v) \in \mathcal{A}$ and $t \in T$. To embed these estimates into the optimization model, we construct a continuous, departure-time-dependent energy function $\mathcal{E}_{ij}^v(\mathcal{T}_i^v)$ by fitting a prediction model to the preprocessed energy series. This function is treated as a fixed input to the optimization formulation but links energy consumption directly to the vehicle's departure-time decision variable, enabling forecast-based, time-dependent energy effects to be captured without introducing additional state variables or re-optimization. Details of the fitted prediction model and data sources used in this study are provided in [Section 6.3](#), and the full integration procedure is summarized in [Algorithm 1](#).

Algorithm 1 Weather-aware energy consumption integration.

- 1: **Initialize Service Area**
 - 2: Define the service area as a directed graph $G(\mathcal{N}, \mathcal{A})$, where $\mathcal{N}$ is the set of locations (nodes), and $\mathcal{A}$ is the set of multi-modal transportation arcs.
 - 3: **Partition Service Area into Voronoi Regions**
 - 4: Obtain weather station locations $S = \{s_1, s_2, \dots, s_n\}$ within G .
 - 5: Construct a Voronoi diagram using S as seed points.
 - 6: Assign each node $p \in \mathcal{N}$ to its closest weather station $s_i \in S$ based on Euclidean distance $d(p, s_i)$.
 - 7: Define Voronoi region $VR_i = \{p \in \mathcal{N} \mid d(p, s_i) \leq d(p, s_j), \forall s_j \in S, j \neq i\}$.
 - 8: **Discretize the Planning Horizon into Time Intervals**
 - 9: Determine the frequency of weather updates.
 - 10: Define the set of time intervals accordingly as $t \in T$.
 - 11: **Calculate Arc-Based Energy Consumption**
 - 12: **for** all arcs $(i, j, v) \in \mathcal{A}$
 - 13: **for** all time intervals $t \in T$
 - 14: Compute energy consumption function F_{ijt}^v considering: Vehicle type, weather-adjusted data, energy consumption formulation, and terrain conditions.
 - 15: **endfor**
 - 16: Update and store energy values.
 - 17: Fit a prediction function to estimate future arc energy consumption using the energy consumption time series F_{ijt}^v .
 - 18: Utilize the fitted function for arc energy consumption $\mathcal{E}_{ij}^v(\mathcal{T}_i^v)$ as a function of departure time.
 - 19: **endfor**
-

5. Solver-based solution methodology

The proposed multi-modal routing and scheduling model can become computationally challenging even for moderate-sized rural networks due to heterogeneous vehicle types, time-dependent energy consumption, and transshipment decisions in geographically dispersed service areas. To improve tractability without altering the core optimization formulation, we introduce a set of preprocessing, tractability-enhancement, and model-control mechanisms. These methodological components are required to translate the theoretical formulation in [Section 4](#) into a practically implementable solution framework.

5.1. Handling multi-role nodes via dummy final depots

In multi-modal delivery networks, a single physical location may serve multiple operational roles, such as a pickup location, operational base, transshipment point, recharging/refueling station, or final depot (see [Fig. 1](#)). While such role sharing reflects operational reality, it can introduce structural challenges in routing formulations, particularly in the presence of sub-tour elimination constraints.

If these roles are not modeled carefully, logical inconsistencies may arise. For example, a node that simultaneously functions as a final depot for one vehicle and as a transshipment point for another may be incorrectly treated as a terminal node, leading to infeasible route terminations or unintended routing loops.

To address this issue, we apply a preprocessing step that introduces dummy nodes corresponding exclusively to final-depot functionality. For each candidate final depot, a dummy node is created and connected to its original node via a zero-distance, zero-time arc. The dummy node serves solely as a terminal node and has no outgoing arcs. This construction preserves the flexibility of original nodes to serve multiple operational roles, while ensuring that sub-tour elimination constraints [\(6\)–\(8\)](#) are enforced consistently, thereby maintaining routing feasibility and network connectivity.

Remark 2. To ensure correct modeling of nodes with multiple roles, we define a set $\mathcal{O}^v$ that includes a dummy node l for each candidate final depot $i \in \mathcal{N}$. A directed dummy arc $(i, l, v) \in \mathcal{A}$ is added to represent one-way movement from the original node to its dummy counterpart, which has no outgoing arcs and is exempt from sub-tour elimination constraints. Dummy arcs incur zero time,

cost, and energy, serving purely as structural elements. This construction preserves routing feasibility and allows final depot selection without interfering with other operational roles of the original node.

5.2. Energy- and command-and-control (C2)-feasibility arc pruning

To improve computational tractability while preserving operational realism, we preprocess the aerial network by pruning drone arcs that are infeasible under forecast-based energy requirements and C2 connectivity conditions. Incorporating C2 feasibility reflects regulatory and operational requirements for safe UAS operations, particularly beyond-visual-line-of-sight (BVLOS) missions, where reliable communication with ground control infrastructure is essential (Barlow et al., 2019; Fontanesi et al., 2023). Explicitly modeling communication performance (e.g., signal strength, latency, or link reliability) or dynamically re-optimizing trajectories would substantially increase model complexity and exceed the scope of the day-ahead planning framework considered in this study; therefore, C2 feasibility is enforced through preprocessing rather than explicit MILP constraints. This approach removes structurally infeasible arcs before optimization, allowing the model to focus on routing and logistics decisions within an operationally feasible aerial network.

Battery feasibility under forecast-based weather: The departure-time-dependent energy consumption function $\mathcal{E}_{ij}^v(\mathcal{T}_i^{lv})$ is considered for each drone arc $(i, j, v) \in \mathcal{A}$. An arc is retained if there exists at least one feasible departure time within the planning horizon for which the required energy does not exceed the drone's battery capacity:

$$(i, j, v) \in \mathcal{A} \iff \exists \mathcal{T}_i^{lv} \in [0, T] \text{ such that } \mathcal{E}_{ij}^v(\mathcal{T}_i^{lv}) \leq E_{\max}^v, \quad \forall v \in \mathcal{V}. \quad (33)$$

Arcs that violate this condition for all feasible departure times are removed prior to optimization.

C2 connectivity feasibility: Drone arc $(i, j, v) \in \mathcal{A}$ is discretized into a finite set of K sampled locations $p_{ijv} = \{p_{ijv}^1, p_{ijv}^2, \dots, p_{ijv}^K\}$, used for connectivity assessment based on publicly available mobile broadband coverage data (e.g., 4G/5G LTE). Because such coverage information is available only at a discrete spatial resolution, the sampled locations provide a tractable approximation of C2 feasibility along the arc. They do not represent routing decisions, trajectory discretization, or control waypoints. Let $C(p_{ijv}^k)$ denote a binary C2 availability indicator at location p_{ijv}^k , defined as

$$C(p_{ijv}^k) = \begin{cases} 1, & \text{if } p_{ijv}^k \text{ lies within the broadband coverage area,} \\ 0, & \text{otherwise.} \end{cases} \quad (34)$$

The C2 coverage fraction of drone arc (i, j, v) is then computed as $\phi_{ijv} = \frac{1}{K} \sum_{k=1}^K C(p_{ijv}^k)$, with $\phi_{ijv} \in [0, 1]$. Given minimum coverage threshold $\tau \in (0, 1]$, drone arc (i, j, v) is considered C2-feasible if $\phi_{ijv} \geq \tau$. Arcs that do not satisfy this condition are removed from the feasible arc set prior to optimization.

5.3. Arc-dominance rules for multi-modal network reduction

The application of arc-dominance rules is used as a preprocessing mechanism to improve computational efficiency by reducing the size of the multi-modal network prior to optimization. Through data-driven screening, arc-mode combinations that are provably suboptimal are identified and eliminated based on simple dominance criteria. This reduction preserves optimality while substantially decreasing the number of routing variables, thereby enhancing the tractability of the proposed optimization model. The following proposition formalizes a sufficient condition under which one arc-mode combination dominates another in both travel time and transportation cost.

Proposition 1. Consider two vehicles of different transportation modes, denoted by $a \in \mathcal{V}$ and $b \in \mathcal{V}$, traversing from node i to node j via arcs (i, j, a) and $(i, j, b) \in \mathcal{A}$, respectively. Let their speeds be V^a and V^b , and costs per mile be κ^a and κ^b . Let d_{ij}^a and d_{ij}^b represent the effective travel distances for modes a and b , respectively (e.g., straight-line for aerial and road distance for ground).

Define the following relative efficiency metrics:

$$CF_{ij} = \frac{d_{ij}^b}{d_{ij}^a}, \quad R = \frac{V^b}{V^a}, \quad K = \frac{\kappa^b}{\kappa^a}.$$

Assume $R \geq 1$ and $K \leq 1$; that is, vehicle b is at least as fast and less costly per unit distance than vehicle a . Then, the arc (i, j, a) is considered dominated and can be removed from the network if both of the following conditions hold:

$$CF_{ij} < R \quad \text{and} \quad CF_{ij} < \frac{1}{K}.$$

Proof. Let the travel times be $\theta_{ij}^a = \frac{d_{ij}^a}{V^a}$ and $\theta_{ij}^b = \frac{d_{ij}^b}{V^b}$, and travel costs be $C_{ij}^a = \kappa^a \cdot d_{ij}^a$ and $C_{ij}^b = \kappa^b \cdot d_{ij}^b$.

(i) *Time Efficiency Condition:* If $CF_{ij} < R$, then:

$$\frac{d_{ij}^b}{d_{ij}^a} < \frac{V^b}{V^a} \quad \Rightarrow \quad \frac{d_{ij}^b}{V^b} < \frac{d_{ij}^a}{V^a} \quad \Rightarrow \quad \theta_{ij}^b < \theta_{ij}^a$$

Thus, vehicle b is faster.

(ii) *Cost Efficiency Condition*: If $CF_{ij} < \frac{1}{K}$, then:

$$\frac{d_{ij}^b}{d_{ij}^a} < \frac{1}{\frac{\kappa^b}{\kappa^a}} = \frac{\kappa^a}{\kappa^b} \Rightarrow \kappa^b \cdot d_{ij}^b < \kappa^a \cdot d_{ij}^a \Rightarrow C_{ij}^b < C_{ij}^a$$

Thus, vehicle b is cheaper.

Since vehicle b is both faster and cheaper under the stated conditions, the arc (i, j, a) is dominated and can be removed without affecting optimality. $\square$

This dominance property is particularly relevant in multi-modal delivery networks where one transportation mode, often ground vehicles, may outperform another mode, such as drones, on certain origin-destination pairs in both time and cost. Arc-mode combinations satisfying the dominance conditions can therefore be systematically removed during preprocessing. This reduction decreases the number of decision variables and constraints without excluding any potentially optimal solutions, improving computational efficiency while preserving model fidelity.

5.4. Model-level controls for waiting-time management

When energy consumption is modeled as a function of vehicles' departure time, the optimization may exhibit unintended waiting behavior at intermediate nodes. In particular, if the predictive energy function decreases with later departures, the model may favor delayed departures to reduce energy use. While this behavior is mathematically consistent with the objective, it may be operationally unrealistic in time-sensitive healthcare delivery settings. The following proposition formalizes this phenomenon and motivates the introduction of modeling controls.

Proposition 2. Let the energy consumption for a vehicle $v \in \mathcal{V}$ traversing arc $(i, j, v) \in \mathcal{A}$ be an affine function of its departure time $\mathcal{T}'_i{}^v$ from node i , given by:

$$\mathcal{E}_{ij}^v(\mathcal{T}'_i{}^v) = a\mathcal{T}'_i{}^v + b,$$

where $a \in \mathbb{R}$ and $b \in \mathbb{R}$ are fixed parameters. Let the latest feasible departure time be:

$$T_{ij}^{\max} = \min(b_r, l_d) - \tau_{jd}^{\text{acc}},$$

where b_r is the latest delivery time for order $r \in \mathcal{R}$, l_d is the latest allowable delivery time at destination $d \in \mathcal{D}$, and τ_{jd}^{acc} is the accumulated travel time from node j to node d .

Assume that the vehicle must depart from node i no earlier than a known lower bound $T_{ij}^{\min}$ and no later than $T_{ij}^{\max}$ to ensure timely delivery. Then, the optimal departure time ($\mathcal{T}'_i{}^{v*}$) that minimizes energy consumption is given by:

$$\mathcal{T}'_i{}^{v*} = T_{ij}^{\max} \text{ if } a \leq 0, \quad \text{else } T_{ij}^{\min}$$

Proof. The proposed model aims to minimize the affine energy function $\mathcal{E}_{ij}^v(\mathcal{T}'_i{}^v) = a\mathcal{T}'_i{}^v + b$ over the feasible interval $\mathcal{T}'_i{}^v \in [T_{ij}^{\min}, T_{ij}^{\max}]$. Since the function is linear, its minimum occurs at one of the two interval boundaries. If $a > 0$, the function is increasing, and the minimum is attained at $T_{ij}^{\min}$. If $a < 0$, the function is decreasing, and the minimum is attained at $T_{ij}^{\max}$. If $a = 0$, the function is constant, and any value in the interval yields the same energy; without loss of generality, we choose $T_{ij}^{\max}$ to preserve scheduling flexibility. Thus, the optimal departure time depends on the sign of a . $\square$

This reveals a key modeling limitation when incorporating a predictive energy consumption function as a linear function of departure time. Specifically, when the slope $a \leq 0$, the model may introduce excessive and impractical waiting at nodes to minimize energy use, especially when delivery time windows are not tight. While mathematically optimal, such behavior may conflict with operational constraints, highlighting the need to regulate or penalize waiting to ensure realistic and implementable solutions. To address this, we propose two potential solutions:

(1) Adding a time penalty to the objective function: To discourage unnecessary waiting driven by energy minimization, we add a penalty term $\sum_{v \in \mathcal{V}} \sum_{d \in \mathcal{D}} \epsilon \mathcal{T}_d^v$ to the original objective function (1). Here, ϵ represents a small penalty coefficient associated with later arrival times. This term incentivizes earlier departures when feasible, balancing energy efficiency and practical considerations, such as service continuity and schedule adherence.

(2) Defining a limit for waiting times: Alternatively, we can explicitly constrain waiting times by defining a maximum allowable waiting threshold for node i , denoted by Ω_i . Also, this waiting time cannot exceed the remaining time from arrival to the latest allowable delivery time. These constraints are formulated as:

$$\mathcal{T}'_i{}^v - \mathcal{T}_i^v - \Omega_i \leq \begin{cases} 0, & \forall v \in \mathcal{V}, i \in \mathcal{N} \setminus (\mathcal{P}^v \cup \mathcal{H}) \\ g^v \mathcal{Z}_i^v + \sum_{r \in \mathcal{R}} h^{vn}(\gamma_{ir}^{vn} + \gamma_{ir}^{nv}), & \forall v, n \in \mathcal{V}, n \neq v, i \in (\mathcal{P}^v \cap \mathcal{H}) \\ \sum_{r \in \mathcal{R}} h^{vn}(\gamma_{ir}^{vn} + \gamma_{ir}^{nv}), & \forall v, n \in \mathcal{V}, n \neq v, i \in \mathcal{H} \setminus \mathcal{P}^v \\ g^v \mathcal{Z}_i^v, & \forall v \in \mathcal{V}, i \in \mathcal{P}^v \setminus \mathcal{H} \end{cases} \quad (35)$$

Table 5
Overview of four operating scenarios by volume and wind conditions.

Operating Scenario	Feature	Operation Date	Number of Pickup Locations	Total Order Weight (lb)
LV-LW	Low Volume - Low Wind	06/27/22	13	6.5
LV-HW	Low Volume - High Wind	02/21/22	7	12.1
HV-LW	High Volume - Low Wind	02/14/22	37	53
HV-HW	High Volume - High Wind	01/24/22	51	40.5

$$\Omega_i \leq \min(b_r, l_i) \sum_{j:(j,i,v) \in A} \mathcal{Y}_{jir}^v + M(1 - \sum_{j:(j,i,v) \in A} \mathcal{Y}_{jir}^v) - \mathcal{T}_i^v, \quad \forall i \in D, v \in \mathcal{V}, r \in R \quad (36)$$

Together, these controls ensure that the optimization model balances energy efficiency with operational realism, without compromising feasibility or significantly increasing computational complexity.

6. Case study

This section presents a case study based on Shield Illinois, a statewide COVID-19 testing initiative with a large-scale logistics network. We use this real-world context to evaluate our proposed multi-modal medical supply network. We begin by describing the decentralized point-to-point delivery system implemented during the Shield Illinois program and present performance metrics from four status quo scenarios to highlight key operational patterns. Next, we outline the development of the case-study network, input parameters, and operating scenarios used to evaluate the proposed optimization model. Finally, we detail the integration of drones into the model, including the estimation of energy consumption influenced by wind conditions.

6.1. Decentralized delivery system of shield illinois

Shield Illinois was a statewide initiative launched by the University of Illinois System to provide rapid, saliva-based COVID-19 testing during the pandemic. From July 2020 to June 2023, the program administered more than 7 million tests across 1,779 K-12 schools, 57 colleges and universities, and 77 community-based sites. Supported by a logistics network spanning a 400-by-200-mile service area and 13 diagnostic laboratories, Shield Illinois consistently achieved an average turnaround time of 16 h, processing between 10,000 and 14,000 samples per day during peak demand periods. For our case study, we focus on rural southern Illinois, where logistical constraints align with the assumptions and structure of our proposed model.

The Shield Illinois logistics system operated as a decentralized, point-to-point delivery network using ground vehicles, primarily cars and SUVs. Each of the 2,000+ testing sites arranged transportation independently to deliver collected samples to designated laboratories for Polymerase Chain Reaction (PCR) processing. While this approach enabled broad access to testing, it also led to operational inefficiencies, particularly in rural areas, where the lack of route optimization and vehicle pooling contributed to delays in diagnostics and public health responses.

To evaluate logistics performance under varying operational conditions, we select four representative days, each serving as an individual operating scenario in our single-period planning framework, where one period corresponds to a single day, and capturing distinct combinations of order volume (high vs. low) and wind conditions (calm vs. windy). While the model can accommodate a range of weather factors, we focus specifically on wind, as it is the most critical environmental variable influencing aerial operations in the rural geography of southern Illinois. Order volume serves as a proxy for system congestion, while wind intensity affects vehicle energy consumption and delivery time reliability. The selected instances are based on actual rural operational data from southern Illinois and provide a structured basis for performance benchmarking and comparison with our proposed optimization framework. Table 5 summarizes the operating scenario characteristics.

To evaluate each instance, we assess total system costs, distance traveled, carbon emissions, weighted average delivery time, and delivery performance across time thresholds of 1, 5, 10, and 15 h (representing same-day delivery), and beyond 15 h (classified as next-day delivery). The total system cost is estimated based on transportation time, using a marginal delivery cost of \$43.5 per hour per car, to account for fuel and labor. Weighted average delivery time is computed by weighting each order's delivery time by its volume, while emission estimates are derived from Rodrigues et al. (2022), using 339.8 g/mile for ground vehicles and 14.4 g/mile for drones.

Table 6 summarizes the performance of the decentralized ground-based delivery system across four baseline instances (B1–B4) representing the status quo operational practice. These instances serve as the baseline for benchmarking the proposed optimization framework, which enables centralized and coordinated multi-modal planning (see Section 7). While the system achieved reasonable coverage within the 10- to 15-h delivery window, it frequently failed to satisfy more stringent same-day delivery targets, particularly within the 1- and 5-h thresholds, highlighting its limited responsiveness. For example, in the low-volume but high-wind scenario (LV-HW, B2), only 37.7% of orders were delivered within 1 h, despite relatively low congestion. Performance deteriorated further in high-volume (e.g., HV-HW, B4) and geographically dispersed scenarios (e.g., LV-LW, B1), where no orders met the 1-h target and fewer than 53% were delivered within 5 h. These shortcomings are amplified in settings with dispersed pickup locations and higher demand, resulting in increased transportation costs, longer travel distances, higher emissions, and failure to consistently meet same-day delivery requirements.

Table 6

Performance summary of implemented ground-based delivery system.

Operating Scenario	Instance	Total Costs (\$)	Total Distance Traveled (mile)	CO ₂ Emissions (g)	Weighted Average Delivery Time (hr)	Orders delivered within 1 h (%)	Orders delivered within 5 h (%)	Orders delivered within 10 h (%)	Orders delivered within 15 h (%)	Orders Delivered Next Day (%)
LV-LW	B1	777	984	6726	6.5	0%	48.21%	93.87%	93.87%	6.13%
LV-HW	B2	329	390	2666	2.81	37.68%	85.22%	98.45%	100%	0%
HV-LW	B3	1682	2026	13848	5.75	18.37%	51.08%	90.93%	95.16%	4.84%
HV-HW	B4	2449	2968	20286	5.48	0%	52.74%	93.41%	99.24%	0.76%

6.2. Network design and operational parameters

For each operating scenario, the transportation network is constructed using the actual geographic coordinates of the pickup sites and a central diagnostic laboratory, which together define the set of network nodes. In the model, selected nodes serve as initial depots and candidate final depots for each vehicle type, as well as candidate locations for recharging or transshipment operations. Ground arcs connect every pair of nodes based on real-world driving distances and travel times obtained from Google Maps, thereby capturing realistic road network conditions. Drone arcs represent direct aerial travel between nodes and are defined using Euclidean distances and constant-speed cruising times.

C2 availability is evaluated using publicly available mobile broadband coverage data from the FCC Broadband Data Collection (Federal Communications Commission, 2025). These coverage classifications are derived from provider-reported propagation models and correspond to standard LTE service definitions and minimum received signal reference power (RSRP) thresholds used to indicate baseline communication availability (3GP, 2023). For each drone arc, connectivity is screened by sampling points along the geodesic path between its endpoints at 700-meter (just under half a mile) intervals, including both endpoints. A strict threshold of $\tau = 1$ is adopted for C2 feasibility, requiring full LTE coverage at all sampled locations. This conservative choice reflects baseline BVLOS safety considerations in the study area and avoids reliance on partial or intermittent connectivity when constructing the feasible drone subnetwork. Spatial visualizations of the resulting C2-feasible and infeasible drone arcs are provided in Appendix A.

We assume a marginal operational cost of \$43.5 per hour for ground vehicles and \$30 per hour for drones, reflecting typical expenses associated with fuel, labor, and equipment usage, based on cost estimates reported in Grote et al. (2024). Ground vehicles are assumed to have unlimited capacity, while drones are limited to carrying up to 15 pounds per trip. Each sample is assumed to weigh 5 g (0.011 pounds). All orders are available for pickup at 8:00 AM and must be delivered by 8:00 PM the same day. Additional operational costs include \$2 per drone recharge and \$1 per transshipment event, regardless of the number of orders handled. These values reflect nominal handling and infrastructure costs, scaled to reflect their relatively minor contribution to overall operational expenses. We assume that each drone recharge/battery-swapping operation takes 10 min, consistent with current rapid-charging capabilities, while each transshipment operation requires 5 min per order to account for basic handling and coordination time.

Fig. 2 illustrates the spatial structure of the logistics network for each operating scenario, highlighting the geographic distribution of pickup nodes around the central diagnostic laboratory, which serves as the destination. The service area includes six weather stations and corresponding Voronoi regions (VRs); additional details are provided in Section 6.3. Instances B1, B3, and B4 exhibit high spatial dispersion, with pickup locations spread across five to six VRs. In contrast, instance B2 shows a more geographically concentrated pattern, with all pickup sites located within three neighboring regions, clustered south of the central diagnostic laboratory. Notably, B1 remains dispersed despite its low order volume, while B3 and B4 combine a wide geographic spread with higher demand levels.

6.3. Drone type and energy consumption estimates

This subsection specifies the representative drone characteristics and the physical abstraction level used to estimate energy consumption in the case study. The model targets operational routing and scheduling decisions over a single planning period (i.e., a day), rather than detailed flight dynamics. In particular, this study assumes the use of long-range drones capable of carrying up to 15 pounds of medical supplies, with a maximum operational range of 100 miles and a constant ground speed of $V_g = 62$ miles per hour (for details on the drone platform, we refer the reader to Appendix B). To estimate energy consumption, we focus on steady, level flight and incorporate hourly wind speed and direction data from six weather stations distributed across the service area. Meteorological data are collected from multiple sources, as detailed in Appendix C.

Drone energy consumption on each arc is considered by estimating the time-dependent function F_{ijt}^v , following the procedure outlined in Algorithm 1 in Section 4.3. This function is derived from the aerodynamic model proposed by Zeng and Zhang (2017), which estimates power requirements based on aerodynamic drag, propulsion efficiency, and wind influence. We assume that drones are fully loaded throughout each flight segment, yielding a conservative upper-bound estimate of energy consumption. The total energy required to traverse an arc is given by:

$$F_{ijt}^v = T \cdot \left(c_1 V^3 + \frac{c_2}{V} \right) \cdot (1 + \rho_{cw}), \quad (37)$$




Fig. 2. Geographic distribution of pickup and delivery locations with VR boundaries across the case studies.

where F_{ijt}^v is the total energy consumed (joules) during flight on the arc (i, j, v) , $T = \frac{d_{ij}^v}{V_g}$ is the time to traverse on this arc with the length of d_{ij}^v (mile), V is the effective airspeed adjusted for wind conditions, ρ_{cw} is the crosswind penalty factor, and the constants $c_1 = 9.26 \times 10^{-4}$ and $c_2 = 2250$ represent aerodynamic drag and propulsion efficiency, respectively. Eq. (37) reflects energy consumption during cruise-dominated flight, which constitutes the primary component of each mission segment between pickup and delivery locations in rural settings. Energy expenditures associated with takeoff, climb, descent, and landing are incorporated implicitly through conservative parameterization and full-payload assumptions. Highly dynamic flight regimes and aggressive maneuvers are not considered, as medical delivery missions are assumed to follow stable, waypoint-based flight profiles governed by onboard flight-control systems.

To capture wind effects and estimate effective airspeed V , we project the wind vector onto the direction of each arc. Let θ denote the angle between the wind direction and the arc bearing. The wind vector is decomposed into longitudinal component: $W \cos(\theta)$, and crosswind component: $W \sin(\theta)$, where W is the wind speed (mph). The effective airspeed V is then adjusted as follows:


$$V = \begin{cases} V_g + |W \cos(\theta)| & \text{if } \theta > 90^\circ \text{ (headwind)} \\ V_g - |W \cos(\theta)| & \text{if } \theta \leq 90^\circ \text{ (tailwind)} \end{cases}$$

To account for the energy required to stabilize the drone against lateral drift, we define the positive crosswind penalty factor as $\rho_{cw} = \frac{|W \sin(\theta)|}{W}$. This simplifies directional wind impacts without adding computational complexity. This formulation realistically

Table 7

Best-fit time intervals for predicting drone energy consumption in each case study.

Case Study	Best R^2 Interval, Mean R^2	Best RMSE Interval, Mean RMSE	Best Balanced Interval, Score
HV-LW	(6:00-18:00), 0.39	(6:00-18:00), 142.45	(6:00-18:00), 2.00
HV-HW	(4:00-18:00), 0.58	(9:00-18:00), 119.78	(9:00-18:00), 1.80
LV-HW	(2:00-18:00), 0.30	(7:00-19:00), 92.55	(7:00-19:00), 1.84
LV-LW	(9:00-24:00), 0.35	(8:00-18:00), 143.72	(9:00-21:00), 1.59

**Fig. 3.** Linear regression models of energy consumption over time for two representative arcs.

incorporates both longitudinal and lateral wind effects, allowing for spatially and temporally accurate estimation of drone energy use under varying wind conditions.

To capture temporal variability in weather, we estimate the hourly energy consumption values F_{ijt}^v for each drone arc by adjusting effective wind speed over a 24-h period. This produces a discrete time series of forecast-based energy consumption values, referred to as the arc's energy profile. The weather forecasts are sampled at discrete time intervals, whereas departure times are continuous decision variables. Therefore, the discrete energy profile is approximated exogenously, prior to optimization, using a continuous departure-time-dependent prediction function. In this case study, each arc-specific energy profile is represented by the linear regression function $\mathcal{E}_{ij}^v(\mathcal{T}'_i{}^v) = \alpha_{ij}^v \mathcal{T}'_i{}^v + \beta_{ij}^v$, where α_{ij}^v and β_{ij}^v are estimated before optimization and subsequently treated as fixed input parameters. The optimization model therefore evaluates the fitted function at the continuous departure time selected by the model rather than directly optimizing over discrete weather intervals. Since the regression coefficients are fixed, the energy-consumption expression remains affine in the continuous departure-time variable and enters the energy-balance constraints (Eq. (28)) without introducing nonlinearities or products of decision variables. Consequently, the implemented case-study model remains an MILP. The fitted energy functions preserve the weather-aware nature of the framework because they are derived from forecast-based energy estimates that explicitly account for localized wind speed, wind direction, arc bearing, effective airspeed, and crosswind effects.

To ensure operational relevance, the fitting process always includes the core delivery window from 8:00 AM to 6:00 PM. For each case study, we evaluated multiple time intervals extending beyond this window and assessed model performance using the coefficient of determination (R^2) and root mean square error (RMSE), as summarized in Table 7. The most representative linear model $\mathcal{E}_{ij}^v(\mathcal{T}'_i{}^v)$ is selected based on the best trade-off between R^2 and RMSE. Although this case study uses a linear regression representation, the proposed framework is not restricted to this functional form. More generally, any exogenously estimated departure-time-dependent energy function may be incorporated. Alternative approximation methods, such as simulation-based fitted functions or piecewise-linear representations, could also be used when greater fidelity is desired. Fig. 3 illustrates examples of fitted regression lines for two arcs with positive and negative slopes. A positive slope indicates increasing energy consumption over time due to intensifying headwinds, while a negative slope reflects improving wind conditions (i.e., stronger tailwinds).

7. Computational results

This section reports the computational experiments used to evaluate the operational performance and decision-support value of the proposed optimization framework in a rural last-mile delivery setting. Using the case study described in Section 6, we analyze cost efficiency, delivery performance, and environmental impacts across different operating scenarios. We first define the experimental categories and summarize the main performance metrics. We then assess the benefits of weather-aware modeling, followed by a sensitivity analysis that examines the isolated effects of key operational parameters on system performance and mode selection. In addition, this section evaluates the computational impact of the proposed preprocessing and tractability-enhancement procedures through a benchmark comparison between the enhanced implementation and an unreduced-network baseline. The section concludes with managerial insights derived from these results.

Table 8

Operating scenario instances categorized by logistics strategy (ground-vehicle-only, drone-only, or multi-modal, as defined by the available fleet), network size, and order volume.

Operating Scenario	Instance	Available Fleet	Number of Pickup Locations	Total Order Weight (lb)	Total Number of Arcs
LV–LW	INS1	1 Car	13	6.5	196
	INS2	1 Drone			138
	INS3	1 Drone, 1 Car			334
LV–HW	INS4	1 Car	7	12.1	64
	INS5	1 Drone			64
	INS6	1 Drone, 1 Car			128
HV–LW	INS7	1 Car	37	53	316
	INS8	4 Drones			652
	INS9	4 Drones, 1 Car			997
HV–HW	INS10	1 Car	51	40.5	436
	INS11	3 Drones			1201
	INS12	3 Drones, 1 Car			1637

Note: The listed vehicles define the available fleet for each representative feasible scenario. The optimization model determines which of these vehicles are activated and how they are used.

7.1. Experimental instances and optimal configurations

We evaluate model performance across a range of logistics strategies by generating multiple instances for each operating scenario described in Section 6.1, grouped into three main categories based on their operational characteristics. Table 8 summarizes key attributes of each instance, including the logistics strategy based on available fleet, number of pickup locations, total order weight, and network size, measured by the total number of arcs.

The fleet configurations reported in Table 8 represent the feasible experimental scenarios considered for evaluating the three logistics strategies. For each operating scenario, the listed vehicles define the resources made available to the model so that the corresponding ground-vehicle-only, drone-only, or multi-modal strategy can be evaluated. These configurations are not intended to represent strategic fleet-size recommendations. Given each available fleet, the optimization model determines the vehicles that are activated and their associated routing, scheduling, recharging, and transshipment decisions. The effect of varying the number of available cars and drones is examined separately through the fleet-size sensitivity analysis presented in Section 7.4.5.

Logistics Strategy 1: Ground Vehicle-Only Network (No Drones). This category of instances evaluates network performance using only ground transportation, with a single car (typically an SUV) available to handle pickups and deliveries, without drones or transshipment. The single ground vehicle is specified as the available fleet limit for the main ground-vehicle-only scenario, reflecting a limited-resource operational planning setting. This category includes instances INS1, INS4, INS7, and INS10, and serves as a benchmark for assessing the value of shared resources, the potential benefit of using drones, and optimization enhancements. The model minimizes transportation and potential recharging costs, subject to constraints (2)–(14) and (18)–(28). For consistency across instances, the car is assumed to start its route at the pickup location farthest from the destination. This design choice yields efficient baseline routes and intentionally favors the ground-vehicle benchmark, thereby providing a conservative and realistic basis for comparison.

Logistics Strategy 2: Drone-Only, Weather-Aware Network. This configuration evaluates system performance using only drones, providing insight into the capabilities and limitations of fully aerial delivery under time-dependent weather conditions (INS2, INS5, INS8, and INS11). As described above, the available drone fleet in each instance is a resource-constrained scenario input. The number made available is bounded below by payload feasibility—given the 15 lb payload, at least four drones are required for INS8 and three for INS11 for the demand to be servable at all, while a single drone suffices for the lighter INS2 and INS5 loads—but this bound is only a feasibility floor, not the operative fleet size. Within the available limit, the optimization model endogenously determines how many drones are actually activated and how they are routed and scheduled, with utilization depending on routing, sequential service opportunities, battery limitations, recharging, energy feasibility, delivery time windows, and network structure rather than payload alone. The effect of relaxing these available-fleet limits is examined directly in the fleet-size sensitivity analysis of Section 7.4.5.

Logistics Strategy 3: Multi-Modal, Weather-Aware Network. This category represents an integrated multi-modal configuration constructed by combining the ground-vehicle and drone subnetworks used in Categories 1 and 2. For each instance (INS3, INS6, INS9, and INS12), the set of nodes, arcs, vehicles, and operational constraints is defined as the union of the corresponding car-only and drone-only benchmarks, with no additional infrastructure, vehicles, or routing options introduced, ensuring a direct, like-for-like comparison across all categories. Within this integrated feasible network, the model determines whether ground vehicles, drones, or their coordination through transshipment and recharging operations provides the lowest-cost feasible solution for the specified available fleet.

Fig. 4 presents the optimal routes and operational decisions for the multi-modal instances in Logistics Strategy 3. Although both transportation modes are available in these instances, the optimization model endogenously determines which vehicles to activate and may select either a single-mode or coordinated multi-modal solution. The results highlight the critical role of spatial structure in




Fig. 4. Optimal delivery routes with recharging and transshipment operations for multi-modal instances (logistics strategy 3).

determining the optimal network configuration. In instances with concentrated pickup locations, such as INS6 (LV-HW), where all pickup sites are located south of the delivery location, a single drone with sufficient capacity can serve all orders efficiently, making multi-modal coordination unnecessary. In contrast, in spatially dispersed networks such as INS3 (LV-LW), multi-modal solutions that combine ground vehicles and drones become advantageous. In these operating scenarios, transshipment and intermediate recharging enable drones to extend their effective service range and reduce total travel distance, improving system performance even at low demand levels. Fig. 4c and d further illustrate that, in high-volume operating scenarios, the availability of multiple vehicles can improve network efficiency by enabling parallel service. In these instances, drone-only configurations emerge as the optimal solution, as recharging allows multiple sorties within the planning horizon and the available drone fleet is sufficient to handle the order volume.

7.2. Numerical results and performance measures

This section presents computational results for the twelve experimental instances in Table 8, covering variations in demand volume, wind conditions, and logistics strategies, including ground-vehicle-only, drone-only, and multi-modal configurations. The analysis examines how routing structure, available fleet configuration, spatial characteristics, and weather-aware energy modeling jointly affect system performance.

Table 9 reports the key performance metrics for all optimized instances, including total system cost, weighted average delivery time, total distance traveled, CO₂ emissions, and solver run time. We applied all tractability-enhancement procedures proposed in Section 5 to all instances. Most instances are solved exactly to optimality, and only minor optimality gaps are observed in a small

Table 9
Comparison of performance metrics across all instances.

Operating Scenario	Instance	Total Costs (\$)	Total Distance Traveled (mile)	CO ₂ Emissions (g)	Weighted Average Delivery Time (h)	Solver Run Time (s) (Gap)
LV-LW	INS1	451	584	3991	10.4	0.2 (0.00)
	INS2	259	502	24.95	9.43	24.61 (0.00)
	INS3	258	474	464	7.7	10.37 (0.00)
LV-HW	INS4	154	157	1073	3.6	0.1 (0.00)
	INS5	68	132	6.6	2.1	0.1 (0.00)
	INS6	68	132	6.6	2.1	0.1 (0.00)
HV-LW	INS7	444	416	2843	10.2	11.9 (0.00)
	INS8	170	343	17.1	2.46	10,800 (0.029)
	INS9	170	343	17.1	2.46	10,800 (0.015)
HV-HW	INS10	437	623	4258	10.03	2.47 (0.00)
	INS11	261	520	25.84	2.98	10,800 (0.0085)
	INS12	261	520	25.84	2.98	10,800 (0.014)

Note: A time limit of 3 h (10,800 s) was imposed; instances not reaching proven optimality within this limit were terminated.

number of more challenging instances. Given the day-ahead planning setting, solution times on the order of hours are operationally acceptable and do not require real-time computation. The computational role of these preprocessing strategies is evaluated separately in Section 7.4.6, where the enhanced implementation is compared with a solver-only baseline under the same solver settings. Across all operating scenarios, the optimized solutions substantially outperform the decentralized ground-based system (status quo baselines B1–B4) reported in Table 6, achieving significant reductions in cost, delivery time, and emissions. These improvements are most pronounced in instances where drones are available, either as the sole transportation mode or in combination with ground vehicles. The following subsections provide a detailed comparative analysis.

7.2.1. Cost analysis

The cost breakdown across the experimental instances (Fig. 5) shows that total system cost varies substantially with transportation mode, spatial dispersion, and demand intensity. Across all operating scenarios, ground-vehicle-only networks (INS1, INS4, INS7, and INS10) incur the highest costs. This outcome is primarily driven by long travel distances, limited routing flexibility, and reliance on a single vehicle to serve geographically dispersed pickup locations in rural settings.

Relative to the decentralized baselines (B1–B4), even optimized ground-vehicle-only solutions achieve substantial cost reductions, underscoring the value of centralized routing and coordinated operational planning. For example, in the LV–LW operating scenario, the optimized car-only configuration (INS1) reduces total cost by approximately 40% compared to the status quo (from \$777 to \$451). Similarly, in the HV–HW scenario, optimization alone reduces total cost by 83%, from \$2,449 in B4 to \$437 in INS10, demonstrating that significant gains can be achieved with optimization even without introducing new transportation modes.

Drone-enabled configurations further improve cost performance by reducing travel distances and enabling parallel service. Across operating scenarios, drone-only solutions achieve cost reductions of approximately 41–62% relative to optimized car-only networks. In high-volume scenarios (HV–LW and HV–HW), drone-only configurations consistently yield the lowest total cost and fastest delivery times, as multiple drones provide efficient point-to-point service and largely eliminate the need for intermediate transshipment or recharging. In these settings, incorporating ground vehicles offers limited additional benefit, even within multi-modal configurations.

In contrast, in low-volume but spatially dispersed networks such as LV–LW, multi-modal solutions emerge as optimal. Although total demand can be served by drones alone, long inter-node distances and limited range make intermediate recharging and transshipment operationally valuable. While the drone-only configuration is a close runner-up in cost, the multi-modal solution achieves substantially shorter delivery times, by nearly two hours on average, highlighting that spatial structure, rather than demand volume alone, plays a critical role in determining the optimal mode configuration.

These results demonstrate that the proposed framework does not favor a single transportation mode a priori. Instead, it systematically identifies the most cost-effective configuration based on the interaction among demand intensity, geographic dispersion, and operational constraints, providing actionable guidance for designing efficient rural last-mile delivery systems.

7.2.2. Delivery schedule analysis

This subsection examines delivery timeliness across the optimized instances. Fig. 6 reports the cumulative percentage of delivered order quantities reaching the diagnostic laboratory within 1-, 5-, 10-, and 15-h time thresholds for each operating scenario. Delivery timeliness is evaluated at the order level; therefore, in operating scenarios involving transshipment, such as LV–LW, only vehicles that perform the final delivery to the laboratory are included in the histograms. All results are benchmarked against the corresponding decentralized baselines (B1–B4).

The LV–LW operating scenario demonstrated the longest delivery times across all fleet configurations, despite its relatively low demand volume. This outcome is primarily driven by the network's sparse spatial structure and the long distances between pickup




Fig. 5. Cost breakdown per instance across all operating scenarios.

locations and the diagnostic laboratory. In this operating scenario, both single-mode fleet sizes have an average delivery time of around 10 h. The multi-modal configuration significantly improves delivery performance, reducing the weighted average delivery time by approximately 26%, from 10.4 h in the car-only configuration to 7.7 h. By contrast, in more spatially concentrated networks such as LV-HW, drone-only delivery substantially outperforms car-only operations, reducing delivery time by approximately 50%, and multi-modality does not improve delivery time.

In the high-volume operating scenarios (HV-LW and HV-HW), drone-only configurations provide the fastest delivery performance. In the HV-HW scenario, the optimized drone solutions (INS11 and INS12) deliver almost all orders within 5 h, while the car-only configuration (INS10) shows much later arrivals, with most deliveries occurring around the 10-h mark. In the HV-LW scenario, where demand is high and pickup locations are more spatially dispersed, the drone-based solutions (INS8 and INS9) still provide fast service, delivering approximately 27% of quantities within 1 h and 73% within 5 h, with complete delivery achieved by 10 h. The car-only solution (INS7), however, results in delayed deliveries, with all orders arriving after 10 h.

These results indicate that the effectiveness of multi-modal coordination in improving delivery timeliness depends strongly on geographic dispersion and order volume. When demand is geographically concentrated or drone fleet capacity is sufficient, drone-only operations provide the best delivery performance. Multi-modal coordination becomes most valuable in spatially dispersed networks, where transshipment and intermediate recharging enable timely and feasible service.

7.2.3. Total distance traveled and CO₂ emissions analysis

Fig. 7 illustrates the total travel distance by vehicle type and the corresponding CO₂ emissions across operating scenarios. Although the objective function minimizes total system cost rather than travel distance directly, variable operating costs are proportional to distance traveled; therefore, distance patterns provide a structural explanation for the observed cost differences.

Across all operating scenarios, the decentralized baseline exhibits the greatest total travel distance. Centralized optimization substantially reduces ground travel in car-only configurations, achieving distance reductions of approximately 40–80%, with the largest improvements observed in high-volume scenarios. Introducing drones yields additional distance savings, with drone-enabled configurations reducing total travel distance by roughly 50–83% relative to the status quo systems.

Comparisons across operating scenarios indicate that the marginal benefit of drone deployment increases with network scale and spatial dispersion. In high-volume scenarios, where travel distances are largest, drone deployment achieves the most pronounced absolute reductions in both distance and emissions, with reductions of approximately 18% in total travel distance and 99.9% in CO₂ emissions. In low-volume scenarios, the corresponding gains are more moderate, at approximately 14%.

In general, aerial service dominates ground routing under a wide range of operating conditions. However, in the LV-LW scenario, the multi-modal configuration reduces total travel distance by approximately 6% relative to the drone-only solution, while simultaneously increasing CO₂ emissions by approximately 1,800%. This trade-off reflects the introduction of ground vehicle travel in the multi-modal configuration and demonstrates that, under specific spatial and demand conditions, reductions in distance do not




Fig. 6. Pickup-to-delivery time distributions for all instances, weighted by delivered quantities.

necessarily translate into lower emissions. However, reductions in travel distance could translate directly into emission savings. That is, drone-enabled multi-modal configurations consistently achieve the lowest CO₂ emissions, with the strongest gains occurring in scenarios characterized by extensive ground travel, where emissions are reduced by more than 99% relative to the status quo baseline.

7.3. Comparison of weather-aware and non-weather-aware models

To assess the impact of incorporating weather conditions into vehicle energy consumption and route planning, we compare our weather-aware time-dependent forecast-based energy consumption model with a simplified range-based benchmark commonly used in the literature (e.g., Enayati et al., 2023). The range-based model assumes that each drone operates with a fixed maximum flight range, irrespective of wind speed or direction. This comparison highlights how accounting for weather-related constraints influences network configuration, vehicle operations, and overall system performance in drone-integrated scenarios.




Fig. 7. Distance traveled by vehicle type and corresponding CO₂ emissions across all operating scenarios, with emissions values shown atop the bars.

For this analysis, we replaced the weather-aware constraints (Eqs. (22) to (28)) with a set of constraints to account for the static drone flight range, as detailed below:

$$\Delta_i^v \leq \zeta^v, \quad \forall i \in \mathcal{N}, v \in \mathcal{V} \quad (38)$$

$$\Delta_{O^v}^v \geq \sum_{(i,j,v) \in \mathcal{A}} \mathcal{X}_{ij}^v \zeta^v, \quad \forall v \in \mathcal{V} \quad (39)$$

$$\mathcal{Z}_i^v \leq \sum_{j:(j,i,v) \in \mathcal{A}} \mathcal{X}_{ji}^v, \quad \forall v \in \mathcal{V}, i \in \mathcal{P}^v \quad (40)$$

$$\Delta_i^v \geq \zeta^v * \mathcal{Z}_i^v, \quad \forall v \in \mathcal{V}, i \in \mathcal{P}^v \quad (41)$$

$$\Delta_i^v - \Delta_i^v \leq \begin{cases} M \mathcal{Z}_i^v, & \forall v \in \mathcal{V}, i \in \mathcal{P}^v \\ 0, & \forall v \in \mathcal{V}, i \in \mathcal{N} \setminus \mathcal{P}^v \end{cases} \quad (42)$$

$$\Delta_i^v - L_{ij}^v \geq \Delta_j^v - M(1 - \mathcal{X}_{ij}^v), \quad \forall (i,j,v) \in \mathcal{A} \quad (43)$$

In this formulation, the variables Δ_i^v and Δ_j^v represent the remaining range of vehicle $v \in \mathcal{V}$ upon departure and arrival at node $i \in \mathcal{N}$, respectively. The parameter ζ^v denotes the maximum range that vehicle v can travel without recharging, and L_{ij}^v is the travel distance between nodes i and j using vehicle v .

Constraints (38) ensure that the remaining range of any vehicle upon departure from a node does not exceed its maximum range. Constraints (39) initialize each vehicle at its full range when departing from its operational base. Constraints (40) enforce that a vehicle can only be recharged or refueled at a designated node $i \in \mathcal{P}^v$ if a feasible arc $(j,i,v) \in \mathcal{A}$ exists. Constraints (41) specify that recharging at node i restores the vehicle to its full range. Constraints (42) maintain continuity by requiring that, unless recharged, a vehicle's departure range from a node equals its arrival range. Finally, constraints (43) update the vehicle's remaining range based on the distance traveled along arc (i,j,v) . Our results based on this range-based formulation revealed three key insights:

(1) **Weather-driven changes in routing topology and modal allocation:** The weather-aware model can actively determine how a multi-modal fleet is used, i.e., selecting the operating mode, shaping the routing structure, and coordinating drones and ground vehicles when time- and location-dependent energy requirements alter their relative efficiency. Rather than committing to a single mode across the planning horizon, the optimizer reallocates service between modes and inserts transshipment-based handoffs




Fig. 8. Comparison of optimal routing and operational decisions for INS3 under (a) weather-aware and (b) static-range models.

Table 10

Performance comparison of weather-aware and static-range models for INS3.

Formulation	Total Cost (\$)	Drone Share of Total Distance (%)	Weighted Average Delivery Time (h)	Drone Distance (mile)	Car Distance (mile)	Recharges	Transshipments
Weather-aware	258	86	7.7	409	65	6	1
Range-based	240	100	8.7	468	0	7	0

Note: Drone share is the ratio of drone distance to total (drone plus ground-vehicle) distance.

wherever doing so improves operational performance. This capability is significant because a static range assumption may settle on a single-mode plan without accounting for how energy consumption varies across arcs and departure times, masking coordinated alternatives that are in fact more efficient under the forecast conditions.

INS3, the LV–LW multi-modal instance, illustrates this behavior (Fig. 8, Table 10). Under the weather-aware formulation, the model relies primarily on drones early in the service period and introduces limited ground-vehicle support only when it becomes operationally beneficial. The result is a genuinely coordinated car-drone plan in which the two modes exchange load through one transshipment, and the drone's share of total distance falls from 100% to 86%. This restructuring is also visible topologically in the route maps: the weather-aware solution (Fig. 8a) introduces a ground-vehicle arc and a transshipment node that are entirely absent from the all-drone range-based solution (Fig. 8b).

This coordination yields a better-balanced plan: shifting work off the drones cuts drone travel from 468 to 409 miles and recharging from seven cycles to six, while weighted average delivery time improves from 8.7 to 7.7 h, at a modest cost increase from \$240 to \$258 for the added 65 miles of ground-vehicle travel. The balance among these competing effects depends on the underlying cost and operational parameters and cannot be identified by inspection; resolving such interdependent trade-offs is precisely the role of the proposed decision-support model.

By contrast, the range-based formulation assigns the entire service to drones here. Relying on a static effective-range assumption, it cannot detect the benefit of coordinated car-drone service and settles on a single-mode plan. Its lower reported cost should not be read as operational superiority: it reflects a simplified treatment of drone energy, and the weather-aware plan's higher cost is not an inefficiency but the price of a more reliable plan, since the range-based savings rest on energy assumptions that may not hold during execution.

(2) Improved resilience under headwind conditions: The weather-aware optimization model enhances network resilience by proactively adjusting drone operations in response to wind-related energy demands, particularly under headwind conditions, which significantly increase energy consumption. This behavior is illustrated in instance INS5, where routing and operational differences between the weather-aware and static range-based models are shown in Fig. 9. In this case, the weather-aware model includes an additional recharge cycle not required under the static model, resulting in a slight increase in total cost and delivery time. However, this added recharge ensures the route remains feasible under real-world wind conditions. By explicitly modeling wind-induced energy variation, the weather-aware approach reduces the risk of in-flight failures and improves the robustness and reliability of the overall delivery network.




Fig. 9. Comparison of optimal routing and operational decisions for INS5 under (a) weather-aware and (b) static-range models.




Fig. 10. Optimal routing and operations for INS8 under (a) flexible and (b) restricted delivery schedules.

(3) Strategic delays under tailwind conditions: The weather-aware model demonstrates the ability to strategically delay drone departures to capitalize on favorable wind conditions, particularly tailwinds occurring later in the day (as discussed in [Proposition 2-Section 5.4](#)). By aligning departures with wind patterns, the model reduces energy consumption and avoids unnecessary mid-route recharging, thereby improving both efficiency and reliability. To ensure practical applicability, the model can constrain waiting times within predefined delivery windows, allowing controlled departure flexibility. Within these limits, it optimizes the trade-off between energy use, recharging, and delivery time, while selecting the most efficient configuration (i.e., drone, ground vehicle, or hybrid). This capability is especially valuable in settings where recharging is costly or limited, enabling more adaptive and energy-efficient delivery strategies.

This behavior is illustrated in [Fig. 10](#) for INS8. Under a flexible delivery schedule ([Fig. 10a](#)), Drone 1 and Drone 2 depart at 15:48 and 17:16, respectively, taking advantage of tailwinds to complete their routes in 1.9 and 2.1 h, both without requiring recharging and before the latest allowable delivery time of 20:00. In contrast, under a stricter delivery schedule ([Fig. 10b](#)), the model dispatches drones earlier, before tailwinds emerge, forcing both to recharge mid-route to complete their deliveries. To ensure comparable time efficiency, we restrict the latest allowable delivery time to match the optimal completion time achievable with a single-car delivery.

Taken together, these findings show that weather-aware planning shapes multiple dimensions of drone-enabled logistics, including departure timing, recharge frequency and placement, routing, transshipment, and modal allocation, with the nature of the effect depending on the network's operational flexibility. Where alternative routes and modes are limited (INS5, INS8), weather information primarily adjusts scheduling and energy-management decisions while leaving route topology largely intact; where flexibility exists

(INS3), it can additionally reshape routing and the coordination of drones and ground vehicles. The value of weather-aware modeling should therefore not be judged by visible route deviations alone: its broader contribution is to account for time- and location-dependent energy requirements when determining how transportation resources are scheduled and coordinated. A static range-based formulation, by omitting this variation, may underestimate operational energy requirements and settle on less coordinated single-mode plans, whereas forecast-based energy consumption provides a more reliable basis for feasible, operationally realistic delivery.

7.4. Sensitivity analysis

This section evaluates the robustness of the proposed multi-modal delivery framework by examining how key economic and operational parameters, beyond the base-case values considered in Section 6, affect *system response*—evaluated using total operational cost, optimal fleet configuration, drone utilization share, and service performance. The analysis focuses on identifying structural shifts in the cost-optimal mode configuration, referred to here as regimes. A regime denotes the dominant transportation configuration selected by the optimization model, including car-only, drone-only, and hybrid (multi-modal) solutions that combine ground vehicles and drones. The objective is to understand how parameter changes induce transitions between these regimes and to quantify the resulting trade-offs between cost efficiency and delivery performance under realistic operating conditions.

7.4.1. Sensitivity to marginal vehicle operating costs

Vehicle operating costs directly affect the marginal cost of serving demand and therefore play a central role in determining the cost-optimal fleet configuration. To assess these effects, drone and ground-vehicle operating costs are varied independently by applying multiplicative factors ranging from 0.4 to 2.0 of the base-case value, in increments of 0.2. When drone costs vary, ground-vehicle costs are held fixed at their base-case levels, and vice versa. The model is solved to optimality for each parameter setting.

7.4.1.1. Effect of Drone Marginal Operating Cost. Across all operating scenarios, increasing drone operating cost leads to a monotonic increase in total system cost. However, the corresponding impact on delivery performance depends strongly on demand volume, fleet composition, and network structure. As shown in Fig. 11, in high-volume scenarios with multiple active vehicles, low drone operating costs relative to car costs lead to drone-only solutions and the shortest delivery times. As drone costs increase, service performance remains largely stable, indicating that moderate cost increases primarily affect total cost rather than delivery time. When drone costs approach parity with car costs, the optimal solution shifts to a multi-modal regime, in which some demand is reassigned to ground vehicles to contain further cost growth, resulting in longer delivery times.

The magnitude of this trade-off depends on fleet size. In the HV–LW operating scenario, where the fleet includes more drones, shifting from drone-only to hybrid configurations reduces total cost by approximately 2–11% relative to drone-only operation, while increasing average delivery time by about 12%. In the HV–HW operating scenario with a smaller fleet, cost reductions of approximately 4–19% are achieved, but at the expense of substantial service degradation, with average delivery times increasing by roughly 74–93%. This suggests that when fleet size is limited, cost-driven changes in configuration have a much stronger effect on service performance. The corresponding average delivery time trends are shown in Fig. D.1 in Appendix D.

In low-volume operating scenarios, system responses are more abrupt. Because a single vehicle type is sufficient to meet demand, the choice of configuration is strongly influenced by spatial structure. In the geographically concentrated LV–HW operating scenario, a single vehicle can efficiently serve all demand, and changes in drone operating cost trigger direct transitions between drone-only and car-only solutions. In contrast, in spatially dispersed networks such as the LV–LW operating scenario, hybrid configurations remain optimal across a wider range of drone costs, as multi-modality reduces total travel distance and recharge-related overhead. Only when the drone operating cost decreases substantially relative to the car cost does the system transition to drone-only operation.

7.4.1.2. Effect of Car Marginal Operating Cost. Fig. 12 confirms and complements the observations from drone cost sensitivity analysis. Increasing car operating costs consistently shifts the system toward greater reliance on drones, reducing average delivery time but increasing total system costs.

In high-volume operating scenarios, drone-only solutions remain cost-optimal even when ground-vehicle operating costs are 80% or higher relative to drone costs. Hybrid or car-dominant configurations emerge only when ground-vehicle costs fall below roughly 70% of drone marginal cost. In these regimes, reallocating part of the workload to ground vehicles reduces total system cost by approximately 3–18%, but at the expense of increased average delivery time. The magnitude of service degradation varies across operating scenarios: in HV–LW, delivery time worsens by approximately 25–75%, while in HV–HW, the shift toward car usage leads to even more pronounced increases in delivery time due to the smaller available fleet. This demonstrates that cost savings achieved by increasing car usage can be accompanied by substantial service degradation, particularly in time-sensitive settings. The corresponding average delivery time trends are shown in Fig. D.2 in Appendix D.

In low-volume scenarios, the effect of car cost variation is more discrete. In the single-fleet LV–HW operating scenario, reducing car operating cost by approximately 50% is sufficient to trigger a regime change from drone-only to car-only operation. In the spatially dispersed LV–LW operating scenario, increasing car operating cost rapidly shifts the system toward drone-only solutions, while decreasing car cost increases the car share within hybrid configurations.

7.4.1.3. Relative Car-Drone Cost Ratio. To integrate the effects of drone and car operating costs, we analyze system behavior as a function of the car-to-drone operating cost ratio. This ratio captures the relative economic attractiveness of the two delivery modes and provides a compact representation of the regime changes observed across the individual marginal cost sensitivity analyses.




Fig. 11. Sensitivity of total network cost and optimal delivery configuration to changes in drone operating cost across the four operating scenarios.

Fig. 13 reveals a clear regime-based structure across all four operating scenarios. In the high-volume scenarios, drone usage increases monotonically with the cost ratio and exhibits a pronounced transition to drone-only operation once the ratio exceeds approximately 1.04, indicating that the marginal cost of car-based delivery is about 4% higher than that of drone delivery. Below this threshold, hybrid configurations dominate; however, the degree of drone utilization differs substantially across operating scenarios. In particular, the HV-LW operating scenario consistently maintains significantly higher drone usage than the HV-HW operating scenario at comparable cost ratios, with differences approaching 40–50 percentage points in hybrid regimes. This divergence reflects greater fleet redundancy in HV-LW, which allows drones to absorb a larger share of deliveries before a system-wide regime switch becomes necessary.

Low-volume operating scenarios exhibit more abrupt responses to changes in the cost ratio. In the LV-HW scenario, a regime shift occurs at a much lower ratio (around 0.58), as a single vehicle type (drone) is sufficient to serve all demand, and cost changes directly trigger transitions between car-only and drone-only solutions. In contrast, the spatially dispersed LV-LW case displays a wider hybrid region. Here, drone usage increases gradually with the cost ratio and stabilizes at approximately 70% over a broad intermediate range, indicating sustained multi-modal operation driven by reduced travel distances and recharge-related overhead.

The implications of these utilization patterns for cost-service trade-offs are illustrated in Fig. 14. Each point in the figure corresponds to a cost-optimal solution generated by scaling either the drone or ground-vehicle operating cost by a multiplier of 0.4–2.0 relative to its base-case value, with the other cost held fixed. Across all scenarios, low cost ratios favor car-dominated solutions with lower total system cost but longer delivery times, while high cost ratios favor drone-intensive solutions that minimize delivery time at a higher cost. Hybrid configurations emerge in intermediate regimes, where combining both modes yields the most favorable balance between cost efficiency and service performance. The relationship between cost and service is highly non-linear. In several regimes, modest reductions in total cost are accompanied by disproportionately large increases in average delivery time, particularly when the system transitions from drone-intensive to car-intensive operation, where reduced parallelism and longer tours amplify service delays. Conversely, increasing reliance on drones can yield substantial service improvements at relatively small additional cost, up to the point where delivery time reaches its minimum achievable level and further increases in cost yield diminishing service returns.

Taken together, these results demonstrate that the relative car-drone cost ratio governs not only vehicle utilization but also the shape of the cost-service frontier. The analysis highlights a central planning trade-off: whether marginal cost savings justify potentially large losses in service performance, and under which operational conditions increased reliance on drones yields meaningful service gains relative to their additional cost.




Fig. 12. Sensitivity of total network cost and optimal delivery configuration to changes in car operating cost across the four operating scenarios.




Fig. 13. Drone usage as a function of the car-to-drone cost ratio across the four demand-volume scenarios.

7.4.2. Sensitivity to drone recharge operations

This section evaluates how drone recharge operations influence the economic and operational performance of the proposed multi-modal delivery system. While recharge cost affects the economic attractiveness of drone operations, recharge duration constrains effective fleet availability and system capacity.

7.4.2.1. Effect of Recharge Cost. Drone recharge operations introduce a cost component that differs from marginal vehicle operating expenses. While marginal operating costs scale directly with vehicle usage, recharge costs accumulate through repeated energy replenishment and associated downtime and may therefore influence both economic performance and operational behavior. To assess these effects, we vary the unit drone recharge cost relative to the base-case value of \$2 by considering higher cost levels of \$5, \$10, \$15, and \$20, while holding all other parameters fixed, and decompose total system cost into transportation, recharge, and transshipment components.

Fig. 15 shows that transportation cost dominates total system cost at low recharge levels, accounting for over 90% of total cost in most scenarios. As recharge cost increases, its contribution grows nonlinearly, particularly under high demand and adverse weather conditions, revealing scenario-dependent differences in economic sensitivity and operational flexibility.




Fig. 14. Cost-service trade-offs across all operating scenarios. Each point represents a cost-optimal solution under a given cost ratio; marker shape and color indicate the optimal configuration, while marker size reflects the magnitude of the cost imbalance.

In scenarios where drone-only configurations are optimal, recharge cost affects system performance primarily through the economic dimension. In both high-volume cases, increasing recharge cost leads to an approximately linear increase in total cost, while delivery time, drone utilization, and the optimal regime remain unchanged. For example, in the LV-HW case, recharge cost grows from about 6% to nearly 32% of total cost, yet the configuration and service performance remain stable. Similarly, in HV-HW and HV-LW scenarios, the recharge-cost share increases from roughly 4% to 28% and from 2% to 20%, respectively. These patterns indicate that, under high demand, drones are operationally indispensable and recharge pricing does not trigger regime shifts.

By contrast, the LV-LW scenario exhibits genuine operational adaptation. In this spatially dispersed, low-demand setting, hybrid configurations remain optimal across a wide range of recharge costs, and increasing recharge cost induces systematic reallocation between modes. As recharge cost rises to moderate levels (up to about \$15), drone usage declines from approximately 86% to 71%, reducing recharge operations by nearly 83% and improving average delivery time by about 28%. Beyond this threshold, further increases in recharge cost lead to sharper reductions in drone usage (to around 68%) and a corresponding increase in delivery time of roughly 5%, as car-dominant operation becomes more prominent. This non-monotonic response illustrates that recharge cost can influence not only total cost but also service performance when sufficient spatial and modal flexibility exists.

These observations highlight the importance of explicitly modeling recharge operations in drone-assisted medical logistics. When multi-modality is feasible, recharge cost becomes a structural driver of vehicle utilization and service outcomes, shaping nonlinear trade-offs between cost and performance.

7.4.2.2. Effect of Recharge Duration. Recharge duration represents a purely operational constraint that affects vehicle availability rather than monetary cost. Longer recharge times increase drone downtime, reduce effective fleet capacity, and may induce congestion at recharge locations. To isolate these effects, we vary the drone recharge duration relative to the base-case value of 10 min by considering recharge times of 30, 45, and 60 min, while holding all cost parameters fixed, and decompose total operational time into travel, recharge, and transshipment components. The base-case value of 10 min reflects rapid battery-swapping operations, whereas longer durations correspond to on-site recharging under progressively more constrained conditions, such as limited charging power, shared infrastructure, automated charging pads in remote settings, or geographically isolated deployment environments.

Fig. 16 shows that increasing recharge duration systematically increases total operational time across all scenarios as expected; however, the magnitude and nature of this increase depend strongly on demand volume, fleet redundancy, and weather conditions.

In low-volume scenarios, the impact of recharge duration is governed primarily by spatial structure. In the LV-HW operating scenario, the system remains drone-only across all tested values, and increasing recharge duration from 10 to 60 min raises the share of recharge time from approximately 14% to nearly 48% of total operational time without inducing a regime change. In contrast, the




Fig. 15. Recharge cost sensitivity analysis. Bars show the total network cost decomposed into transportation, recharge, and transshipment components. Percentage labels indicate each component's share of the total cost.

spatially dispersed LV-LW case exhibits threshold behavior: recharge durations up to 45 min degrade service performance (delivery time increases from 7.7 to 11.2 h) without affecting cost or vehicle mix, while a recharge duration of 60 min triggers a reallocation toward ground vehicles. This shift reduces drone congestion and lowers delivery time to 9.9 h at a marginal increase in total cost.

High-volume scenarios display clearer capacity effects. In the HV-LW operating scenario, recharge duration has a muted impact on performance, with recharge time accounting for less than 27% of total operational time even at the longest duration. By contrast, in the HV-HW case, recharge time increases from approximately 9% to nearly 37% of total operational time as recharge duration rises, reflecting the compounding effect of adverse weather and increased energy demand.

These observations demonstrate that recharge duration acts as a capacity and congestion constraint rather than a cost driver. Moderate increases primarily reduce service quality without inducing structural changes, while extreme delays force rebalancing of the vehicle mix to preserve operational feasibility. This behavior contrasts with recharge cost sensitivity, where adjustments are driven by economic trade-offs, highlighting the importance of explicitly modeling recharge duration in drone-integrated delivery systems.

7.4.3. Sensitivity to drone capacity

Drone payload capacity is constrained by medical-specific handling requirements, such as cold-chain packaging, protective containers, or regulatory safety margins, which reduce the usable payload available for order consolidation. Unlike operating or recharging costs, capacity constraints act as feasibility drivers by limiting how much demand can be served per trip and by shaping how deliveries are consolidated across vehicles. To examine these effects, we vary drone payload capacity relative to the base-case value of 15 lb by progressively reducing it in increments of 2.5 lb, corresponding to reductions of approximately 17%, 33%, and 50%.

Across operating scenarios, the impact of capacity constraints depends on whether aggregate demand can be accommodated within available payload limits. When capacity is sufficient, it is non-binding and does not alter the optimal modal structure. In these regimes, system behavior mirrors the operating-cost sensitivity results, with drone-only or drone-dominated configurations remaining optimal and service performance largely unchanged.

In low-volume scenarios, capacity constraints primarily affect feasibility rather than inducing gradual regime shifts. In the LV-HW case, where demand is spatially concentrated, a single drone can serve all orders when capacity exceeds the aggregate payload requirement; once capacity falls below this threshold, the optimal solution transitions directly to a car-only configuration. In the LV-LW case, where demand is geographically dispersed but relatively small, hybrid configurations remain optimal across feasible capacity levels.

High-volume scenarios exhibit clearer structural transitions driven by payload limitations. Fig. 17 illustrates this behavior for the HV-HW operating scenario; other cases are discussed in the text and omitted for brevity. At the base-case capacity (15 lb),




Fig. 16. Recharge duration sensitivity analysis. Bars show total network time decomposed into transportation, recharge, and transshipment components. Percentage labels indicate each component's share of total operational time.

the system adopts a drone-only configuration, with drones handling 100% of delivered quantities and traveled arcs, achieving the lowest total cost (\$261) and the shortest average delivery time (2.98 h). As capacity is reduced to 12.5–7.5 lb, drone-only operation becomes infeasible. Rather than eliminating drone routes, the model reallocates deliveries across multiple drones and introduces ground-vehicle support, yielding hybrid configurations.

This transition is marked by a divergence between quantity-based and arc-based drone usage. While arc-based drone participation remains high (approximately 61–62%), indicating continued reliance on drones for spatial coverage, quantity-based drone usage drops sharply (to 27–47%), reflecting payload limitations. Consequently, total cost increases from \$261 to \$326–\$344, and average delivery time rises from 2.98 to 3.24–3.67 h due to reduced consolidation efficiency and greater reliance on slower ground transport.

A similar pattern emerges in the HV-LW scenario, but the transition is smoother. The availability of additional vehicles allows demand to be redistributed across a larger fleet, delaying binding capacity effects and resulting in more gradual changes in cost, delivery time, and modal composition.

These findings highlight that drone capacity constraints differ fundamentally from economic or recharging parameters. While cost and recharge sensitivities primarily affect operating efficiency and service performance, payload limitations directly determine feasibility and consolidation structure. When capacity is non-binding, modal selection is governed by relative operating costs; once capacity becomes binding, multi-modal operation emerges as a necessary mechanism to preserve feasibility and acceptable service levels.

7.4.4. Sensitivity to delivery time window constraints

Delivery time windows are a critical operational constraint in medical logistics, as delayed deliveries can compromise service quality or patient outcomes. To evaluate the system's responsiveness to time pressure, we progressively tighten the latest allowable delivery time to the most restrictive feasible level, while holding all cost and fleet parameters fixed at their base values.

Our results reveal strong robustness of drone-only configurations to tightening of the latest allowable delivery time window. While stricter deadlines do not alter the optimal delivery mode in drone-only settings, they affect operational scheduling by compressing completion times and modifying recharge operations under high-wind conditions. In contrast, under multi-modal configurations, tighter delivery windows increase reliance on the overall fleet, leading to a more balanced use of vehicles.

In operating scenarios where drone-only operation is already cost-optimal, and weather conditions do not significantly intensify recharge requirements (e.g., LV-HW and HV-LW), tightening the delivery time window leaves total system cost, delivery mode, and average delivery time unchanged. In these cases, drones complete deliveries at the earliest possible time, and stricter deadlines do not impose binding constraints on the system.

By contrast, under high-wind conditions, even when drone-only operation remains optimal and scheduling explicitly accounts for weather effects, tightening the delivery window can influence operational timing and recharge decisions. In particular, the HV-HW




Fig. 17. Vehicle-level quantity allocation under drone capacity sensitivity for the HV-HW operating scenario. Bars show how delivered quantities are redistributed across individual drones and the ground vehicle as drone payload capacity varies.




Fig. 18. Sensitivity of drone operations to delivery time-window tightening in the HV-HW operating scenario.

scenario provides a more informative operational response under time pressure. Although drone-only operation remains optimal, tightening the latest allowable delivery time below the baseline induces schedule compression rather than mode switching. Fig. 18 illustrates this mechanism: under a loose delivery deadline, Drone 2 adjusts its operation to prevailing high-wind conditions, departs later, and the whole operation completes after approximately 8.5 h with fewer recharging operations. When the delivery deadline is tightened, the model initiates Drone 2 at the earliest pickup time (8:00 AM) and introduces an additional, earlier recharge, enabling completion after approximately 4.6 h at a modest increase in operating cost (about 1%). This behavior highlights how, under high-wind conditions, delivery time-window tightening can increase operational intensity through recharging even when the optimal configuration remains unchanged.

In contrast, the LV-LW case exhibits a structural shift under tighter delivery time windows. As deadlines tighten, the system moves toward greater multi-modality and more balanced vehicle utilization. Under the base setting, the model selects a drone-dominant configuration as the cost-optimal mode. As delivery deadlines tighten, the model-expectedly-shifts toward greater ground-vehicle use to enable parallel service, improving timeliness at higher cost. This reveals a cost-service trade-off driven by delivery constraints and spatial dispersion rather than operating costs alone. When an additional drone is available (two drones and one ground vehicle), the model selects the second drone, indicating that the shift toward ground-vehicle use in the one-drone-one-car fleet is driven partly by binding fleet availability rather than intrinsic drone inefficiency.

7.4.5. Sensitivity to available fleet size

To examine how available fleet resources affect system performance, we conduct a sensitivity analysis on the number of available ground vehicles and drones in two representative low-wind scenarios (chosen to hold wind constant), LV-LW and HV-LW, allowing us to isolate the effect of demand volume. The analysis does not determine a strategic fleet size; it evaluates how the operational model responds under alternative available-fleet limits. Fig. 19 reports total system cost and total operational time across the evaluated configurations, with each cell annotated by the numbers of cars and drones the model activates—distinguishing availability from actual utilization. Each cell additionally displays an icon denoting the logistics strategy selected for that configuration (ground-vehicle-only, drone-only, or multi-modal), so that the optimal mode and the gap between available and active vehicles can be read directly from the heatmap. Because both scenarios share calm conditions, the reported fleet behavior reflects demand and spatial structure rather than favorable wind; under stronger headwinds, drone energy feasibility tightens, which would if anything shift the activated mix toward greater ground-vehicle use and reinforce the value of multi-modal coordination.

Fig. 19a shows that increasing the number of available vehicles beyond a relatively small threshold yields diminishing cost improvements in the LV-LW scenario. When no drones are available, the model relies entirely on ground transportation, resulting




Fig. 19. Sensitivity of system performance to available fleet limits in the LV-LW and HV-LW operating scenarios. The heatmaps report total system cost and total operational time under alternative numbers of available cars and drones. Cell annotations indicate the numbers of active drones (D) and cars (C) selected by the optimization model.

in comparatively high total costs, particularly when only one car is available. The total cost reaches \$451 with one car. Introducing one drone reduces the cost to approximately \$258 across the evaluated car limits, indicating that even limited drone availability can reduce reliance on long ground-vehicle routes. When two drones are available, the cost decreases further to approximately \$210, while configurations with three or more drones reach a minimum-cost level of approximately \$186. This pattern indicates that the marginal cost benefit becomes negligible once three drones are available, even when additional vehicles remain available to the model.

The effect of additional available fleet resources on service performance is more pronounced than their effect on cost. As shown in Fig. 19b, total operational time decreases with drone availability, falling from more than 10.7 h in the car-only configurations to approximately 6.6 h when three or more drones are available. The reduction in operational time reflects the value of parallel service and reduced waiting enabled by multiple drones. In contrast, increasing the number of available ground vehicles provides relatively limited time savings once drone service is available, indicating that drone availability is the primary driver of service-time improvement in this low-volume, spatially dispersed scenario.

Table 11

Benchmark comparison between the unreduced network and the reduced network.

Operating scenario	Instance	Configuration	Network	Number of arcs	Arc reduction (%)	Binary variables	Binary reduction (%)	Status (final gap)
HV–LW	INS8	Drone-only	Full	3392	–	136,623	–	No feasible incumbent
HV–LW	INS8	Drone-only	Reduced	652	80.8	32,063	76.5	Feasible incumbent (0.029)
HV–LW	INS9	Hybrid	Full	4345	–	176,129	–	No feasible incumbent
HV–LW	INS9	Hybrid	Reduced	997	77.1	45,557	74.1	Feasible incumbent (0.015)
HV–HW	INS11	Drone-only	Full	3291	–	176,434	–	No feasible incumbent
HV–HW	INS11	Drone-only	Reduced	1201	63.5	68,102	61.4	Feasible incumbent (0.008)
HV–HW	INS12	Hybrid	Full	4429	–	238,181	–	No feasible incumbent
HV–HW	INS12	Hybrid	Reduced	1637	63.0	90,205	62.1	Feasible incumbent (0.014)

Note: The unreduced (i.e., full) network contains all arcs generated during network construction prior to preprocessing and arc reduction. The reduced network corresponds to the implementation used in the main computational experiments. All runs were subject to a 3-h time limit. Arc and binary-variable reductions are computed relative to the unreduced network.

The HV–LW scenario shows a different pattern, as higher demand creates more scope for coordinated car–drone service. In the ground-vehicle-only configurations (Fig. 19c), total cost falls from \$444 with one car to \$392 with three, after which a fourth car stays inactive. As drones become available, the model activates one to three drones alongside a single car, transitioning to a drone-only solution (\$170, all cars inactive) only at four available drones. Unlike LV–LW, ground vehicles therefore remain in active use across most configurations, reflecting the value of multi-modal coordination under high demand. Operational time (Fig. 19d) falls sharply with drone availability, from 10.2 to 5.95 h, again driven primarily by drones rather than additional cars.

These results demonstrate that the operational value of ground vehicles depends on the number of drones available to the model. When drone availability is limited, coordinated car–drone operation allows the model to distribute the workload between transportation modes and improve system performance. As additional drones become available, reliance on ground transportation decreases. Once the available drone fleet is sufficient to serve the demand efficiently, the model may select a drone-only configuration even though ground vehicles remain available.

The comparison between the two operating scenarios further shows that the effect of available fleet resources depends on demand volume and network requirements. In the lower-volume LV–LW scenario, the model transitions to a drone-only solution when two drones are available, and additional ground vehicles provide no operational benefit. In the higher-volume HV–LW scenario, coordinated car–drone operation remains beneficial when one to three drones are available, and the transition to drone-only service occurs only when enough drones are available to handle the total demand. Thus, the higher-volume scenario requires greater drone availability before ground-vehicle support is no longer selected.

Overall, the sensitivity analysis shows that the operational value of additional available vehicles depends jointly on demand volume, network structure, transportation mode, and the planning objective. Additional drone availability can substantially reduce cost and operational time, while the value of additional ground vehicles depends on whether coordinated multi-modal service remains beneficial. The results also show that increasing the available fleet does not necessarily increase active vehicle utilization, because the model activates only the resources that contribute to the optimal operating plan. Accordingly, these findings should be interpreted as an operational evaluation of alternative available-fleet limits rather than as a strategic recommendation regarding long-term fleet acquisition or fleet size.

7.4.6. Scalability and computational behavior

This subsection evaluates the computational scalability of the proposed optimization framework and identifies the primary drivers of solver run time. The analysis includes two complementary components. First, a benchmark comparison directly quantifies the computational impact of the proposed preprocessing and arc-reduction procedures. Second, scalability trends are examined using the twelve test instances (INS1–INS12) derived from the four operating scenarios (LV–HW, LV–LW, HV–HW, and HV–LW), which span increasing levels of problem size, fleet complexity, and energy-feasibility interactions.

Solver performance is measured using wall-clock run time (seconds), with a uniform time limit of 10,800 s. Instances reaching the time limit are retained to avoid selection bias and to reflect practical computational constraints. To examine growth trends, instances are grouped by (i) number of pickup locations, (ii) total number of network arcs, and (iii) fleet size. For each group, the average solver run time is computed across all instances.

To isolate the effect of preprocessing, we compare the reduced network used in the main computational experiments with an unreduced network containing all arcs generated during network construction prior to preprocessing and arc reduction. Both networks are solved using the same mathematical formulation, hardware, solver settings, and 3-h time limit. The benchmark focuses on the four largest high-volume instances because these cases represent the primary computational bottlenecks of the test set and therefore provide the most informative setting for evaluating scalability improvements. For each instance, we report the number of arcs, the arc-reduction percentage, the number of binary variables, the binary-variable reduction percentage, and the solver status with final optimality gap.

Table 11 demonstrates the computational impact of the proposed preprocessing and arc-reduction procedures. Across the four benchmark instances, the reduced network decreases the number of arcs by 63.0%–80.8% and the number of binary variables by




Fig. 20. Average solver runtime growth as a function of instance size and fleet composition. Panel (a) shows average runtime by number of nodes, panel (b) by number of arcs, and panel (c) by number of vehicles. Each point represents the average runtime across all instances within the same size category, including time-limited runs. Runtime is shown on a logarithmic scale.

61.4%–76.5%. Under the same 3-h computational budget, the unreduced-network baseline did not identify a feasible incumbent in any of the reported cases, whereas the reduced network produced feasible solutions with final optimality gaps below 0.029. Increasing the solver time limit for the unreduced-network runs from 3 h to 6 h did not improve the solver outcome, as no feasible incumbent was identified in any instance. These results indicate that the proposed preprocessing and arc-reduction procedures substantially reduce model size and significantly improve tractability for the larger weather-aware instances considered in this study, with the reported reductions reflecting the combined effect of the preprocessing and arc-reduction procedures. Having isolated the computational impact of preprocessing, we next examine the broader scalability patterns across all twelve instances.

Fig. 20a and b report average run time as a function of node and arc-count levels. Although run time increases with both metrics, the growth is highly nonlinear and spans several orders of magnitude. Notably, instances with moderate node counts but large arc sets exhibit run times comparable to much larger instances, indicating that solver difficulty is driven not only by problem size but also by the density of feasible routing alternatives captured by the arc set.

The arc set, in turn, reflects spatial dispersion and energy feasibility. As pickup locations become more widely distributed, feasible drone paths require additional intermediate connections, recharging options, and weather-dependent feasibility checks, leading to larger and denser arc sets. Consequently, distance-related effects manifest computationally through arc expansion rather than node growth, explaining why some instances with relatively few nodes but many arcs are computationally demanding.

Fig. 20c examines solver run time as a function of fleet size and reveals a pronounced threshold effect. Instances with one or two vehicles are typically solved within seconds, whereas instances with three or more vehicles experience a sharp increase in run time and frequently reach the time limit. This transition reflects a combinatorial expansion of routing, scheduling, and recharging decisions: adding vehicles increases the feasible task assignments, the overlap of service windows, and the interaction among recharge schedules under time-dependent weather and energy constraints. Beyond this point, solver effort is dominated by cross-vehicle feasibility interactions rather than by problem size metrics alone.

These observations indicate that computational difficulty in weather-aware, energy-constrained drone logistics systems is governed by the interaction between spatial dispersion, arc-level feasibility, and fleet complexity. Arc count and fleet size, therefore, emerge as practical indicators of computational stress, while node count alone is insufficient. The operating scenarios considered in this study thus provide realistic stress tests that capture the dominant sources of computational complexity encountered when deploying the proposed optimization framework in practical medical logistics settings.

7.5. Managerial insights

The computational results provide practical guidance for designing and operating drone-assisted medical delivery systems in rural environments. Rather than favoring a single technology, the proposed framework reveals how optimal operational configurations and active vehicle-utilization patterns emerge from the interaction between demand intensity, spatial dispersion, weather conditions, operational constraints, and the available fleet. Together, these factors determine not only which available vehicles should be activated, but also how multi-modal operations should be coordinated to balance cost efficiency and service performance.

First, the value of drones is strongly influenced by demand volume, spatial dispersion, and the number of drones available for operation. When sufficient drone resources are available, aerial point-to-point service and parallel operations can produce substantial reductions in both cost and delivery time. However, when drone availability is limited, coordinated car–drone operations may remain beneficial, particularly in higher-volume or geographically dispersed networks.

In these settings, ground vehicles can complement drones by sharing the service workload, reducing drone travel and recharge requirements, and supporting operational feasibility. As additional drones become available, the model may progressively reduce its reliance on ground transportation and ultimately select a drone-only configuration.

Importantly, the transition between operational regimes is nonlinear: additional vehicle availability does not necessarily lead to proportional improvements in performance or to the activation of every available vehicle. The preferred logistics strategy depends

on the operational contribution of each mode under the specific demand, network, weather, and resource conditions. These findings imply that effective operation of rural medical delivery systems requires adaptive mode allocation driven by spatial structure, energy feasibility, available resources, and service-level requirements, rather than fixed technology preferences or cost comparisons alone.

The results also highlight the role of charging infrastructure as a critical enabler of drone-based medical logistics. Recharge costs and recharge durations do not merely affect operational expenses; they also reflect the availability, accessibility, and efficiency of charging infrastructure in rural areas. When charging is fast and inexpensive, such as in settings with battery-swapping capabilities or well-distributed charging facilities, drones retain their operational advantage across a wide range of demand and weather conditions. Conversely, when charging is slow or costly, drone feasibility becomes constrained, and multi-modal configurations emerge as a more resilient alternative. This finding implies that investments in charging infrastructure can be as important as increasing available vehicle resources in unlocking the full potential of drone-assisted delivery systems, particularly in rural regions where infrastructure scarcity can limit the scalability of aerial logistics.

Second, sensitivity analyses reveal nonlinear trade-offs between cost and service performance. Moderate increases in drone operating or recharge costs often affect total cost without altering the selected operational configuration, indicating that the preferred logistics strategy may remain stable across a range of cost conditions. However, when multi-modality is feasible, changes in cost, recharge duration, payload capacity, or available vehicle resources can induce regime shifts in mode allocation. These transitions demonstrate that small cost savings may come at disproportionate service degradation, underscoring the importance of explicitly balancing cost and timeliness in operational decisions.

Third, weather-aware planning can affect departure timing, recharge placement and frequency, route selection, vehicle activation, and modal allocation. In rural networks with limited alternative paths, its main effects may appear through scheduling and energy-management decisions while the broad route structure remains similar. When alternative routes or transportation modes are available, forecast-dependent energy requirements may also produce noticeable changes in routing and the coordination of drones and ground vehicles. Integrating forecast-based energy modeling, therefore, enhances operational robustness by ensuring that operational decisions reflect the energy conditions expected during execution, even when visible route changes are limited.

Finally, the available-fleet sensitivity analysis shows that initial increases in drone availability can substantially improve both operating cost and service time, while the marginal benefit of additional vehicles diminishes once sufficient operational flexibility is available. The results also demonstrate that the optimization model may leave some available vehicles inactive when they do not contribute to the optimal operating plan. In lower-volume scenarios, the model can transition relatively quickly toward drone-only service, whereas higher-volume scenarios may benefit from coordinated car-drone operations across a wider range of available-fleet configurations. Available-resource planning should therefore consider demand volume, spatial characteristics, cost, and service-level objectives rather than relying on cost minimization alone.

The proposed framework serves as an operational decision-support tool that enables healthcare providers and logistics operators to determine how a specified set of available drones and ground vehicles should be activated, routed, scheduled, recharged, and coordinated under realistic rural and medical constraints.

8. Conclusion and future research

This study developed a weather-aware, multi-modal optimization framework for medical supply delivery in rural regions using drones and ground vehicles over a single-period operational planning horizon. The proposed model jointly integrates routing and scheduling decisions with spatiotemporal energy consumption, explicit recharge operations, and transshipment options, enabling operationally realistic planning under weather variability and infrastructural constraints. Using real-world data from the Shield Illinois testing program, the framework demonstrates how drone-enabled logistics can be systematically coordinated with ground transportation to support flexible, efficient, and reliable rural healthcare delivery.

Computational experiments across twelve instances spanning four operating scenarios demonstrate substantial and quantitatively significant performance improvements achieved through the proposed optimization framework and drone integration, across multiple operational and environmental performance metrics. For instance, relative to the decentralized ground-based status quo, centralized optimization of ground vehicle routing alone reduces total system cost by approximately 40–83%. Incorporating drones through multi-modal coordination yields an additional 38–62% cost reduction, while drone-enabled configurations reduce weighted average delivery time by up to 85% and CO₂ emissions by more than 99%. In the primary high-volume instances with sufficient drone availability, drone-only operation achieves the lowest cost and fastest service. However, the available-fleet sensitivity analysis shows that coordinated car-drone operation remains beneficial when drone availability is limited. In low-volume but spatially dispersed networks, multi-modal configurations may also emerge as optimal.

Sensitivity analyses further reveal regime-based behavior governed by the car-drone operating cost ratio, with highly non-linear cost-service trade-offs. The available-fleet sensitivity results additionally demonstrate that increasing the number of available vehicles does not necessarily increase active vehicle utilization. Instead, the model activates only the subset of vehicles that contributes to the optimal operating plan, and the transition from coordinated multi-modal service to drone-only operation depends on demand volume and drone availability.

A key methodological contribution of this work is the integration of spatio-temporal weather effects through departure-time-dependent energy functions. This formulation enables the optimization model to adapt scheduling and energy-feasibility decisions based on wind forecasts without introducing additional state variables or requiring real-time re-optimization. The results show that weather-aware planning enhances operational robustness by introducing additional recharging under adverse headwinds and exploiting delayed departures under favorable tailwinds to reduce energy consumption and avoid unnecessary mid-route recharging.

The operational effects of weather information depend on the flexibility of the underlying network. When alternative paths and transportation modes are limited, weather awareness primarily affects departure timing, recharge placement and frequency, and energy-feasibility margins while the broad routing structure remains similar. When greater routing or multi-modal flexibility is available, forecast-dependent energy requirements may also alter route selection, vehicle activation, transshipment, and modal allocation.

The framework is intentionally designed as a deterministic, day-ahead decision-support tool, in which forecast-based energy functions are preprocessed and treated as fixed inputs during optimization. This design reflects practical constraints in healthcare logistics, including regulatory requirements, coordination complexity, and limited ability to re-optimize in real time. While richer vehicle dynamics and adaptive recourse actions could be incorporated, doing so would introduce additional complexity and calibration challenges that may limit deployability in operational healthcare settings. By abstracting vehicle dynamics into arc-level energy functions, the framework balances modeling realism with tractability, positioning it as an operational decision-support tool rather than a high-fidelity digital twin of UAS platforms. In addition, extending the framework to incorporate dynamic airspace constraints and authorization processes, such as evolving regulatory approvals, LAANC ceilings, and UTM interactions, would require a fundamentally different modeling layer and is identified as an important direction for future research as regulatory frameworks and operational infrastructures continue to mature. More dynamic communication-aware formulations, including time-varying connectivity constraints or adaptive routing under uncertain command-and-control (C2) conditions, likewise represent a complementary extension beyond the day-ahead planning scope considered here. Future work could also examine sensitivity to alternative spatial interpolation schemes beyond the Voronoi-based mapping used in this study, as well as to denser sensing infrastructures for weather and communication data.

Several additional directions for future research emerge from this work. Extending the framework to dynamic or multi-stage formulations, including stochastic, robust, or rolling-horizon models would enable adaptive routing and energy-management decisions under evolving information. In parallel, tailored solution methodologies, such as decomposition-based exact approaches and heuristics, could support larger-scale networks or applications that require near-real-time decision-making. These extensions are particularly relevant for dense urban logistics systems with tighter operational coupling, more complex airspace constraints, and higher demand intensity.

Moreover, the present study focuses on short-term operational planning in settings where fleet ownership and long-term investment decisions are exogenous, and transportation services are often provided by third-party logistics operators. Hence, fixed vehicle acquisition costs are not modeled explicitly because the objective focuses on the operational costs associated with utilizing an available fleet during the planning horizon. In contexts where fleets are owned and operated internally (such as large hospital systems, public health agencies, emergency medical services, or vertically integrated healthcare logistics networks), explicit integration of fleet sizing and capital investment decisions alongside operational routing represents a promising avenue for future work.

Beyond its application to the Shield Illinois case study, the proposed framework serves as a generalizable decision-support tool for governments and health systems seeking to modernize medical logistics. Its design is particularly well-suited for rural settings, where infrastructure constraints, sparse populations, and limited access to medical care create unique delivery challenges. The model's assumptions, such as uncongested roads, unrestricted airspace, and limited built environments, reflect rural conditions, making the framework both realistic and implementable in such regions. As climate variability and healthcare demands increase, adaptive delivery systems like ours will be crucial for enhancing global health security and promoting equitable access to healthcare in underserved rural areas. Finally, the framework is readily extensible to other time-critical domains, including disaster response, humanitarian relief, and agricultural logistics, where operational feasibility, energy efficiency, and environmental conditions are central to system performance.

CRedit authorship contribution statement

Pardis Bahmani: Writing – review & editing, Writing – original draft, Visualization, Methodology, Formal analysis, Conceptualization; **Shakiba Enayati:** Writing – review & editing, Supervision, Methodology, Funding acquisition, Formal analysis, Conceptualization; **Sina Ansari:** Writing – review & editing, Visualization, Formal analysis, Data curation; **Srikanth Gururajan:** Methodology; **Ziteng Wang:** Data curation, Writing – review & editing.

Acknowledgments

This material is based upon work supported by the [National Science Foundation](#) under Grant No. [CMMI-2347150](#). Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. The case study presented in this work was supported by SHIELD Illinois (NIU Grant No. G3A63106, No. G3A63118). Preliminary work contributing to this research was supported by the U.S. Department of Transportation through the University of Nebraska-Lincoln Mid-America Transportation Center for Transportation Safety and Equity (MATC-TSE) under contract number 69A3552348307.

Data availability

Data will be made available on request.

Appendix A. C2 connectivity screening results

This appendix presents spatial visualizations and summary statistics associated with the command-and-control (C2) connectivity screening applied to drone arcs in the case study. Fig. A.1 illustrates the spatial distribution of FCC LTE coverage together with the resulting C2-feasible and -infeasible drone arcs for each operating scenario. Applying this preprocessing approach, approximately 38% of the drone arcs in the distributed LV–LW network and about 30% of the drone arcs in the remaining network scenarios are pruned.




Fig. A.1. Command-and-control (C2) connectivity screening for the four operating scenarios. Green arcs indicate arcs retained after preprocessing, while red arcs indicate the removed arcs due to insufficient LTE coverage. The coverage shown reflects the availability of FCC broadband data collection and is used solely for preprocessing; the figures do not represent optimized routes or operational flight plans.

Appendix B. UAS platform characterization and flight-energy modeling framework

We describe a representative fixed-wing UAS platform considered as a candidate for field deployment. The platform belongs to the class of lightweight fixed-wing UAS suitable for long-range medical logistics missions in rural environments. Fig. B.1 illustrates a representative UAS platform used to contextualize the level of vehicle abstraction adopted in this study.

The representative UAS is constructed using lightweight composite materials (e.g., fiberglass and carbon fiber) and powered by brushless electric propulsion systems with lithium-polymer battery packs. The platform is expected to be equipped with onboard avionics and sensors enabling autonomous waypoint navigation, with flight control implemented through widely used autopilot systems (e.g., Pixhawk-class controllers) (Quidiello and Gururajan, 2024). Key physical and operational characteristics of the candidate UAS are summarized in Table B.1.



Fig. B.1. Representative fixed-wing long-range UAS platform considered for operational deployments.

Table B.1

Representative physical and operational parameters of the candidate fixed-wing UAS.

Parameter	Value
Weight	~28 lb
Wing chord length	14 in
Wingspan	104 in
Fuselage length	84 in
Operational altitude limit	< 400 ft AGL
Payload volume	~500 in ³
Takeoff/landing distance	~100 ft

To provide further insight into the aerodynamic characteristics of the platform, a first-order approximate aerodynamic model was developed around a steady, wings-level cruise condition. This cruise condition represents the dominant flight regime during sample collection missions between pickup locations.

The aerodynamic coefficients were obtained through preliminary analysis of the aircraft geometry using the *Advanced Aircraft Analysis* (AAA) software suite (DARcorporation, n.d.). Representative stability and control derivatives are reported in Table B.2.

Table B.2

Representative aerodynamic stability and control derivatives of the candidate UAS.

Category	Coefficient	Value
Drag	c_{D0}	≈ 0.054
	c_{Da}	≈ 0.165
Lift	c_{La}	≈ 5.23
	$c_{L\dot{\alpha}}$	≈ 0.92
	c_{Lq}	≈ 15.04
	$c_{L\delta_s}$	≈ 0.096
Pitching moment	$c_{m\alpha}$	≈ -4.94
	$c_{m\dot{\alpha}}$	≈ -3.77
	c_{mq}	≈ -24.17
	$c_{m\delta_s}$	≈ -0.40

The negative value of the pitching moment derivative ($c_{m\alpha} < 0$) indicates that the aircraft is statically stable about the cruise operating point. This stability property is relevant for operational energy modeling, as marginally stable or unstable configurations would require increased control effort to maintain steady flight, thereby increasing power consumption. In contrast, a stable platform allows throttle inputs to dominate power usage during cruise, which supports the use of arc-based, cruise-energy approximations in the optimization model.

Building on this stability characterization, we adopt a simplified but conservative energy modeling structure that reflects typical flight profiles of long-endurance UAS platforms used in medical logistics. Each flight segment between two service locations is decomposed into five phases: takeoff, climb, cruise, descent, and landing. Consistent with fixed-wing and hybrid UAS operations in routine delivery missions, we assume that cruise flight dominates the mission duration and energy expenditure. Let E_{cruise} denote the energy consumed during the cruise phase. The total energy consumption of a flight segment is modeled as:

$$E_{\text{total}} = E_{\text{cruise}} + E_{\text{to}} + E_{\text{c}} + E_{\text{d}} + E_{\text{l}}, \quad (\text{B.1})$$

where E_{to} , E_c , E_d , E_l represent the energy consumed during takeoff, climb, descent, and landing, respectively. Following conservative assumptions commonly adopted in the UAS logistics and operations research literature, we approximate the aggregate energy of non-cruise phases as a fixed fraction of cruise energy:

$$E_{total} = 1.1 E_{cruise}. \quad (B.2)$$

This approximation intentionally biases energy estimates upward, ensuring that routing decisions remain feasible under realistic operational conditions. Furthermore, we assume that flight operations are conducted under nominal maneuvering conditions without aggressive acceleration or high-angle-of-attack maneuvers, which is consistent with routine medical delivery missions and the intended deployment context of the candidate UAS platforms. These assumptions provide a tractable and conservative abstraction of UAS energy consumption that is suitable for network-level routing and scheduling optimization while remaining aligned with physically plausible flight behavior.

Appendix C. Weather data integration

To incorporate localized weather effects into our case study, we identified six weather stations located within the broader service region, which are listed in Table C.1. We then partitioned the area into six Voronoi regions, each defined by proximity to one of the weather stations. This spatial segmentation allowed each logistics node to be mapped to its nearest station, enabling region-specific modeling of drone performance under local wind conditions.

We sourced historical hourly weather data from multiple providers, including AccuWeather and Weather.com. After evaluating data quality and consistency, we selected Visual Crossing as our primary source due to its high-resolution records aggregated from reputable sources such as the National Weather Service (NWS) and the Deutscher Wetterdienst (DWD). To validate the reliability of wind data, we compared Visual Crossing outputs with local weather station records. The datasets were highly consistent, with a systematic one-hour time offset observed across all stations. In cases where local station data was missing or incomplete, Visual Crossing was used to fill the gaps. Minor discrepancies (e.g., wind speed values of 0 vs. 0.1 or 0.2) were retained to maintain data fidelity. For larger missing intervals, we applied time-based interpolation and flagged the imputed values during preprocessing.

Table C.1 lists the six weather stations used in our analysis, including their Visual Crossing identifiers and approximate pairwise distances with the official local weather stations.

Table C.1
Weather station mapping to visual crossing locations.

Official Weather Station Location	Visual Crossing Location	ID	Distance Between the Weather Stations (miles)
Cahokia St Louis AP, IL	Cahokia, East St Louis, IL	St3960	2
Peoria GTR Peoria AP, IL	N Peoria Ave, Peoria, IL	St14842	6
St Charles Co AP, MO	Airport Rd, St Charles, MO	St53904	10
Springfield Capital AP, IL	E Capitol Ave, Springfield, IL	St93822	4
Quincy RGNL AP, IL	Quincy, IL	St93989	11
Champaign Willard AP, IL	Airport Rd, Urbana, IL	St94870	9


Appendix D. Sensitivity of average delivery time to vehicle operating costs

This appendix reports the sensitivity of average delivery time to variations in drone and car operating costs across the four operating scenarios.


Fig. D.1 reports average delivery time as a function of drone operating cost across the four operating scenarios. Delivery-time trends remain largely stable over wide cost ranges and change only when the cost-optimal delivery configuration shifts. In low-volume scenarios, delivery time varies in a more discrete manner, while in spatially dispersed settings changes occur more gradually across cost multipliers.

Fig. D.2 illustrates the sensitivity of average delivery time to variations in car operating cost across the four operating scenarios. In high-volume scenarios, average delivery time decreases monotonically as car operating cost increases until a lower bound is reached, beyond which further cost increases result in limited additional change. Across these scenarios, delivery-time reductions of approximately 76% and 71% are observed in the HV–HW and HV–LW cases, respectively.


In low-volume scenarios, delivery-time responses to car operating cost exhibit more discrete behavior. In the LV–HW case, reductions in car operating cost are associated with abrupt increases in average delivery time of approximately 70%. In contrast, in the LV–LW scenario, delivery times fluctuate within a narrow range of approximately 1.5 h across a broader span of cost multipliers.




(a) LV-LW operating scenario



(b) LV-HW operating scenario



(c) HV-LW operating scenario



(d) HV-HW operating scenario

Fig. D.1. Sensitivity of average delivery time to changes in drone operating cost across the four operating scenarios. Shaded regions indicate the cost-optimal delivery configuration.




Fig. D.2. Sensitivity of average delivery time to changes in car operating cost across the four operating scenarios. Shaded regions indicate the cost-optimal delivery configuration.

References

- Alyassi, R., Khonji, M., Karapetyan, A., Chau, S. C.-K., Elbassioni, K., Tseng, C.-M., 2023. Autonomous recharging and flight mission planning for battery-operated autonomous drones. *IEEE Trans. Autom. Sci. Eng.* 20 (2), 1034–1046. <https://doi.org/10.1109/TASE.2022.3175565>
- Backman, G., Hunt, P., Khosla, R., Jaramillo-Strouss, C., Fikre, B.M., Rumble, C., Pevalin, D., Pérez, D.A., Pineda, M.A., Frisancho, A., Tarco, D., Motlagh, M., Farcasanu, D., Vladescu, C., 2008. Health systems and the right to health: an assessment of 194 countries. *Lancet* 372 (9655), 2047–2085. [https://doi.org/10.1016/S0140-6736\(08\)61781-X](https://doi.org/10.1016/S0140-6736(08)61781-X)
- Barlow, Z., Jashami, H., Sova, A., Hurwitz, D.S., Olsen, M.J., 2019. Policy processes and recommendations for unmanned aerial system operations near roadways based on visual attention of drivers. *Trans. Res. C Emerg. Technol.* 108, 207–222. <https://doi.org/10.1016/j.trc.2019.09.012>
- Beigi, P., Rajabi, M.S., Aghakhani, S., 2022. An Overview of Drone Energy Consumption Factors and Models. pp. 1–20. https://doi.org/10.1007/978-3-030-72322-4_200-1
- Betti Sorbelli, F., 2024. UAV-based delivery systems: a systematic review, current trends, and research challenges. *ACM J. Auton. Transp. Syst.* <https://doi.org/10.1145/3649224>
- Bocewicz, G., Radzki, G., Nielsen, P., Banaszak, Z., 2022. Uavs' dynamic routing, subject to time windows variation. *IFAC-Pap.* 55 (2), 457–462. 14th IFAC Workshop on Intelligent Manufacturing Systems IMS 2022. <https://doi.org/10.1016/j.ifacol.2022.04.236>
- Bosona, T., 2020. Urban freight last mile logistics-challenges and opportunities to improve sustainability: a literature review. *Sustainability* 12, 8769. <https://doi.org/10.3390/su12218769>
- Chen, H., Hu, Z., Solak, S., 2021. Improved delivery policies for future drone-based delivery systems. *Eur. J. Oper. Res.* 294. <https://doi.org/10.1016/j.ejor.2021.02.039>
- Chen, M., Smyth, A.W., Giometto, M.G., Li, M.Z., 2023. Drone delivery routing with stochastic urban wind. In: 2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC), pp. 2260–2267. <https://doi.org/10.1109/ITSC57777.2023.10422533>
- Cheng, C., Adulyasak, Y., Rousseau, L.-M., 2024. Robust drone delivery with weather information. *Manuf. Serv. Oper. Manag.* 26. <https://doi.org/10.1287/msom.2022.0339>
- DARcorporation, . Advanced aircraft analysis (AAA) software package, version 5.0. Computer software. Accessed January 30, 2026. <https://www.darcorp.com>
- Dorling, K., Heinrichs, J., Messier, G.G., Magierowski, S., 2017. Vehicle routing problems for drone delivery. *IEEE Trans. Syst. Man Cybern. Syst.* 47 (1), 70–85. <https://doi.org/10.1109/TSMC.2016.2582745>
- Dukkanci, O., Kara, B.Y., Bektaş, T., 2021. Minimizing energy and cost in range-limited drone deliveries with speed optimization. *Transp. Res. Part C Emerging Technol.* 125, 102985. <https://doi.org/10.1016/j.trc.2021.102985>
- Enayati, S., Li, H., Campbell, J., Pan, D., 2023. Multimodal vaccine distribution network design with drones. *Transp. Sci.* 57. <https://doi.org/10.1287/trsc.2023.1205>
- Evolved Universal Terrestrial Radio Access (E-UTRA); 2023. Requirements for Support of Radio Resource Management. Technical Report TS 36.133. 3GPP.
- Federal Communications Commission, 2025. Fcc broadband data collection: mobile broadband coverage. Accessed: January 2026. <https://broadbandmap.fcc.gov>
- Fontanesi, G., Zhu, A., Arvaneh, M., Ahmadi, H., 2023. A transfer learning approach for UAV path design with connectivity outage constraint. *IEEE Internet Things J.* 10 (6), 4998–5012. <https://doi.org/10.1109/JIOT.2022.3220981>
- Goldschmidt, K., Tasko, K., 2022. The downstream effects of the COVID-19 pandemic: the supply chain failure, a wicked problem. *J. Pediatr. Nurs.* 65, 29–32. <https://doi.org/10.1016/j.pedn.2022.04.001>

- Griffith, E., Schurer, J.M., Mawindo, B., Kwibuka, R., Turibyarive, T., Amuguni, J., 2023. The use of drones to deliver rift valley fever vaccines in Rwanda: perceptions and recommendations. *Vaccines* 11, 605. <https://doi.org/10.3390/vaccines11030605>
- Grote, M., Oakley, A., Pilko, A., Krol, J., Blakesley, A., Cherrett, T., Scanlan, J., Anvari, B., Martinez-Sykora, A., 2024. The effects of costs on drone uptake in multi-modal logistics systems within a healthcare setting. *Transp. Econ. Manag.* 2, 58–75. <https://doi.org/10.1016/j.team.2024.03.001>
- Guerriero, F., Surace, R., Loscri, V., Natalizio, E., 2014. A multi-objective approach for unmanned aerial vehicle routing problem with soft time windows constraints. *Appl. Math. Model.* 38 (3), 839–852. <https://doi.org/10.1016/j.apm.2013.07.002>
- Gunaratne, K., Thibbotuwawa, A., Vasegaard, A.E., Nielsen, P., Perera, H.N., 2022. Unmanned aerial vehicle adaptation to facilitate healthcare supply chains in low-income countries. *Drones* 6 (11). <https://doi.org/10.3390/drones6110321>
- Gürel, Ö., Serdarasan, S., 2024. Drone-assisted last-mile delivery under windy conditions: zero pollution solutions. *Smart Cities* <https://api.semanticscholar.org/CorpusID:274002149>.
- Jeon, H., Lucarelli, C., Mazarati, J., Ngabo, D., Song, H., 2022. Leapfrogging for last-mile delivery in health care. *SSRN Electron. J.* <https://doi.org/10.2139/ssrn.4214918>
- Kim, K., Kim, S., Kim, J., Jung, H., 2024. Drone-assisted multimodal logistics: trends and research issues. *Drones* 8 (9). <https://doi.org/10.3390/drones8090468>
- Kim, S., Kwak, J.H., Oh, B., Lee, D.-H., Lee, D., 2021. An optimal routing algorithm for unmanned aerial vehicles. *Sensors* 21 (4). <https://doi.org/10.3390/s21041219>
- Koshita, N., Devi, Y., Chauhan, C., 2024. Evaluating barriers to the adoption of delivery drones in rural healthcare supply chains: preparing the healthcare system for the future. *IEEE Trans. Eng. Manag.* 71, 13096–13108. <https://doi.org/10.1109/TEM.2022.3210121>
- Kruk, M.E., Gage, A.D., Arsenault, C., Jordan, K., Leslie, H.H., Roder-DeWan, S., Adeyi, O., Barker, P., Daelmans, B., Doubova, S.V., English, M., García-Elorrio, E., Guanais, F., Gureje, O., Hirschhorn, L.R., Jiang, L., Kelley, E., Lemango, E.T., Liljestrand, J., Malata, A., Marchant, T., Matoso, M.P., Meara, J.G., Mohanan, M., Ndiaye, Y., Norheim, O.F., Reddy, K.S., Rowe, A.K., Salomon, J.A., Thapa, G., Twum-Danso, N. A.Y., Pate, M., 2018. High-quality health systems in the sustainable development goals era: time for a revolution. *Lancet Glob. Health* 6 (11), e1196–e1252. [https://doi.org/10.1016/S2214-109X\(18\)30386-3](https://doi.org/10.1016/S2214-109X(18)30386-3)
- Lakhwani, T.S., Sinjana, Y., Kapoor, A.P., 2025. Erosion-enhanced rain optimization algorithm for efficient blood bag transport in multimodal healthcare logistics. *Int. J. Transp. Sci. Technol.* <https://doi.org/10.1016/j.ijst.2025.05.002>
- Lamprey, E., Serwaa, D., 2020. Hightech and innovation journal the use of zipline drones technology for COVID-19 samples transportation in Ghana. *HighTech Innov. J.* 1. <https://doi.org/10.28991/HIJ-2020-01-02-03>
- Liu, Z., Yu, B., Chen, T., Zhang, L., 2025. Integrating battery-related decisions into truck-drone tandem delivery problem with limited battery resources. *Transp. Res. C Emerg. Technol.* 174, 105082. <https://doi.org/10.1016/j.trc.2025.105082>
- Macrina, G., Di Puglia Pugliese, L., Guerriero, F., Laporte, G., 2020. Drone-aided routing: a literature review. *Transp. Res. C Emerg. Technol.* 120, 102762. <https://doi.org/10.1016/j.trc.2020.102762>
- Moadab, A., Farajzadeh, F., Valilai, O., 2022. Drone routing problem model for last-mile delivery using the public transportation capacity as moving charging stations. *Sci. Rep.* 12. <https://doi.org/10.1038/s41598-022-10408-4>
- Pinto, R., Lagorio, A., 2022. Point-to-point drone-based delivery network design with intermediate charging stations. *Transp. Res. C Emerg. Technol.* 135, 103506. <https://doi.org/10.1016/j.trc.2021.103506>
- Quidiello, A., Gururajan, S., 2024. Development of virtual sensors using a compact computer cluster for fault tolerant control on UAVs. <https://doi.org/10.2514/6.2024-2904>
- Radzki, G., THIBBOTUWAWA, A., BOCEWICZ, G., 2019. Uavs flight routes optimization in changing weather conditions - constraint programming approach. *Appl. Comput. Sci.* 15 (3), 5–20. <https://doi.org/10.23743/acs-2019-17>
- Ramos, T. R.P., Vigo, D., 2023. A new hybrid distribution paradigm: integrating drones in medicines delivery. *Expert Syst. Appl.* 234, 120992. <https://doi.org/10.1016/j.eswa.2023.120992>
- Rodrigues, T.A., Patrikar, J., Oliveira, N.L., Matthews, H.S., Scherer, S., Samaras, C., 2022. Drone flight data reveal energy and greenhouse gas emissions savings for very small package delivery. *Patterns* 3 (8), 100569. <https://doi.org/10.1016/j.patter.2022.100569>
- Ruehle, F., 2020. Data science applications to string theory. *Phys. Rep.* 839, 1–117. <https://doi.org/10.1016/j.physrep.2019.09.005>
- Schermer, D., Moeini, M., Wendt, O., 2019. A matheuristic for the vehicle routing problem with drones and its variants. *Transp. Res. C Emerg. Technol.* 106, 166–204. <https://doi.org/10.1016/j.trc.2019.06.016>
- Shi, Y., Lin, Y., Li, B., Yi Man Li, R., 2022. A bi-objective optimization model for the medical supplies' simultaneous pickup and delivery with drones. *Comput. Ind. Eng.* 171, 108389. <https://doi.org/10.1016/j.cie.2022.108389>
- Stierlin, N., Risch, M., Risch, L., 2024. Current advancements in drone technology for medical sample transportation. *Logistics* 8 (4). <https://doi.org/10.3390/logistics8040104>
- The U.S. health service administration, <https://toolkit.ncats.nih.gov/glossary/underserved-group/#:-:text=The%20U.S.%20Health%20Services%20Administration,or%20a%20high%20elderly%20population.>
- Thibbotuwawa, A., Bocewicz, G., Nielsen, P., Banaszak, Z., 2019a. Planning deliveries with UAV routing under weather forecast and energy consumption constraints. *IFAC-Pap.* 52 (13), 820–825. <https://doi.org/10.1016/j.ifacol.2019.11.231>
- Thibbotuwawa, A., Bocewicz, G., Radzki, G., Nielsen, P., Banaszak, Z., 2020. UAV mission planning resistant to weather uncertainty. *Sensors* 20 (2). <https://doi.org/10.3390/s20020515>
- Thibbotuwawa, A., Bocewicz, G., Zbigniew, B., Nielsen, P., 2019b. A solution approach for UAV fleet mission planning in changing weather conditions. *Appl. Sci.* 9 (19). <https://doi.org/10.3390/app9193972>
- Troudi, A., Addouche, S.-A., Dellagi, S., Mhamedi, A.E., 2018. Sizing of the drone delivery fleet considering energy autonomy. *Sustainability* 10 (9). <https://doi.org/10.3390/su10093344>
- Wolfinger, D., Tricoire, F., Doerner, K.F., 2019. A matheuristic for a multimodal long haul routing problem. *EURO J. Transp. Logist.* 8 (4), 397–433. <https://doi.org/10.1007/s13676-018-0131-1>
- Zeng, Y., Zhang, R., 2017. Energy-efficient UAV communication with trajectory optimization. *IEEE Trans. Wirel. Commun.* 16 (6), 3747–3760.
- Zhang, G., Jia, N., Zhu, N., Adulyasak, Y., Ma, S., 2023. Robust drone selective routing in humanitarian transportation network assessment. *Eur. J. Oper. Res.* 305 (1), 400–428. <https://doi.org/10.1016/j.ejor.2022.05.046>
- Zhang, J., Campbell, J., Sweeney, II, D., Hupman, A., 2021. Energy consumption models for delivery drones: a comparison and assessment. *Transp. Res. D Transp. Environ.* 90, 102668. <https://doi.org/10.1016/j.trd.2020.102668>