

Article

Bayesian Network-Driven Demand Prediction and Multi-Trip Two-Echelon Routing for Fleet-Constrained Metropolitan Logistics

Ming Liu , Xiangye Yao and Lihua Sun * 

School of Economics & Management, Tongji University, Shanghai 200092, China; mingliu@tongji.edu.cn (M.L.); 1750578@tongji.edu.cn (X.Y.)

* Correspondence: sunlihua@tongji.edu.cn

Abstract

Urban logistics in metropolitan areas faces mounting pressure to deliver faster while controlling operational costs under strict fleet size constraints. Traditional vehicle routing models assume unlimited vehicle availability, overlooking realistic fleet utilization and spatial-temporal demand imbalances. This paper introduces the fleet-constrained metropolitan logistics problem (FCMLP), a novel framework integrating trunk linehaul scheduling, two-echelon routing, multi-trip operations, and anticipatory fleet positioning. We model the FCMLP as a Markov Decision Process capturing the stochastic and dynamic nature of metropolitan delivery flows. Our solution framework combines interpretable Bayesian Network-based demand forecasting for transparent proactive vehicle relocation decisions, parameterized cost-function approximation for dynamic order-to-linehaul assignment, and Adaptive Large Neighborhood Search for multi-trip vehicle routing. Computational experiments on synthetic instances and real-world data from a major e-commerce platform in Jakarta demonstrate 20–26% total cost reduction. Multi-trip operations alone reduce fleet size by 23%, while interpretable predictive relocation further improves performance by 7% through a 20% reduction in emergency deployments. The framework's interpretability enhances operator trust and facilitates practical adoption, offering logistics platforms a path to improve vehicle utilization through operational efficiency and transparent predictive intelligence without expanding fleet size.



Academic Editor: Marco Bortolini

Received: 20 October 2025

Revised: 17 November 2025

Accepted: 25 November 2025

Published: 28 November 2025

Citation: Liu, M.; Yao, X.; Sun, L. Bayesian Network-Driven Demand Prediction and Multi-Trip Two-Echelon Routing for Fleet-Constrained Metropolitan Logistics. *Appl. Sci.* **2025**, *15*, 12609. <https://doi.org/10.3390/app152312609>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: two-echelon routing; multi-trip vehicles; fleet constraints; dynamic order assignment; Markov Decision Process; Bayesian networks

1. Introduction

The rapid growth of e-commerce has transformed urban logistics operations in metropolitan areas worldwide. With global parcel volumes growing over 50% annually, logistics providers face mounting pressure to deliver faster while managing operational costs. This challenge is particularly acute in densely populated metropolitan regions, where multiple cities interconnect through complex transportation networks. While major e-commerce platforms have established sophisticated logistics networks, they face increasing pressure to optimize fleet utilization as vehicle acquisition and maintenance represent substantial capital investments. The dynamic nature of e-commerce demand, with characteristic spatial-temporal imbalances across urban districts, further complicates efficient resource allocation and necessitates intelligent fleet management strategies.

Efficient fleet utilization has emerged as a critical operational priority. While existing literature often assumes unlimited vehicle availability, practical operations require careful

management of fleet resources. For large operators, optimizing fleet size yields significant cost savings and environmental benefits. For small and medium-sized providers, working within fixed fleet constraints is necessity-driven by capital limitations. This convergence creates a universal need for strategies maximizing existing vehicle asset productivity. The challenge extends beyond static optimization to encompass dynamic adaptation, where predictive capabilities and proactive resource positioning become essential for maintaining service quality while ensuring transparent reasoning that operators can understand and trust.

Two-echelon distribution structures, commonly adopted in metropolitan areas, feature parcels flowing from suppliers to city hubs, then to satellite facilities, before final customer delivery. This hub-satellite-customer structure enables consolidation for efficient middle-mile transport and deploys smaller vehicles for congested urban streets. The two-echelon vehicle routing problem (2E-VRP) has been comprehensively reviewed by Sluijk et al. [1]. By placing reload points near customer clusters, satellites dramatically reduce reload time, synergistically enabling multi-trip routing where vehicles complete multiple pickup and delivery cycles within operational periods.

Multi-trip mechanisms offer powerful solutions for improving fleet efficiency. By relaxing traditional “one vehicle, one route per day” assumptions, the Multi-Trip Vehicle Routing Problem (VRPMT) enables operators to serve more customers with existing fleets, as surveyed by Cattaruzza et al. [2]. Recent research has integrated multi-trip operations into two-echelon models [3,4]. However, implementing multi-trip operations within metropolitan logistics presents complex coordination challenges requiring sophisticated planning that considers vehicle capacity, time windows, and dynamic fleet repositioning, with effectiveness relying on providing clear operational decision explanations.

This paper addresses these challenges by proposing an integrated optimization framework for metropolitan logistics that maximizes vehicle utilization through multi-trip strategies and predictive fleet management. We formalize this as the *fleet-constrained metropolitan logistics problem* (FCMLP), modeling it as a Markov Decision Process (MDP) that captures stochastic order arrivals and sequential decision-making. Our approach integrates interpretable Bayesian Network-based demand prediction for proactive relocation, parameterized Cost Function Approximation (CFA) for dynamic order assignment, and Adaptive Large Neighborhood Search (ALNS) for multi-trip routing. To the best of our knowledge, this is the first work integrating (i) explicit fleet constraints in two-echelon multi-trip frameworks; (ii) interpretable Bayesian Network-based prediction for proactive relocation; (iii) dynamic order-to-linehaul assignment with multi-trip delivery under stochastic demand; and (iv) unified MDP formulation capturing component interactions.

The main contributions are (1) formulating the FCMLP with explicit fleet constraints and multi-trip operations while incorporating interpretable predictive analytics; (2) developing an adaptive MDP-based framework coordinating Bayesian Networks for demand prediction, parameterized cost functions for order assignment, enhanced ALNS for multi-trip routing, and hierarchical Bayesian optimization; and (3) demonstrating through experiments on synthetic instances and real-world Jakarta e-commerce data that multi-trip operations achieve 22.6% fleet reduction and 20.5% cost reduction, with proactive relocation providing additional benefits.

The paper is organized as follows: Section 2 reviews relevant literature. Section 3 defines the problem and mathematical formulation. Section 4 details our solution methodology. Section 5 presents computational experiments. Section 6 discusses findings and future research.

2. Literature Review

Modern metropolitan logistics complexity demands sophisticated optimization frameworks. This review analyzes three critical streams: (1) multi-trip vehicle routing under fleet constraints; (2) two-echelon vehicle routing with multi-trip capabilities; and (3) dynamic stochastic dispatching. We systematically examine these domains to identify significant gaps at their intersection.

2.1. Multi-Trip Vehicle Routing Under Fleet Constraints

The VRPMT relaxes single-tour constraints, allowing vehicles to perform subsequent trips after reloading. Cattaruzza et al. [2] provide foundational VRPMT formulation overviews. Recent research increasingly recognizes explicit fleet constraints as operators optimize fleet size. Hernandez et al. [5] demonstrated significant fleet reductions through multi-trip strategies. Zhen et al. [6] extended this to multi-depot settings with release dates, combining features “seldom considered simultaneously.”

Strategic fleet planning integration has emerged as crucial. Fermín Cueto et al. [7] divide problems into strategic (depot allocation, fleet sizing) and operational planning for robustness. Chen et al. [8] address heterogeneous fleets under budget constraints, showing 17.54% cost reduction. Karademir et al. [9] demonstrate multi-trip benefits in integrated water-land transportation. Pan et al. [10] incorporate time-dependent speeds through adaptive large neighborhood search.

However, these studies remain limited to single-echelon networks or static planning, unable to capture unique fleet utilization opportunities that two-echelon structures provide through satellites near customer clusters.

2.2. Two-Echelon Routing with Satellite Reload Points

Two-echelon systems enable first-echelon consolidation and second-echelon agility. Sluijk et al. [1] comprehensively review 2E-VRP literature. Satellites near customer clusters reduce reload time versus single-echelon systems, making multiple trips economically attractive. Grangier et al. [11] pioneered two-echelon multi-trip VRP with satellite synchronization.

Satellite concepts evolved beyond traditional transshipment. Enthoven et al. [12] introduce parcel lockers as covering locations. Li et al. [13] address bi-synchronization at origin and destination satellites. Dos Santos et al. [14] use lockers with occasional couriers for last-mile flexibility. Dahimi et al. [15] innovate with mobile satellites serving dual vehicle-transshipment roles.

Multi-trip integration advances include Dumez et al. [3] allowing first-echelon multi-trips with reverse flows, and Lehmann & Winkenbach [4] focusing on second-echelon multi-trips. Tadaros & Kyriakakis [16] handle hierarchical multi-switch systems. Gutierrez et al. [17] incorporate time-dependent travel times. Zamal et al. [18] merge first and last-mile operations. Şahin & Yaman [19] explicitly address limited fleet size in two-echelon systems. Guo et al. [20] study front warehouse models with transshipment constraints.

Despite structural sophistication, these models assume static demand and rarely treat fleet constraints as primary objectives.

2.3. Dynamic and Stochastic Dispatch Under Fleet Limitations

Dynamic modeling addresses uncertainty through Markov Decision Processes. Crainic et al. [21] address capacity planning through stochastic bin packing. Ghilas et al. [22] combine stochastic demands with scheduled lines.

Real-time dispatching requires sophisticated strategies. Klapp et al. [23] formulate Dynamic Dispatch Waves Problems where reactive optimization reduces costs while im-

Table 1. Cont.

Literature	Problem Characteristics				Environment		Operational Features		Analytics			Model	Solution Method
	2E	MT	FC	E2E	DE	SE	RT	MA	BN	PD	PA		
Dynamic and Stochastic Dispatching													
Ghilas et al. [22]						✓						SP	SAA
Klapp et al. [23]					✓		✓					MIP	Heuristic
Van Heeswijk et al. [24]					✓		✓	✓				MDP	ADP
Zamal et al. [27]	✓			✓	✓	✓	✓	✓				MDP	CFA+ALNS
Sluijk et al. [28]	✓					✓						CC-SP	CG
Bayesian Networks													
Garvey et al. [31]									✓			BN	Simulation
Hosseini et al. [32]					✓				✓			DBN-DTMC	Analytical
Liu et al. [33]					✓				✓			R-DBN	SA
Liu et al. [34]									✓			R-DBN	SP
Hosseini & Ivanov [35]					✓				✓			ML-BN	Simulation
Liu et al. [36]					✓				✓			SV-DBN	DC
Attar et al. [37]									✓			BN	Learning
Harikrishnakumar & Nannapaneni [38]									✓	✓		QBN	Quantum
Liu et al. [39]					✓				✓			RA-DBN	QCQP
Our work	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	MDP+BN	CFA+ALNS

Legend: 2E: Two-Echelon; MT: Multi-Trip; FC: Fleet Constraints; E2E: End-to-End Integration; DE: Dynamic Environment; SE: Stochastic Environment; RT: Real-Time Decision; MA: Metropolitan Area Focus; BN: Bayesian Network; PD: Predictive Demand; PA: Proactive Actions; MIP: Mixed Integer Programming; MDP: Markov Decision Process; SP: Stochastic Programming; CC-SP: Chance-Constrained Stochastic Programming; B&P: Branch-and-Price; ALNS: Adaptive Large Neighborhood Search; CFA: Cost Function Approximation; ADP: Approximate Dynamic Programming; LBB: Logic-Based Benders Decomposition; SAA: Sample Average Approximation; CG: Column Generation; DBN: Dynamic Bayesian Network; R-DBN: Robust DBN; ML-BN: Multi-Layer BN; QBN: Quantum BN; RA-DBN: Risk-Averse DBN; SV-DBN: Structure-Variable DBN; DTMC: Discrete-Time Markov Chain; SA: Simulated Annealing; QCQP: Quadratically Constrained Quadratic Programming; DC: Decomposition-and-Clustering.

3. Model

This section presents the mathematical formulation of the FCMLP introduced above. We begin with a detailed description of the system components and operational characteristics, followed by the formulation as a Markov Decision Process (MDP) following the framework introduced by Powell [41]. A complete list of notation is provided in Supplementary Materials S1.

3.1. Problem Description

The FCMLP considers logistics operations across multiple cities $z \in \mathcal{Z}$ within a metropolitan area, where predictive analytics and proactive resource management are essential for maintaining service quality under fleet constraints. The core challenge extends beyond traditional reactive routing to encompass anticipatory decision-making: leveraging interpretable machine learning models to predict spatial-temporal demand patterns and proactively repositioning idle vehicles before demand materializes.

Each stochastic and dynamic order $o \in \mathcal{O}$ is characterized by demand size q_o , pickup location p_o at (X_{p_o}, Y_{p_o}) in origin city z_o^{fm} , and delivery location d_o at (X_{d_o}, Y_{d_o}) in destination city z_o^{lm} . Both locations have hard time windows $[E_{p_o}, L_{p_o}]$ and $[E_{d_o}, L_{d_o}]$, respectively.

The logistics system employs two-echelon vehicle routing within each city. First-echelon vehicles with capacity Q_z^1 operate between hubs and satellites, while second-echelon vehicles with capacity Q_z^2 handle final distribution from satellites to customers. Each city z maintains a limited fleet of $|V_z^2|$ homogeneous second-echelon vehicles distributed across satellites. This fleet size is a strategic decision variable balancing operational costs against service requirements. Routing decisions are made independently within each city.

Vehicle operations require service time φ at each customer/supplier location and reload time ρ at satellites for multi-trip operations, after which vehicles are immediately available for subsequent trips. Second-echelon vehicles start each trip from their current

satellite location, which can change over time through proactive relocation. Pre-scheduled linehaul services $\ell \in \mathcal{L}$ connect city hubs with departure time $\bar{t}_\ell$, capacity Q_ℓ , fixed cost F_ℓ , and variable cost U_ℓ per unit weight.

The dynamic nature of metropolitan demand creates significant spatial–temporal imbalances across satellites. To address this, we incorporate side information $\mathcal{I}_k$, including historical patterns, temporal features, and area characteristics. This feeds into an interpretable Bayesian Network predicting future demand distribution, enabling proactive vehicle relocation that positions resources where they will be needed rather than where they were last used.

To maximize limited fleet productivity, vehicles perform multiple trips within the planning horizon $\mathcal{T} = [0, T]$. Strategic satellite placement near customer clusters reduces reload-related travel compared to central depot returns, making multi-trip operations economically viable.

Operating costs include daily fixed cost c_z^{fixed} per regular vehicle and variable cost c_z^{var} per distance unit. When the regular fleet cannot meet time windows, emergency vehicles are deployed at significantly higher costs ($c^{\text{em, fixed}}$ and $c^{\text{em, var}}$), creating the fundamental trade-off between fleet size and emergency deployment frequency.

The objective is to determine (1) the optimal second-echelon fleet size $|V_z^2|$ for each city and (2) a real-time operational policy coordinating order-to-linehaul assignments and multi-trip routing to minimize the total expected cost over the planning horizon.

Key Assumptions:

1. **Unlimited First-Echelon Fleet Size:** The number of first-echelon vehicles operating between hubs and satellites is assumed to be unlimited. This reflects the fundamental difference in operational characteristics between the two echelons: first-echelon vehicles operate on predictable routes between fixed facilities (hubs and satellites) with consolidated loads, enabling efficient capacity planning and flexible fleet adjustment through third-party logistics providers or on-demand vehicle leasing. This assumption transforms the dynamic first-echelon routing into a sequence of independent static optimization problems, where each epoch's routing decisions are made without tracking vehicle states or availability constraints. In contrast, second-echelon vehicles must navigate dynamic customer-specific routes with strict time windows, requiring dedicated local resources and specialized knowledge of the service area.
2. **Fixed Second-Echelon Fleet During Operations:** The second-echelon fleet size remains constant throughout the operational period. Vehicle acquisition and disposal decisions cannot be made in real-time due to capital investment lead times and labor contracts.
3. **Deterministic Travel Times:** Travel times between locations are known and constant. While this simplification is common in strategic fleet sizing problems and aligns with standard VRP literature, we partially capture real-world traffic variations through time-dependent linehaul capacity adjustments during peak (80% capacity) and off-peak (120% capacity) periods. This approach balances model tractability with operational realism for the core fleet sizing and multi-trip coordination challenges.
4. **No Split Deliveries:** Each order must be handled entirely by a single vehicle. This maintains operational simplicity and package tracking integrity, reflecting standard practice in parcel delivery operations.
5. **Predictable Demand Patterns:** Order arrivals exhibit spatial-temporal patterns that can be learned from historical data and side information. While individual orders remain stochastic, aggregate demand at the satellite level shows sufficient regularity to enable meaningful predictions through our Bayesian Network approach.

6. **Relocation Time Windows:** Vehicle relocations can only occur during idle periods and must be complete before the vehicle's next anticipated assignment. This ensures that proactive positioning does not interfere with active service commitments.

These assumptions balance model realism with analytical tractability, focusing on the critical challenge of maximizing limited second-echelon fleet utilization through integrated multi-trip strategies and predictive fleet management in two-echelon metropolitan logistics networks.

3.2. Markov Decision Process

This section formulates the problem as a hierarchical optimization problem. In the first stage, we determine the second-echelon fleet sizes for each city. In the second stage, we model the operational decisions as an MDP given the fixed fleet sizes.

Stage 1—Strategic Fleet Sizing: Before the operational period begins, determine the second-echelon fleet size $|V_z^2|$ for each city $z \in \mathcal{Z}$. This decision incurs a fixed daily operating cost $C^{\text{fleet}}(|V_z^2|)$ that includes driver wages, insurance, and maintenance. Specifically,

$$C^{\text{fleet}}(|V_z^2|) = \sum_{z \in \mathcal{Z}} c_z^{\text{fixed}} \cdot |V_z^2| \quad (1)$$

Stage 2—Operational Decisions: Given the fleet sizes from Stage 1, we model the dynamic operational decisions as an MDP. The system is observed at discrete decision epochs $k \in \mathcal{K} := \{0, \dots, K\}$, occurring at equidistant time intervals δ between all epochs.

State Variable: At each decision epoch k , the system state captures both operational status and predictive information to enable proactive fleet management. The state variable $S_k \in \mathcal{S}$ is defined as

$$S_k = (t_k, \mathbf{o}_k^{\text{fm}}, \mathbf{o}_k^{\text{mm}}, \mathbf{t}_k^{\text{mm}}, \mathbf{o}_k^{\text{lm}}, \mathbf{t}_k^{\text{lm}}, \mathcal{L}_{zk}, \mathbf{v}_k, \mathcal{I}_k) \quad (2)$$

where the components are organized into the following three categories:

Temporal and Order Information:

- $t_k \in \mathcal{T}$ is the current time at decision epoch k .
- $\mathbf{o}_k^{\text{fm}}$ is a vector of order indices ready for first-mile pickup in each city.
- $\mathbf{o}_k^{\text{mm}}$ and $\mathbf{t}_k^{\text{mm}}$ are vectors of orders awaiting middle-mile linehaul shipment and their arrival times at origin hubs.
- $\mathbf{o}_k^{\text{lm}}$ and $\mathbf{t}_k^{\text{lm}}$ are vectors of orders awaiting last-mile delivery and their arrival times at destination hubs.
- $\mathcal{L}_{zk}$ denotes the set of linehauls departing after time t_k from city z .

Vehicle Status Information:

- $\mathbf{v}_k = \{(\gamma_v^{\text{current}}, \text{status}_v, t_v^{\text{avail}})\}_{v \in V_z^2, z \in \mathcal{Z}}$ represents the status of second-echelon vehicles, where the following apply:
 - $\gamma_v^{\text{current}} \in \mathcal{G}_z$ is the current satellite location of vehicle v .
 - $\text{status}_v \in \{\text{idle}, \text{routing}, \text{relocating}\}$ is the operational status.
 - t_v^{avail} is the time when the vehicle becomes available for new assignments.

Side Information for Predictive Analytics:

- $\mathcal{I}_k = \{\mathcal{I}_{\gamma k}\}_{\gamma \in \mathcal{G}}$ contains satellite-specific side information, where each $\mathcal{I}_{\gamma k} = (H_{\gamma k}, A_{\gamma}, N_{\gamma k}^{\text{nd}}, E_{\gamma k})$ includes the following:
 - $H_{\gamma k} \in \mathbb{R}^+$: historical average demand at satellite γ for similar time periods over the past few days.
 - $A_{\gamma} \in \{\text{Type1}, \text{Type2}, \text{Type3}\}$: satellite area classification.
 - $N_{\gamma k}^{\text{nd}} \in \mathbb{Z}^+$: number of pending next-day orders currently assigned to satellite γ .

- $E_{\gamma k} \in \{0, 1\}$: binary indicator for special events near satellite γ .

This enhanced state representation enables the system to leverage historical patterns and current operational status for anticipatory decision-making, particularly for proactive vehicle positioning based on predicted demand surges.

Decision Variables: The decision $x_k \in \mathcal{X}(S_k)$ comprises the following three key components for each city z :

1. **Order Assignment:** Binary variables $\hat{a}_{o\ell zk} \in \{0, 1\}$ indicating if order $o \in \mathbf{o}_k^{\text{fm}} \cup \mathbf{o}_k^{\text{mm}}$ is assigned to linehaul $\ell \in \mathcal{L}_{zk}$.
2. **Vehicle Dispatch and Routing:** Selecting subset $\mathbf{o}_k^r \subseteq \mathbf{o}_k^{\text{fm}} \cup \mathbf{o}_k^{\text{lm}}$ and determining multi-trip routes for available second-echelon vehicles, with binary variables $\hat{y}_{rzk} \in \{0, 1\}$ indicating if route r_{zk} is executed.
3. **Proactive Vehicle Relocation:** To address predicted demand imbalances across satellites, the system can proactively relocate idle vehicles. Binary variables $\hat{r}_{v\gamma zk} \in \{0, 1\}$ indicate whether idle vehicle v is relocated from its current satellite to satellite γ in city z at epoch k . These relocation decisions are subject to the following operational constraints:
 - Each idle vehicle can be relocated to at most one destination: $\sum_{\gamma \in \mathcal{G}_z} \hat{r}_{v\gamma zk} \leq 1, \forall v \in V_z^{\text{idle}}$.
 - Only vehicles with $\text{status}_v = \text{idle}$ are eligible for relocation.
 - Relocations occur only between satellites within the same city: $\hat{r}_{v\gamma zk} = 0$ if $\gamma \notin \mathcal{G}_z$ where $v \in V_z^2$.
 - To prevent cascading relocations, vehicles relocated in the previous epoch cannot be relocated again: if $\sum_{\gamma} \hat{r}_{v\gamma, z, k-1} = 1$, then $\sum_{\gamma} \hat{r}_{v\gamma zk} = 0$.

Immediate Cost Function: The cost incurred at epoch k given state S_k and decision x_k comprises four components,

$$C(S_k, x_k) = \hat{c}_k^{\text{AC}} + \hat{c}_k^{\text{RC}} + \hat{c}_k^{\text{EC}} + \hat{c}_k^{\text{RLC}} \quad (3)$$

where the cost components are defined as follows:

- $\hat{c}_k^{\text{AC}}$ is the linehaul assignment cost,

$$\hat{c}_k^{\text{AC}} = \sum_{z \in \mathcal{Z}} \left[\sum_{\ell \in \mathcal{L}_{zk}} F_{\ell} \hat{w}_{\ell zk} + \sum_{o \in \mathbf{o}_k^l} \sum_{\ell \in \mathcal{L}_{zk}} q_o U_{\ell} \hat{a}_{o\ell zk} \right] \quad (4)$$

where $\mathbf{o}_k^l = \mathbf{o}_k^{\text{fm}} \cup \mathbf{o}_k^{\text{mm}}$ represents all orders eligible for linehaul assignment, and $\hat{w}_{\ell zk} \in \{0, 1\}$ indicates whether linehaul ℓ is used.

- $\hat{c}_k^{\text{RC}}$ is the regular vehicle routing cost,

$$\hat{c}_k^{\text{RC}} = \sum_{z \in \mathcal{Z}} \sum_{r_{zk} \in R_{zk}^2} c_z^{\text{var}} \cdot \text{dist}(r_{zk}) \cdot \hat{y}_{rzk} \quad (5)$$

where R_{zk}^2 is the set of second-echelon routes in city z at epoch k , and $\text{dist}(r_{zk})$ denotes the total distance of route r_{zk} .

- $\hat{c}_k^{\text{EC}}$ is the emergency transportation cost,

$$\hat{c}_k^{\text{EC}} = \sum_{o \in \mathcal{O}_k^{\text{violated}}} (c^{\text{em, fixed}} + c^{\text{em, var}} \cdot \text{dist}_o^{\text{em}}) \quad (6)$$

where $\mathcal{O}_k^{\text{violated}}$ is the set of orders that cannot be served within time windows by the regular fleet, and $\text{dist}_o^{\text{em}}$ is the emergency vehicle travel distance for order o .

- $\hat{c}_k^{RLC}$ is the proactive relocation cost, representing the operational expense of repositioning vehicles between satellites,

$$\hat{c}_k^{RLC} = \sum_{z \in \mathcal{Z}} \sum_{v \in V_z^{\text{idle}}} \sum_{\gamma \in \mathcal{G}_z} c_z^{\text{var}} \cdot \text{dist}(\gamma_v^{\text{current}}, \gamma) \cdot \hat{r}_{v\gamma zk} \quad (7)$$

where $V_z^{\text{idle}} = \{v \in V_z^2 : \text{status}_v = \text{idle}\}$ represents the set of currently idle vehicles in city z , and $\text{dist}(\gamma_v^{\text{current}}, \gamma)$ is the distance between the vehicle's current satellite and the target satellite.

The routing decision must satisfy the following constraints for fleet size and multi-trip operations:

$$\sum_{r_{zk} \in R_{zk}^2} \hat{y}_{r_{zk}} \leq |V_z^2|, \quad \forall z \in \mathcal{Z} \quad (8)$$

For multi-trip operations, vehicles can perform multiple routes within the planning horizon, subject to time feasibility constraints. A vehicle v can be assigned to a route r only if

$$s_r \geq t_v^{\text{avail}}, \quad \forall r \in \mathcal{R}_{vk}, v \in V_z^2 \quad (9)$$

where $\mathcal{R}_{vk}$ is the set of feasible routes for vehicle v at epoch k , and s_r is the start time of route r . The end time of route r , denoted as e_r , includes the travel time between all nodes in the route plus the service time φ at each customer or supplier location visited. After completing route r , vehicle v becomes available at time

$$t_v^{\text{avail}} = e_r + \rho \quad (10)$$

where ρ is the reload time at the satellite. Furthermore, all routes must be completed within the planning horizon

$$e_r \leq T, \quad \forall r \in R_{zk}^2 \quad (11)$$

When Equations (8), (9), (10), and (11) cannot accommodate all orders within their time windows, the excess orders in $\mathcal{O}_k^{\text{violated}}$ must be served by emergency vehicles at significantly higher costs. This cost differential—with $c^{\text{em, fixed}} \gg c_z^{\text{fixed}}$ and $c^{\text{em, var}} > c_z^{\text{var}}$ —creates the fundamental trade-off: larger regular fleets reduce expensive emergency deployments but increase daily fixed costs, while smaller fleets save on fixed costs but may trigger frequent costly emergency vehicle usage.

Additionally, vehicle operations must respect location constraints

$$\gamma_v^i = \gamma_v^{\text{current}}, \quad \forall v \in V_z^2, \forall \text{ new trips } T_v^i \quad (12)$$

where γ_v^i is the departure satellite for vehicle v 's i -th trip. This ensures vehicles can only start new trips from their current location.

For vehicle relocation, the following constraints apply:

$$\sum_{\gamma \in \mathcal{G}_z} \hat{r}_{v\gamma zk} \leq 1, \quad \forall v \in V_z^{\text{idle}}, z \in \mathcal{Z} \quad (13)$$

ensuring each idle vehicle is relocated to at most one satellite per epoch.

The relocation must complete before the vehicle's next potential assignment,

$$t_k + \frac{\text{dist}(\gamma_v^{\text{current}}, \gamma)}{\text{speed}} \leq t_v^{\text{next}}, \quad \forall v : \hat{r}_{v\gamma zk} = 1 \quad (14)$$

where t_v^{next} is the earliest time the vehicle might be needed based on predicted demand.

Exogenous Information: The exogenous information variable $W_{k+1}(S_k, x_k)$ captures the arrival of new orders $\mathbf{o}_{k+1}^{\text{new}}$ in the system at decision epoch $k + 1$.

Transition Function: The transition function $S_{k+1} = S^M(S_k, x_k, W_{k+1})$ describes the system's evolution, with vehicle status updates now incorporating relocation decisions,

$$\mathbf{v}_{k+1} = \mathcal{U}^v(\mathbf{v}_k, x_k, t_{k+1}, \{\hat{r}_{v\gamma zk}\}) \quad (15)$$

where $\mathcal{U}^v$ updates vehicle locations based on completed routes and relocation decisions. Specifically, if $\hat{r}_{v\gamma zk} = 1$, then

- $\gamma_v^{\text{current}} \leftarrow \gamma$ (update satellite assignment).
- $t_v^{\text{avail}} \leftarrow t_k + \frac{\text{dist}(\gamma_v^{\text{old}}, \gamma)}{\text{speed}}$ (update availability time).
- $\text{status}_v \leftarrow$ relocating during transit, then idle upon arrival.

The status transition follows a two-stage process:

- At epoch k , when relocation is initiated, the following apply:
 - Vehicle remains physically at origin satellite: $\gamma_v^{\text{current}}$ unchanged.
 - Status updates to prevent new assignments: $\text{status}_v \leftarrow$ relocating.
 - Availability time set based on travel time: $t_v^{\text{avail}} \leftarrow t_k + \tau_{\text{reloc}}(\gamma_v^{\text{current}}, \gamma)$.
- At subsequent epoch, when $t_{k'} \geq t_v^{\text{avail}}$, the following apply:
 - Vehicle arrives at destination: $\gamma_v^{\text{current}} \leftarrow \gamma$.
 - Status returns to idle: $\text{status}_v \leftarrow$ idle.
 - Vehicle becomes available for routing from new location.

where $\tau_{\text{reloc}}(\gamma_1, \gamma_2) = \frac{\text{dist}(\gamma_1, \gamma_2)}{\text{speed}}$ represents the travel time between satellites.

Objective Function: The proposed framework seeks to find optimal fleet sizes and an operational policy that minimizes the expected total cost

$$\min_{\{|V_z^2|\}_{z \in \mathcal{Z}}, \pi \in \Pi} \left\{ \sum_{z \in \mathcal{Z}} C^{\text{fleet}}(|V_z^2|) + \mathbb{E} \left[\sum_{k \in \mathcal{K}} C(S_k, X^\pi(S_k)) \middle| S_0, \{|V_z^2|\} \right] \right\} \quad (16)$$

This two-stage stochastic optimization problem captures the essential trade-offs in fleet-constrained metropolitan logistics as follows:

- **First Stage (Strategic Fleet Sizing):** The term $\sum_{z \in \mathcal{Z}} C^{\text{fleet}}(|V_z^2|)$ represents the total fleet operating costs across all cities, where $C^{\text{fleet}}(|V_z^2|) = c_z^{\text{fixed}} \cdot |V_z^2|$ scales linearly with the number of second-echelon vehicles. This captures the fixed daily costs including driver wages, insurance, and maintenance that must be committed before operations begin.
- **Second Stage (Operational Decisions):** The expectation term $\mathbb{E}[\sum_{k \in \mathcal{K}} C(S_k, X^\pi(S_k))]$ represents the expected cumulative operational costs over the planning horizon under policy π . At each epoch k , the immediate cost $C(S_k, x_k)$ comprises the following four components as defined in Equation (3):
 - Linehaul assignment costs $\hat{c}_k^{\text{AC}}$ for middle-mile transportation.
 - Regular vehicle routing costs $\hat{c}_k^{\text{RC}}$ for first and second-echelon operations.
 - Emergency transportation costs $\hat{c}_k^{\text{EC}}$ when regular fleet capacity is exceeded.
 - Proactive relocation costs $\hat{c}_k^{\text{RLC}}$ for anticipatory vehicle repositioning.
- **Stochastic Nature:** The expectation is taken over the stochastic order arrival process W_{k+1} , which captures the uncertainty in demand patterns across time and space. The policy π must adapt to these realizations while working within the constraints imposed by the first-stage fleet size decisions.

This formulation reveals the fundamental challenge: insufficient fleet capacity leads to frequent expensive emergency deployments (high $\hat{c}_k^{EC}$), while excess capacity results in unnecessary fixed costs (high C^{fleet}). The optimal solution balances these competing pressures by determining both the right fleet size and an adaptive operational policy that maximizes the productivity of the chosen fleet through intelligent routing, multi-trip operations, and proactive repositioning.

4. Solution Method

The MDP formulation of the FCMLP requires a solution approach that balances computational tractability with interpretable, proactive decision-making. Figure 1 illustrates the overall system architecture integrating predictive analytics (Bayesian Network, BN) with operational optimization (CFA and ALNS) to support proactive fleet management and efficient multi-trip routing under hard fleet constraints. Figure 2 depicts the integrated decision process within each epoch: the BN predicts spatial-temporal demand, vehicles are proactively relocated, and hierarchical optimization through CFA and ALNS determines order-to-linehaul assignments and two-echelon routes. This modular pipeline keeps real-time decisions computationally manageable while grounding them in transparent predictions.

We use a Bayesian Network for demand prediction because it provides interpretable conditional dependencies that operators can inspect and adjust, while naturally representing prediction uncertainty through probability distributions. The Cost Function Approximation (CFA) of Zamal et al. [27] enables real-time order-to-linehaul decisions by decomposing the stochastic problem into tractable deterministic subproblems parameterized by a few policy coefficients. Adaptive Large Neighborhood Search (ALNS) is adopted for routing due to its robustness on rich VRP variants and its destroy-and-repair paradigm, which is well suited to exploring trip consolidation and reassignment opportunities in multi-trip, two-echelon settings.

Overall, our approach consists of four interconnected components: (i) an interpretable BN for demand prediction using side information, (ii) a proactive vehicle relocation strategy driven by predicted demand imbalances, (iii) a CFA-based policy for dynamic order-to-linehaul assignment, and (iv) an ALNS-based solver for two-echelon routing with multi-trip operations under dynamically adjusted fleet distributions. The remainder of this section details these components and their interactions within the MDP framework.




Figure 1. System architecture of the fleet-constrained metropolitan logistics problem.




Figure 2. MDP decision process flow.

4.1. Demand Prediction and Proactive Vehicle Relocation

This section presents our approach for leveraging side information through an interpretable Bayesian Network to predict demand patterns and make proactive relocation decisions. This predictive-proactive framework represents the key innovation that transforms reactive fleet management into anticipatory resource allocation.

4.1.1. Bayesian Network for Demand Prediction

We employ a Bayesian Network (BN) to model the probabilistic dependencies between side information and future demand patterns at each satellite.

As shown in Figure 3, the network uses a three-layer hierarchical architecture with observable variables (temporal features, historical demand, area type, pending orders, and event indicators) as inputs, two latent variables capturing abstract demand drivers, and predicted demand pressure levels as outputs. The network parameters are learned from historical data using maximum likelihood estimation with Laplace smoothing.




Figure 3. Bayesian Network architecture for satellite-level demand prediction.

For operational decision-making, the probabilistic output is converted into a continuous Demand Pressure Index (DPI),

$$\text{DPI}_{\gamma,k+1} = \kappa_L \cdot P(\mathcal{D}_{\gamma,k+1}^{\text{pred}} = \text{Low}) + \kappa_M \cdot P(\mathcal{D}_{\gamma,k+1}^{\text{pred}} = \text{Medium}) + \kappa_H \cdot P(\mathcal{D}_{\gamma,k+1}^{\text{pred}} = \text{High}) \quad (17)$$

where weights $\kappa_L = 0.2$, $\kappa_M = 0.5$, and $\kappa_H = 0.8$ reflect the relative operational impact of each demand level. The BN provides interpretability through influence analysis and contrastive explanations, enabling operators to understand and trust the relocation decisions. Full technical details, including the probability model, parameter learning procedures, and interpretability mechanisms, are provided in Supplementary Materials S2.

4.1.2. Proactive Relocation Strategy

Based on BN predictions, we implement a threshold-based relocation strategy that balances responsiveness with stability. At each decision epoch k , the strategy evaluates demand-supply imbalances and triggers relocations when beneficial.

Demand-Supply Imbalance Metric: For each satellite γ , we compute the predicted imbalance

$$\mathcal{B}_{\gamma k} = \frac{\text{DPI}_{\gamma,k+1}}{\max(|V_{\gamma}^{\text{available}}|, 1)} \quad (18)$$

where $|V_{\gamma}^{\text{available}}|$ is the number of idle vehicles currently at satellite γ .

System-Wide Imbalance Assessment: We evaluate the system-wide demand distribution to identify relocation opportunities

$$\text{Imbalance Ratio} = \frac{\max_{\gamma \in \mathcal{G}_z} \mathcal{B}_{\gamma k}}{\min_{\gamma \in \mathcal{G}_z} \mathcal{B}_{\gamma k}} \quad (19)$$

Relocation is triggered when this ratio exceeds the threshold η^{reloc} .

Relocation Optimization: When triggered, vehicles are relocated to minimize the expected system imbalance while accounting for relocation costs

$$\min \sum_{v \in V_z^{\text{idle}}} \sum_{\gamma \in \mathcal{G}_z} [c_z^{\text{var}} \cdot \text{dist}(\gamma_v^{\text{current}}, \gamma) - \xi \cdot \Delta \mathcal{B}_{\gamma}(v)] \cdot \hat{r}_{v\gamma zk} \quad (20)$$

where $\Delta \mathcal{B}_{\gamma}(v)$ represents the improvement in system-wide imbalance from relocating vehicle v to satellite γ ,

$$\Delta \mathcal{B}_{\gamma}(v) = \mathcal{B}_{\gamma k} - \frac{\text{DPI}_{\gamma,k+1}}{|V_{\gamma}^{\text{available}}| + 1} + \mathcal{B}_{\gamma_v^{\text{current}},k} - \frac{\text{DPI}_{\gamma_v^{\text{current}},k+1}}{\max(|V_{\gamma_v^{\text{current}}}^{\text{available}}| - 1, 1)} \quad (21)$$

Subject to operational constraints,

$$\sum_{\gamma \in \mathcal{G}_z} \hat{r}_{v\gamma zk} \leq 1, \quad \forall v \in V_z^{\text{idle}} \quad (\text{each vehicle relocated at most once}) \quad (22)$$

$$|V_{\gamma}^{\text{available}}| - \sum_{v: \gamma_v^{\text{current}} = \gamma} \sum_{\gamma' \neq \gamma} \hat{r}_{v\gamma' zk} \geq \nu, \quad \forall \gamma \in \mathcal{G}_z \quad (\text{minimum service level}) \quad (23)$$

$$t_k + \frac{\text{dist}(\gamma_v^{\text{current}}, \gamma)}{\text{speed}} \leq t_v^{\text{next}}, \quad \forall v, \gamma : \hat{r}_{v\gamma zk} = 1 \quad (\text{time feasibility}) \quad (24)$$

Explanation Generation: Each relocation decision is accompanied by an interpretable explanation as follows:

“Vehicle V12 relocated from Satellite A to Satellite B because of the following:

- Predicted demand at B is HIGH (78% confidence) due to evening residential deliveries and 2 pending next-day orders.
- Satellite B has only 2 available vehicles vs. predicted need for 5–6.

- Satellite A has LOW predicted demand (82% confidence) with 4 excess vehicles.
- Relocation distance 12 km, arrival before 14:30 peak period."

4.1.3. Initial Fleet Distribution

At the beginning of the planning horizon ($k = 0$), vehicles must be allocated across satellites to provide balanced coverage. The initial distribution considers both residual next-day orders from the previous operational period and historical demand patterns

$$|V_{\gamma}^2(0)| = \left\lfloor |V_z^2| \cdot \frac{w_{\text{nd}} \cdot N_{\gamma}^{\text{nd}} + w_{\text{hist}} \cdot \bar{H}_{\gamma}^{\text{base}}}{\sum_{\gamma' \in \mathcal{G}_z} (w_{\text{nd}} \cdot N_{\gamma'}^{\text{nd}} + w_{\text{hist}} \cdot \bar{H}_{\gamma'}^{\text{base}})} \right\rfloor + \epsilon_{\gamma} \quad (25)$$

where the following apply:

- N_{γ}^{nd} represents pending next-day orders at satellite γ .
- $\bar{H}_{\gamma}^{\text{base}}$ is the historical average morning demand for the satellite.
- $w_{\text{nd}} = 0.7$ and $w_{\text{hist}} = 0.3$ are weights prioritizing immediate needs over historical patterns.
- $\epsilon_{\gamma} \in \{0, 1\}$ allocates remaining vehicles after floor operation to highest-priority satellites.

This initial distribution ensures that the total fleet size constraint $\sum_{\gamma \in \mathcal{G}_z} |V_{\gamma}^2(0)| = |V_z^2|$ is satisfied while providing coverage aligned with anticipated demand patterns.

4.2. Cost Function Approximation for Linehaul Assignment

At each decision epoch k , we assign orders $\mathbf{o}_{zk}^l = \mathbf{o}_{zk}^{\text{fm}} \cup \mathbf{o}_{zk}^{\text{mm}}$ to available linehails using a parameterized Mixed Integer Programming (MIP) formulation, known as Cost Function Approximation (CFA), where the city index z is now explicitly included to denote city-specific order sets (extending the notation $\mathbf{o}_k^l$ from Section 3). Here, $\mathbf{o}_{zk}^{\text{fm}}$ represents first-mile orders awaiting pickup in city z at epoch k , and $\mathbf{o}_{zk}^{\text{mm}}$ represents middle-mile orders already at the hub awaiting linehaul assignment. The CFA approach focuses on middle-mile transportation between city hubs, anticipating the effectiveness of vehicle routing operations that occur before and after linehaul transport.

To enable effective coordination between linehaul assignment and vehicle routing, we define several critical time points for each order. The latest delivery start time (t_o^{ldst}) represents when a vehicle must depart from the destination hub to meet the delivery time window, accounting for travel to the satellite, reload time, and final delivery. The latest linehaul departure time (t_o^{lldt}) identifies the last feasible linehaul that can deliver the order on time. For first-mile orders assigned to specific linehails, the latest and earliest pickup start times (t_o^{lpst} and t_o^{epat}) determine the temporal constraints for pickup operations, considering synchronization requirements and time windows. These temporal concepts, formally defined in Supplementary Materials S2, establish the feasibility boundaries for order-to-linehaul assignments and subsequent routing decisions.

To capture the spatial and temporal variations in our dynamic system, we extend the notation from Section 3 to include city and time indices where appropriate. For instance, the linehaul capacity Q_{ℓ} is extended to $Q_{\ell zk}$ to reflect potential capacity variations across different cities z and time epochs k . Similarly, fixed costs F_{ℓ} and variable costs U_{ℓ} become $F_{\ell zk}$ and $U_{\ell zk}$, and order weights q_o become q_{ozk} to account for city-specific and time-dependent characteristics. This notation extension allows us to model the heterogeneous operational conditions across different cities and time periods in the metropolitan logistics network.

4.2.1. Modified Assignment Cost with Slack Time

The concept of slack time is introduced to represent the temporal flexibility available for first-mile pickup and last-mile delivery operations. Slack time represents the buffer between when an order could theoretically complete its pickup/delivery operations and when

it must complete them to meet linehaul schedules and customer time windows. Greater slack time provides more routing flexibility in both origin and destination cities, which can help reduce time window violations and the need for emergency vehicle deployments.

For each order o assigned to linehaul ℓ departing at time $\hat{t}_o$, we define origin slack as the time buffer at the origin hub and destination slack as the time buffer at the destination hub. The modified assignment cost incorporating slack time is then given by

$$\tilde{c}_{o\ell zk}^{cfa} = q_{ozk}U_{\ell zk} - \alpha[(s_o^{\text{origin}})^\varepsilon + (s_o^{\text{dest}})^\varepsilon] \quad (26)$$

where q_{ozk} is the weight of order o in city z at epoch k , $U_{\ell zk}$ is the cost per unit weight for linehaul ℓ in city z at epoch k , $\alpha \geq 0$ is a weight parameter controlling the importance of slack time, $\varepsilon \geq 1$ is an exponent that modulates the impact of slack time, and s_o^{origin} and s_o^{dest} represent the origin and destination slack times, respectively (detailed calculations in Supplementary Materials S2). When $\alpha = 0$, we recover pure cost minimization. Positive values of α encourage assignments that provide more operational flexibility for routing.

4.2.2. Deterministic Assignment Formulation

The linehaul assignment problem at epoch k is formulated as

$$\min \sum_{z \in \mathcal{Z}} \left[\sum_{\ell \in \mathcal{L}_{zk}} F_{\ell zk} \hat{w}_{\ell zk} + \sum_{o \in \mathbf{o}_{zk}^I} \sum_{\ell \in \mathcal{L}_{zk}} \tilde{c}_{o\ell zk}^{cfa} \hat{a}_{o\ell zk} \right] \quad (27)$$

$$\text{s.t.} \quad \sum_{o \in \mathbf{o}_{zk}^I} q_{ozk} \hat{a}_{o\ell zk} \leq Q_{\ell zk} \hat{w}_{\ell zk}, \quad \forall \ell \in \mathcal{L}_{zk}, z \in \mathcal{Z} \quad (28)$$

$$\sum_{\ell \in \tilde{\mathcal{L}}_{ozk}} \hat{a}_{o\ell zk} = 1, \quad \forall o \in \mathbf{o}_{zk}^I, z \in \mathcal{Z} \quad (29)$$

$$\hat{a}_{o\ell zk}, \hat{w}_{\ell zk} \in \{0, 1\} \quad (30)$$

where $F_{\ell zk}$ is the fixed cost for utilizing linehaul ℓ in city z at epoch k , $\hat{w}_{\ell zk}$ is a binary variable indicating whether linehaul ℓ is utilized in city z at epoch k , $\hat{a}_{o\ell zk}$ is a binary variable for order-to-linehaul assignment, $Q_{\ell zk}$ is the capacity of linehaul ℓ in city z at epoch k , and $\tilde{\mathcal{L}}_{ozk} \subseteq \mathcal{L}_{zk}$ represents the subset of time-feasible linehauls for order o .

This deterministic formulation balances immediate costs with the temporal flexibility needed for efficient vehicle routing operations, improving overall system performance while maintaining computational tractability.

4.3. Parameterized Adaptive Large Neighborhood Search

This section presents our solution approach for the vehicle routing component of the FCMLP. The routing decisions must be made in real-time at each decision epoch k , taking as input the linehaul assignments determined by the CFA module and producing executable vehicle routes that respect fleet constraints and enable multi-trip operations.

The interface between the CFA and routing modules is carefully designed to maintain computational tractability while ensuring operational feasibility. The CFA module provides two critical inputs to the routing algorithm: (1) the assignment decisions $\hat{a}_{o\ell zk}$ that specify which orders are assigned to which linehauls, and (2) the implied deadlines for first-mile pickups based on the assigned linehaul departure times. These assignments establish hard constraints that the routing algorithm must respect, creating a hierarchical decision structure where tactical linehaul decisions guide operational routing decisions.

Throughout this section, we maintain the city-specific and time-specific notation introduced in the CFA formulation. For instance, operational parameters such as linehaul capacities are denoted as $Q_{\ell zk}$ rather than Q_ℓ , reflecting the heterogeneous operational

conditions across different cities and time periods in the metropolitan logistics network. Additionally, we define the ready-for-routing order set $\mathbf{o}_{zk}^R$ as a subset of orders requiring immediate routing attention, with the specific selection criteria detailed in the dispatching policy.

4.3.1. Algorithm Overview and Positioning

We employ a parameterized Adaptive Large Neighborhood Search (ALNS) algorithm triggered at each decision epoch to solve the fleet-constrained two-echelon vehicle routing problem with multi-trip operations. Extending Zamal et al. [18]’s static framework and building upon Zamal et al. [27]’s dynamic order assignment principles, we incorporate a parameterized dispatching policy where parameter β controls route planning activation timing, balancing routing efficiency with service responsiveness throughout the planning horizon.

The ALNS operates within our broader Markov Decision Process (MDP) framework, serving as the routing optimization module. At each decision epoch k , the algorithm receives a set of orders $\mathbf{o}_{zk}^R$ that have been deemed ready for routing based on the dispatching policy. These orders may include both first-mile pickups and last-mile deliveries, with their linehaul assignments (and corresponding deadlines) already determined by the Cost Function Approximation (CFA) module described in Section 4.2. The integration of multi-trip capabilities allows our ALNS to maximize fleet utilization under the hard constraint of limited vehicle availability.

Before executing the ALNS algorithm, the dynamic order data undergoes preprocessing to transform it into a static problem instance suitable for optimization. This conversion normalizes all temporal data relative to the current decision epoch t_k , ensuring that each ALNS execution begins at time $t = 0$ while maintaining consistency with the broader dynamic system. The detailed preprocessing steps are provided in Supplementary Materials S2.

4.3.2. Parameterized Dispatching Policy

The dispatching policy determines which orders from the system state should be included in the ready-for-routing set $\mathbf{o}_{zk}^R$ at each decision epoch. This policy operates on the post-relocation fleet distribution, ensuring that routing decisions account for the proactively positioned vehicles. The policy is governed by a core parameter β , which represents a time buffer threshold.

$$\mathbf{o}_{zk}^R = \{o \in \mathbf{o}_{zk} : \min(t_o^{\text{lpst}} - t_k, t_o^{\text{ldst}} - t_k) < \beta\} \quad (31)$$

where t_o^{lpst} and t_o^{ldst} are the latest pickup start time and latest delivery start time defined in Section 4.2. The dispatching mechanism is thus triggered if there exists at least one order satisfying this temporal urgency criterion.

Small β values dispatch only urgent orders with limited routing flexibility, while large values enable frequent dispatching with greater consolidation opportunities. When regular fleet routing fails, emergency vehicles ensure service commitments at significantly higher cost, incentivizing efficient regular routing.

The dispatching policy operates on the post-relocation fleet distribution, ensuring that routing decisions account for the proactively positioned vehicles. Vehicles marked as “relocating” (i.e., $\text{status}_v = \text{relocating}$) are excluded from the available vehicle set until their relocation completes at t_v^{avail} . This temporal coordination ensures that routing decisions are based on accurate vehicle availability.

The value of β interacts with the relocation decisions: aggressive relocation (low η^{reloc}) pairs well with conservative dispatching (high β) to allow relocated vehicles time to reach their destinations, while conservative relocation (high η^{reloc}) can accommodate aggressive

dispatching (low β) since vehicles remain near their original positions. This interaction is captured in our hierarchical Bayesian optimization framework.

4.3.3. Core ALNS Framework

The ALNS follows a standard destroy-and-repair framework adapted for fleet-constrained multi-trip operations. The algorithm employs both small and large destroy operations, with small operations focused on local improvements and large operations for diversification. In each iteration, a destroy operator removes a subset of customers from the current solution, followed by a repair operator that reinserts these customers while respecting operational constraints.

The algorithm incorporates a three-tier hierarchy for customer insertion: (1) insertion into available vehicles, (2) creation of new trips for active vehicles if feasible, and (3) use of emergency vehicles as a last resort. This hierarchy ensures maximum utilization of the regular fleet before resorting to expensive emergency options. The acceptance mechanism employs simulated annealing with geometric cooling, balancing exploration and exploitation throughout the search process.

The objective function minimizes the total operational cost

$$f(sol) = \sum_{z \in \mathcal{Z}} \left[\sum_{r \in R_{zk}^1} c_r + \sum_{v \in V_z^{\text{active}}} c_{R_v} + c_z^{\text{fleet}} \cdot |V_z^{\text{active}}| \right] + \sum_{o \in \mathcal{O}_k^{\text{violated}}} c_o^{\text{emergency}} \quad (32)$$

where c_r denotes the total routing cost for route r and c_{R_v} represents the total routing cost for vehicle v across all its trips.

For multi-trip second-echelon routes, the route cost c_{R_v} for vehicle v includes all trip-specific components

$$c_{R_v} = \sum_{i=1}^{m_v} \left[\sum_{j=1}^{n_i} \text{dist}_{h_{j-1}^i, h_j^i} \cdot c_z^{\text{var}} + \text{dist}_{h_{n_i}^i, \gamma_v^i} \cdot c_z^{\text{var}} \right] \quad (33)$$

where $\text{dist}_{i,j}$ denotes the distance between locations i and j .

Large destroy operations result in unconditional acceptance to ensure diversification, while the standard simulated annealing mechanism balances exploitation and exploration throughout the search process.

The complete set of destroy operators (10 variants) and repair operators (8 variants), along with their detailed implementation, data structures for multi-trip representation, and parameter settings, are provided in Supplementary Materials S2. The algorithm's integration with the dynamic system ensures that each static optimization contributes to the overall dynamic objectives while maintaining computational tractability for real-time operations.

4.4. Hierarchical Bayesian Optimization for Fleet Sizing and Operational Parameters

The joint optimization of fleet sizes and operational parameters presents a challenging mixed-integer stochastic optimization problem. We develop a hierarchical Bayesian optimization framework that decomposes the problem into strategic fleet sizing at the outer level and operational parameter tuning at the inner level.

4.4.1. Problem Decomposition

Let $\mathbf{x} = (\mathbf{V}^2, \chi)$ denote the complete decision vector, where $\mathbf{V}^2 = (|V_1^2|, |V_2^2|, \dots, |V_{|Z|}^2|)$ represents the fleet sizes and $\chi = (\alpha, \beta, \varepsilon, \eta^{reloc}, \xi)$ denotes the operational parameters. The objective function can be expressed as

$$f(\mathbf{x}) = f(\mathbf{V}^2, \chi) = \sum_{z \in Z} c_z^{\text{fixed}} \cdot |V_z^2| + \mathbb{E}[C^{\text{operational}}(\chi | \mathbf{V}^2)] \quad (34)$$

This structure suggests a hierarchical optimization

$$\mathbf{V}^{2*} = \arg \min_{\mathbf{V}^2 \in \mathcal{V}} \left\{ \sum_{z \in Z} c_z^{\text{fixed}} \cdot |V_z^2| + \min_{\chi \in \mathcal{X}} \mathbb{E}[C^{\text{operational}}(\chi | \mathbf{V}^2)] \right\} \quad (35)$$

4.4.2. Outer Level: Fleet Size Optimization

At the outer level, we model the minimum achievable total cost for each fleet configuration as a Gaussian Process. The next fleet configuration is selected by maximizing the Probability of Improvement (PI) acquisition function.

4.4.3. Inner Level: Operational Parameter Optimization

For each fleet configuration $\mathbf{V}^2$, we optimize operational parameters using standard Bayesian optimization with normalized parameter space and PI acquisition function.

4.4.4. Knowledge Transfer Mechanism

A key advantage of the hierarchical framework is the ability to transfer knowledge between similar fleet configurations. When initiating the inner optimization for a new fleet configuration $\mathbf{V}_{\text{new}}^2$, we leverage the optimization history from previously evaluated configurations.

For each previously evaluated configuration $\mathbf{V}_j^2$ with optimal parameters χ_j^* , we compute a similarity weight

$$w_j = \exp \left(-\varsigma \sum_{z \in Z} \left(\frac{|V_{z,\text{new}}^2| - |V_{z,j}^2|}{\Delta_z} \right)^2 \right) \quad (36)$$

where ς controls the influence decay rate and Δ_z normalizes the fleet size differences. The initial parameter estimate for the new configuration is:

$$\chi_{\text{init}} = \frac{\sum_j w_j \chi_j^*}{\sum_j w_j} \quad (37)$$

This warm-start strategy significantly reduces the number of inner iterations required for convergence.

4.4.5. Implementation Details

The hierarchical optimization alternates between outer (10–15 iterations) and inner levels (15–25 evaluations each), with each configuration evaluated using $M = 10$ simulation runs. Early iterations use fewer inner evaluations for rapid exploration, while later iterations refine promising configurations, ensuring efficient convergence within practical computational limits.

5. Computational Experiments

This section presents computational experiments to evaluate the performance of our fleet-constrained metropolitan logistics problem (FCMLP) framework. We assess the effectiveness of multi-trip operations under fleet constraints and analyze the trade-offs

between fleet size, operational costs, and service quality through both controlled simulation experiments and real-world case studies. Particular emphasis is placed on evaluating the impact of our Bayesian Network-based demand prediction and proactive vehicle relocation strategies.

5.1. Experimental Setup

We implement our solution framework in Python and conduct all computational experiments on a system with an AMD Ryzen 5 5600X 6-Core Processor (3.70 GHz) and 16 GB of RAM. For solving the mixed-integer programming formulations in our Cost Function Approximation module, we employ IBM ILOG CPLEX 22.1.1.0 as the optimization solver. All code is implemented in Python 3.9 with NumPy for numerical computations and multiprocessing libraries for parallel simulation runs. The datasets are available in GitHub repository: <https://github.com/xyyao1116/FCMLP> (accessed on 24 November 2025).

5.1.1. System Configuration and Test Instances

We construct a metropolitan logistics network comprising interconnected cities, each with one central hub at coordinates (50, 50) and two satellites positioned 20–30 units away in opposite directions on a 100×100 UTM grid. This two-echelon architecture balances consolidation efficiency with last-mile delivery flexibility.

The system operates over a 10 h planning horizon (600 min) divided into 10 hourly decision epochs. Orders arrive uniformly with weights between one and four units and differentiated service commitments: same-day orders require pickup within 3 h and delivery 2–5 h after arrival, while next-day orders allow pickup until minute 420 and concentrated morning delivery (540–600 min). Daily order volume ranges from 150 to 200 per city.

Vehicle fleets operate with distinct capacities: first-echelon vehicles (75 units) shuttle between hubs and satellites with unlimited availability, while second-echelon vehicles (10 units) handle urban delivery with fleet size as a decision variable. Operations include 15 min satellite reload times and 5 min customer service times, with trips capped at 180 min.

The linehaul system incorporates temporal variations: peak hours (0–120, 480–600 min) reduce capacity to 80% and increase costs/travel times to 120%; off-peak hours (240–360 min) expand capacity to 120% with 80% costs/times; regular hours maintain baseline parameters. This captures metropolitan freight dynamics without demand-based adjustments.

Cost parameters reflect operational realities: regular vehicles incur 150 units daily fixed cost plus 2 units/km variable cost; emergency deployments cost 500 units fixed plus 5 units/km variable. These differentials drive the fleet size versus emergency deployment trade-off. Table 2 summarizes all configuration parameters.

Table 2. System configuration parameters.

Parameter Category	Configuration
<i>Geographic Layout</i>	
City structure	1 hub + 2 satellites per city
Coordinate system	$[0, 100] \times [0, 100]$ UTM grid
Hub location	City center (50, 50)
Satellite distance from hub	20–30 units
<i>Temporal Parameters</i>	
Planning horizon (T)	600 min (10 h)
Decision epoch interval (δ)	60 min
Number of epochs (K)	10
Order arrival window	$[0, 600]$ min (uniform)

Table 2. Cont.

Parameter Category	Configuration
<i>Order Characteristics</i>	
Daily order volume per city	150–200 orders
Order weight	1–4 units (uniform random integer)
Same-day pickup window	$[t_{\text{arrival}}, t_{\text{arrival}} + 180]$ min
Same-day delivery window	$[t_{\text{arrival}} + 120, t_{\text{arrival}} + 300]$ min
Next-day pickup window	$[t_{\text{arrival}}, 420]$ min
Next-day delivery window	$[540, 600]$ min
<i>Vehicle Specifications</i>	
First-echelon capacity (Q^1)	75 units
First-echelon fleet size	Unlimited
Second-echelon capacity (Q^2)	10 units
Second-echelon fleet size	To be optimized
Reload time (ρ)	15 min
Service time (φ)	5 min
Maximum trip duration	180 min
<i>Cost Parameters</i>	
Fixed cost (c^{fixed})	150 per vehicle per day
Variable cost (c^{var})	2 per km
Emergency fixed cost ($c^{\text{em, fixed}}$)	500 per deployment
Emergency variable cost ($c^{\text{em, var}}$)	5 per km
<i>Linehaul Configuration</i>	
<i>Peak hours (0–120, 480–600 min)</i>	
Capacity	72 units (80% of baseline)
Travel time	108–130 min (120% of baseline)
Fixed/variable cost	90–108/12–14.4 (120% of baseline)
<i>Off-peak hours (240–360 min)</i>	
Capacity	108 units (120% of baseline)
Travel time	72–86 min (80% of baseline)
Fixed/variable cost	60–72/8–9.6 (80% of baseline)
<i>Regular hours (120–240, 360–480 min)</i>	
Capacity	90 units (baseline)
Travel time	90–108 min (baseline)
Fixed/variable cost	75–90/10–12 (baseline)

5.1.2. Experimental Design

We conduct experiments in two phases. Phase 1 employs controlled simulations using a 2×3 factorial design with hub configurations (two or three hubs) and demand scenarios: More Next-Day (MND: 25% same-day, 75% next-day), Balanced (BAL: 50–50%), and More Same-Day (MSD: 75% same-day, 25% next-day). Phase 2 validates our framework using 14 days of operational data from a major e-commerce platform in Jakarta, comparing against their current Latest Departure Schedule (LDS) heuristic.

5.1.3. Solution Methods

We compare three variants to isolate the impact of our contributions:

OF-ST (Optimized Fleet, Single-Trip): Adapts Zamal et al. [27]’s framework by adding fleet constraints while restricting vehicles to single trips. This baseline jointly optimizes fleet size and operational parameters $(\alpha, \beta, \epsilon)$ without multi-trip capabilities.

OF-MT (Multi-Trip): Enables multi-trip operations with the same optimization framework, allowing vehicles to perform multiple delivery cycles.

OF-MT-R (Multi-Trip with Relocation): Adds Bayesian Network demand prediction and proactive vehicle relocation to OF-MT, using the same fleet size but repositioning idle vehicles based on predicted demand.

All methods use identical Bayesian optimization with 10 iterations evaluating 30-day simulations each. The hierarchical framework optimizes fleet sizes $|V_z^2| \in [10, 30]$ and operational parameters: slack weight $\alpha \in [0, 1]$, dispatching threshold $\beta \in [30, 180]$ min, slack exponent $\epsilon \in [1, 3]$, and relocation threshold $\eta^{\text{reloc}} \in [1.5, 4.0]$ and balance weight $\zeta \in [0.5, 2.0]$ for OF-MT-R only. All methods use a Probability-of-Improvement acquisition function, initialized with a small random design.

5.2. Simulation Results

5.2.1. Overall Performance Comparison

Table 3 summarizes the optimized fleet sizes and total costs across all test instances, highlighting the systematic advantages of multi-trip operations and the incremental benefits of proactive relocation.

Table 3. Overall performance comparison across three approaches.

Hubs	Scenario	OF-ST			OF-MT			OF-MT-R				Imb
		Fleet	Cost	ER (%)	Fleet	Cost	ER (%)	Fleet	Cost	ER (%)	Reloc/Day	
2	MND	45	185,324	3.2	36	156,875	4.1	36	149,231	3.2	3.2	0.28
	BAL	52	198,657	5.8	41	158,926	6.5	41	148,789	5.1	5.8	0.35
	MSD	58	216,482	7.5	44	167,774	8.2	44	154,072	6.8	8.4	0.42
3	MND	68	278,543	3.5	54	231,192	4.3	54	219,632	3.4	4.5	0.30
	BAL	78	298,316	6.2	60	232,687	7.0	60	216,398	5.5	7.2	0.38
	MSD	85	319,624	8.0	64	243,374	8.8	64	223,083	7.1	10.3	0.45
Average		64.3	249,491	5.7	49.8	198,471	6.5	49.8	185,201	5.2	6.6	0.36
vs. OF-ST		-	-	-	-22.6%	-20.5%	+14.0%	-22.6%	-25.8%	-8.8%	-	-
vs. OF-MT		-	-	-	-	-	-	0%	-6.7%	-20.0%	-	-

Note: Fleet represents total second-echelon vehicles across all cities. ER denotes Emergency Rate (percentage of orders requiring emergency vehicles). Cost values are daily averages over a 30-day simulation. Imb represents the average standard deviation of $B_{\gamma k}$ across satellites. Reloc/Day is the average number of vehicle relocations per day.

The results reveal substantial performance improvements through multi-trip operations, with further gains from proactive relocation. Across all scenarios, OF-MT achieves an average fleet reduction of 22.6% (from 64.3 to 49.8 vehicles) while reducing total costs by 20.5%. The addition of proactive relocation in OF-MT-R, while maintaining the same fleet size as OF-MT, delivers an additional 6.7% cost reduction and reduces emergency rates from 6.5% to 5.2% on average.

The cost reduction pattern varies with demand characteristics. MSD scenarios show the highest absolute cost improvements, while the relative benefit of relocation is most pronounced in BAL scenarios, where the mix of same-day and next-day orders creates more complex demand patterns that benefit from anticipatory positioning. The average daily relocation frequency increases with same-day demand proportion, ranging from 3.2 to 4.5 relocations in MND scenarios and 8.4 to 10.3 in MSD scenarios, reflecting the greater spatial-temporal variability in urgent delivery requirements.

5.2.2. Impact of Proactive Relocation

Figure 4 illustrates the effectiveness of proactive relocation in balancing fleet distribution and reducing emergency vehicle deployments across different scenarios. Left panel shows the average standard deviation of demand-supply imbalance scores across satellites. Right panel displays the percentage of orders requiring emergency vehicles. Percentages above bars indicate the relative improvement achieved through relocation.




Figure 4. Impact of proactive relocation on demand-supply imbalance and emergency rate reduction.

The proactive relocation strategy demonstrates consistent effectiveness in reducing demand-supply imbalances across all scenarios. Without relocation (OF-MT), the imbalance scores range from 0.42 to 0.75, indicating significant disparities in vehicle availability relative to demand across satellites. The implementation of BN-based prediction and proactive relocation (OF-MT-R) reduces these imbalances by 33–41%, with the most substantial improvements observed in MSD scenarios where demand patterns are most volatile.

The reduction in imbalance directly translates to improved service reliability. Emergency rates decrease by 17–22% across all scenarios when relocation is enabled, with the largest absolute reductions occurring in high-demand scenarios. This improvement stems from the system's ability to anticipate demand surges and preposition vehicles accordingly, reducing instances where local fleet capacity is overwhelmed. The effectiveness is particularly notable given that no additional vehicles are added—the same fleet size achieves superior performance through better spatial distribution.

5.2.3. Fleet Reduction Analysis

Figure 5 illustrates the fleet size reduction achieved through multi-trip operations across different demand scenarios.




Figure 5. Fleet size reduction achieved by multi-trip operations across different demand scenarios.

The fleet reduction pattern exhibits clear demand dependency. MSD scenarios achieve the highest fleet reductions (24.1% for 2-hub, 24.7% for 3-hub), as the urgency of same-day deliveries aligns well with the increased vehicle circulation enabled by multi-trip operations. Conversely, MND scenarios show smaller reductions (20.0% for 2-hub, 20.6% for 3-hub), as next-day orders' concentrated morning delivery windows limit opportunities for vehicle reuse. The consistent reduction across different hub configurations suggests that multi-trip benefits are robust to network scale variations.

5.2.4. Cost Structure Analysis

Table 4 decomposes the total cost savings into component categories, revealing how multi-trip operations and proactive relocation impact different cost drivers.

The cost decomposition reveals several key insights. Routing costs dominate the total (67.1–81.0%), making their reduction through efficient operations particularly impactful. The introduction of proactive relocation (OF-MT-R) reduces routing costs by an additional 2.3–3.9% compared to OF-MT, demonstrating that better vehicle positioning leads to shorter travel distances. Emergency costs show the most dramatic improvement with relocation, decreasing by 27–32% relative to OF-MT. While relocation incurs additional costs (averaging 670–2147 units across scenarios), these are vastly outweighed by the savings in emergency deployments and routing efficiency. The combined linehaul and relocation costs in OF-MT-R remain below 3.5% of total costs across all scenarios, confirming the economic viability of the proactive strategy.

Table 4. Cost component analysis across all scenarios and approaches.

Scenario	Component	OF-ST				OF-MT				OF-MT-R			
		Fleet	Routing	Emerg	LH	Fleet	Routing	Emerg	LH	Fleet	Routing	Emerg	LH+Reloc
2H-MND	Amount	6750	141,854	6720	3500	5400	118,564	8610	3426	5400	115,879	5856	4096
	% of Total	3.6	76.6	3.6	1.9	3.4	75.6	5.5	2.2	3.6	77.6	3.9	2.7
2H-BAL	Amount	7800	142,567	12,180	3510	6150	113,726	13,650	3474	6150	110,542	9828	4269
	% of Total	3.9	71.8	6.1	1.8	3.9	71.6	8.6	2.2	4.1	74.3	6.6	2.9
2H-MSD	Amount	8700	146,032	15,750	3518	6600	112,554	17,220	3400	6600	108,126	12,546	4800
	% of Total	4.0	67.5	7.3	1.6	3.9	67.1	10.3	2.0	4.3	70.2	8.1	3.1
3H-MND	Amount	10,200	217,318	11,025	5457	8100	182,567	13,545	5380	8100	177,865	9207	6460
	% of Total	3.7	78.0	4.0	2.0	3.5	79.0	5.9	2.3	3.7	81.0	4.2	2.9
3H-BAL	Amount	11,700	217,366	19,530	5420	9000	178,332	22,050	5305	9000	172,146	15,876	6376
	% of Total	3.9	72.9	6.5	1.8	3.9	76.7	9.5	2.3	4.2	79.5	7.3	2.9
3H-MSD	Amount	12,750	219,774	25,200	5400	9600	171,274	27,720	5346	9600	164,532	19,404	7547
	% of Total	4.0	68.8	7.9	1.7	3.9	70.4	11.4	2.2	4.3	73.7	8.7	3.4

Note: Fleet represents fixed vehicle costs; Routing includes first- and second-echelon variable costs; Emerg denotes emergency vehicle costs; LH is linehaul transportation cost; LH+Reloc combines linehaul and proactive relocation costs. Percentages show each component's share of total cost.

5.2.5. Relocation Pattern Analysis

Table 5 provides detailed insights into the temporal distribution and effectiveness of proactive relocations throughout the operational day.

The temporal analysis reveals sophisticated relocation patterns aligned with demand dynamics. Afternoon periods (200–400 min) consistently show the highest relocation frequency across all scenarios, coinciding with the preparation for evening delivery peaks. The Bayesian Network demonstrates robust predictive performance with accuracy ranging from 72.5% to 87.1%, with higher accuracy in afternoon periods when demand patterns are more established. Notably, BN accuracy improves with demand complexity—MSD scenarios achieve 4–8% higher accuracy than MND scenarios, suggesting that the model effectively learns from the richer spatial-temporal patterns in same-day heavy operations.

Cost savings from relocations show a clear correlation with both relocation frequency and BN accuracy. Each relocation generates average savings of 1500–2000 cost units, with

higher values in complex demand scenarios. The return on investment is substantial: relocation costs average 100–150 units per move, yielding benefit–cost ratios exceeding 10:1 in all scenarios.

Table 5. Detailed relocation impact analysis.

Hubs	Scenario	Metric	Time Period			Daily Avg
			Morning (0–200)	Afternoon (200–400)	Evening (400–600)	
2	MND	Relocations	0.8	1.2	1.2	3.2
		BN Accuracy (%)	72.5	78.3	75.6	75.5
		Cost Savings	1823	2456	2365	6644
	BAL	Relocations	1.5	2.3	2.0	5.8
		BN Accuracy (%)	76.2	82.1	79.4	79.2
		Cost Savings	2987	4123	3027	10,137
	MSD	Relocations	2.5	3.4	2.5	8.4
		BN Accuracy (%)	81.3	85.7	82.9	83.3
		Cost Savings	4234	5867	3601	13,702
3	MND	Relocations	1.0	1.8	1.7	4.5
		BN Accuracy (%)	74.2	79.5	77.1	76.9
		Cost Savings	2456	3897	3207	11,560
	BAL	Relocations	1.8	3.0	2.4	7.2
		BN Accuracy (%)	77.8	83.4	80.9	80.7
		Cost Savings	3987	5431	3871	16,289
	MSD	Relocations	3.2	4.3	2.8	10.3
		BN Accuracy (%)	82.7	87.1	84.2	84.7
		Cost Savings	5675	7234	4382	20,291

5.2.6. Temporal Relocation Patterns

Figure 6 visualizes the hourly distribution of vehicle relocations across different scenarios and time periods, revealing how the system adapts its proactive positioning strategy to anticipated demand patterns.




Figure 6. Vehicle relocation frequency throughout the day. Darker colors indicate higher relocation activity. The heatmap shows average relocations per hour across 30 simulation days, demonstrating how relocation intensity varies with scenario complexity and time of day.

The heatmap reveals distinct temporal patterns in relocation activity. The most intense relocation periods occur during hours 240–360 (off-peak afternoon), when the system prepares for evening delivery rushes by repositioning vehicles from low-demand to high-demand satellites. MSD scenarios exhibit the highest relocation intensity (up to 1.6 relocations per hour), reflecting the continuous spatial rebalancing required for same-day deliveries. Conversely, MND scenarios show more concentrated relocation patterns, primarily during mid-day periods when preparing for next-morning deliveries. This temporal intelligence demonstrates the system’s ability to anticipate future demand based on historical patterns and current order flows.

5.2.7. Multi-Trip Utilization Patterns

Figure 7 shows the distribution of vehicles by number of trips performed, demonstrating how the system leverages multi-trip capabilities under different operational conditions.




Figure 7. Distribution of vehicles by number of trips performed daily.

The trip distribution reveals how proactive relocation enhances multi-trip utilization. While OF-MT shows standard utilization patterns, OF-MT-R demonstrates a shift toward higher trip counts, particularly in the 3+ trip category. This improvement is most pronounced in BAL and MSD scenarios, where proactive positioning enables vehicles to capture additional delivery opportunities that would otherwise be missed due to poor initial positioning. The enhanced utilization contributes to the cost reductions observed in OF-MT-R without requiring additional vehicle resources.

5.2.8. Optimized Parameter Analysis

Table 6 presents the optimized operational parameters for each configuration, revealing how the system adapts to different demand patterns and operational strategies.

The parameter optimization reveals systematic adaptations across different operational strategies. The relocation parameters in OF-MT-R show interesting patterns: the threshold η^{reloc} decreases with demand complexity (3.2–3.5 for MND down to 2.0–2.2 for MSD), indicating more aggressive relocation in volatile environments. Conversely, the balance weight ξ increases with same-day proportion (0.7–0.8 for MND up to 1.1–1.2 for MSD), suggesting greater willingness to invest in repositioning when demand urgency is high. These parameter adaptations demonstrate the system’s ability to self-tune based on operational characteristics.


Table 6. Optimized operational parameters across three approaches.

Method	Hubs	Scenario	Core Parameters			Relocation	
			α	β	ε	η^{reloc}	ζ
OF-ST	2	MND	0.259	158	2.2	—	—
		BAL	0.406	120	2.0	—	—
		MSD	0.551	99	1.8	—	—
	3	MND	0.201	162	2.3	—	—
		BAL	0.350	131	2.1	—	—
		MSD	0.507	106	1.9	—	—
OF-MT	2	MND	0.454	135	1.8	—	—
		BAL	0.607	118	1.9	—	—
		MSD	0.754	89	2.0	—	—
	3	MND	0.407	148	1.7	—	—
		BAL	0.555	123	1.8	—	—
		MSD	0.702	97	1.9	—	—
OF-MT-R	2	MND	0.454	135	1.8	3.2	0.8
		BAL	0.607	118	1.9	2.5	1.0
		MSD	0.754	89	2.0	2.0	1.2
	3	MND	0.407	148	1.7	3.5	0.7
		BAL	0.555	123	1.8	2.8	0.9
		MSD	0.702	97	1.9	2.2	1.1

5.3. Sensitivity Analysis

Impact of Fleet Size on Operational Strategy

To understand the relationship between resource constraints and operational strategies, we analyze how the optimal slack weight parameter α varies with fleet size. Figure 8 illustrates this relationship across different operational modes.

**Figure 8.** Relationship between fleet size and optimal slack weight parameter α .

The sensitivity analysis reveals an inverse relationship between fleet size and the importance of slack time in assignment decisions. As fleet size decreases from 30 to 10 vehicles per city, the optimal α value increases from approximately 0.4 to 0.8, indicating that constrained fleets require greater temporal flexibility to maintain service quality. This

pattern is more pronounced for multi-trip operations (OF-MT and OF-MT-R), where the slope is steeper, suggesting that temporal coordination becomes critical when combining limited resources with vehicle reuse strategies. Notably, OF-MT-R shows the steepest slope, indicating that proactive relocation further amplifies the importance of temporal flexibility in resource-constrained environments.

5.4. Case Study

5.4.1. Dataset Description

We validate our approach using operational data from a major e-commerce platform's logistics network in the Jakarta Metropolitan Area. To enable direct comparison with our simulation experiments, we organize the real-world data into the same 2×3 factorial design as follows:

- Time period: 14 operational days.
- Network configurations: 2-hub and 3-hub systems.
- The demand scenarios are as follows:
 - More Next-Day (MND): 25% same-day, 75% next-day orders.
 - Balanced (BAL): 50% same-day, 50% next-day orders.
 - More Same-Day (MSD): 75% same-day, 25% next-day orders.
- Order volume: 2000–3000 orders per configuration.
- Current operation: Latest Departure Schedule (LDS) heuristic with predetermined fleet allocation.

Each hub configuration operates with co-located hub-satellite facilities, reflecting the common practice of integrating transshipment points with distribution centers in dense urban environments. The 2-hub system covers the central metropolitan area, while the 3-hub system extends coverage to peripheral regions. Order characteristics and operational parameters match those used in our simulation experiments, ensuring consistent evaluation conditions.

5.4.2. Performance Comparison

Table 7 compares our approaches against the company's current system, demonstrating the practical impact of our optimization framework.

Table 7. Real-world dataset performance comparison.

Hubs	Scenario	LDS (Current)			OF-ST			OF-MT			OF-MT-R (Ours)		
		Fleet	Cost	ER (%)	Fleet	Cost	ER (%)	Fleet	Cost	ER (%)	Fleet	Cost	ER (%)
2	MND	180	856,200	12.5	135	599,340	8.2	108	489,468	9.5	108	456,327	7.2
	BAL	195	918,540	15.8	150	642,978	10.1	120	514,382	11.8	120	472,230	8.9
	MSD	215	987,450	18.5	165	691,215	11.8	130	529,300	13.2	130	476,370	10.3
3	MND	260	1,286,400	13.2	195	900,480	8.7	155	738,393	10.1	155	687,246	7.8
	BAL	285	1,377,900	16.5	220	964,530	10.8	175	750,823	12.3	175	683,249	9.5
	MSD	310	1,475,640	19.2	240	1,032,948	12.5	190	784,674	13.8	190	702,294	10.8
<i>Average Improvement over LDS:</i>													
OF-ST: −23.5% fleet, −30.0% cost, −35.2% emergency rate													
OF-MT: −39.1% fleet, −43.8% cost, −24.5% emergency rate													
OF-MT-R: −39.1% fleet, −49.5% cost, −42.8% emergency rate													

The real-world validation demonstrates even stronger performance improvements than simulation results. Our OF-MT-R approach maintains the same fleet size as OF-MT (39.1% reduction from current LDS) while achieving an additional 10.3% cost reduction through proactive relocation. Most significantly, emergency rates drop from 13.4%

(LDS average) to 9.3% with OF-MT-R, a 42.8% reduction that substantially improves service reliability.

The superior real-world performance stems from several factors. The current LDS system operates with conservative rules that fail to exploit multi-trip opportunities and lacks dynamic adaptation to demand patterns. Real-world demand exhibits stronger spatial-temporal clustering than our uniformly distributed simulation demand, creating more opportunities for effective vehicle relocation. Additionally, the BN model effectively learns location-specific demand patterns from historical data, achieving prediction accuracies of 78–86% in the real-world setting compared to 72–87% in simulations.

6. Conclusions

This section synthesizes the key insights from our computational experiments and examines their implications for fleet-constrained metropolitan logistics. We analyze the performance drivers of our approach, assess its computational characteristics, and acknowledge the limitations that guide future research directions.

6.1. Key Findings

Our comprehensive experiments reveal four critical insights into fleet-constrained metropolitan logistics:

First, the integration of Bayesian Network-based demand prediction with proactive vehicle relocation delivers substantial operational improvements without requiring additional fleet resources. The OF-MT-R approach achieves 6.7% cost reduction and 20% emergency rate reduction compared to OF-MT in simulations, with even greater benefits (10.3% cost reduction, 21.2% emergency rate reduction) in real-world applications. These improvements demonstrate that predictive analytics can effectively substitute for physical assets in logistics operations.

Second, multi-trip operations deliver substantial efficiency gains when properly implemented, achieving 20.5% cost reduction and 22.6% fleet reduction in controlled simulations, with even greater improvements (49.5% cost reduction when combined with relocation) in real-world applications. These benefits stem from the fundamental synergy between two-echelon architecture and multi-trip strategies—satellites positioned near customer clusters dramatically reduce reload time compared to central depot returns, making multiple trips economically viable even under tight time constraints.

Third, the effectiveness of both multi-trip operations and proactive relocation exhibits strong demand dependency. Same-day heavy scenarios achieve the highest improvements, as urgent, continuously arriving orders create natural opportunities for vehicle circulation and benefit most from anticipatory positioning. The system demonstrates systematic adaptation through parameter tuning, with relocation thresholds decreasing and balance weights increasing as demand urgency rises. Our demand prediction analysis based on the interpretable Bayesian Network further reveals how temporal and spatial features jointly drive these improvements, providing a structured explanation for the observed performance differences across scenarios.

Fourth, successful implementation requires fundamental shifts in operational parameters and management philosophy. Our analysis shows that systems must increase emphasis on temporal flexibility, adopt earlier dispatching triggers, and embrace dynamic fleet distribution based on predictive insights. The interpretability of the Bayesian Network enables operators to understand and trust relocation decisions, facilitating the transition from reactive to proactive logistics management.

6.2. Computational Performance

Our framework demonstrates strong computational performance within metropolitan logistics operational constraints. The hourly decision epochs provide sufficient time for problem instances with 150–200 orders, with our implementation requiring 5–10 min. The hierarchical Bayesian optimization converges within 10–15 outer iterations, each inner optimization completing in under 25 evaluations. The modular architecture enables parallel processing across cities, meeting real-time requirements for typical operations, while larger networks may benefit from decomposition strategies.

6.3. Limitations and Future Directions

While our results demonstrate significant benefits, several limitations warrant acknowledgment. Our deterministic travel time assumption simplifies routing but may underestimate congestion impacts. The prohibition of split deliveries maintains operational simplicity but sacrifices potential efficiency gains. The fixed satellite assignment policy reflects practical constraints but reduces flexibility. Additionally, our Bayesian Network operates at satellite-level aggregation, which limits demand prediction granularity and may omit finer spatial patterns; finer-grained predictions could enable more precise positioning.

Future research could address these limitations through stochastic travel time modeling, relaxing split delivery constraints for specific order types, developing distributed optimization for larger networks, and enhancing demand prediction granularity through hierarchical or higher-resolution models to support more fine-grained relocation decisions. The modular architecture facilitates such extensions without restructuring the entire framework, providing a practical foundation for advancing fleet-constrained metropolitan logistics.

Supplementary Materials: The following supporting information can be downloaded at <https://www.mdpi.com/article/10.3390/app152312609/s1>, Supplementary Materials S1: Complete list of notation used in this paper; Supplementary Materials S2: Technical details including Bayesian Network implementation details and ALNS algorithm details.

Author Contributions: Conceptualization, M.L.; methodology, L.S.; software, X.Y.; validation, L.S.; writing—original draft, X.Y.; supervision, M.L. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the National Natural Science Foundation of China (NSFC) under Grant 72071144.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The datasets used in this study are available at: <https://github.com/xyyao1116/FCMLP> (accessed on 24 November 2025).

Acknowledgments: The authors are grateful for the valuable comments from the editor and reviewers.

Conflicts of Interest: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

1. Sluijk, N.; Florio, A.M.; Kinable, J.; Dellaert, N.; Van Woensel, T. Two-echelon vehicle routing problems: A literature review. *Eur. J. Oper. Res.* **2023**, *304*, 865–886. [[CrossRef](#)]
2. Cattaruzza, D.; Absi, N.; Feillet, D. Vehicle routing problems with multiple trips. *Ann. Oper. Res.* **2018**, *271*, 127–159. [[CrossRef](#)]
3. Dumez, D.; Tilk, C.; Irnich, S.; Lehuédé, F.; Olkis, K.; Péton, O. A matheuristic for a 2-echelon vehicle routing problem with capacitated satellites and reverse flows. *Eur. J. Oper. Res.* **2023**, *305*, 64–84. [[CrossRef](#)]

4. Lehmann, J.; Winkenbach, M. A matheuristic for the Two-Echelon Multi-Trip Vehicle Routing Problem with mixed pickup and delivery demand and time windows. *Transp. Res. Part C Emerg. Technol.* **2024**, *160*, 104522. [\[CrossRef\]](#)
5. Hernandez, F.; Feillet, D.; Giroudeau, R.; Naud, O. Branch-and-price algorithms for the solution of the multi-trip vehicle routing problem with time windows. *Eur. J. Oper. Res.* **2016**, *249*, 551–559. [\[CrossRef\]](#)
6. Zhen, L.; Ma, C.; Wang, K.; Xiao, L.; Zhang, W. Multi-depot multi-trip vehicle routing problem with time windows and release dates. *Transp. Res. Part E Logist. Transp. Rev.* **2020**, *135*, 101866. [\[CrossRef\]](#)
7. Fermín Cueto, P.; Gjeroska, I.; Solà Vilalta, A.; Anjos, M.F. A solution approach for multi-trip vehicle routing problems with time windows, fleet sizing, and depot location. *Networks* **2021**, *78*, 503–522. [\[CrossRef\]](#)
8. Chen, E.; Zhou, Z.; Li, R.; Chang, Z.; Shi, J. The multi-fleet delivery problem combined with trucks, tricycles, and drones for last-mile logistics efficiency requirements under multiple budget constraints. *Transp. Res. Part E Logist. Transp. Rev.* **2024**, *187*, 103573. [\[CrossRef\]](#)
9. Karademir, C.; Beirigo, B.A.; Atasoy, B. A two-echelon multi-trip vehicle routing problem with synchronization for an integrated water- and land-based transportation system. *Eur. J. Oper. Res.* **2025**, *322*, 480–499. [\[CrossRef\]](#)
10. Pan, B.; Zhang, Z.; Lim, A. Multi-trip time-dependent vehicle routing problem with time windows. *Eur. J. Oper. Res.* **2021**, *291*, 218–231. [\[CrossRef\]](#)
11. Grangier, P.; Gendreau, M.; Lehuédé, F.; Rousseau, L.-M. An adaptive large neighborhood search for the two-echelon multiple-trip vehicle routing problem with satellite synchronization. *Eur. J. Oper. Res.* **2016**, *254*, 80–91. [\[CrossRef\]](#)
12. Enthoven, D.L.J.U.; Jargalsaikhan, B.; Roodbergen, K.J.; uit het Broek, M.A.J.; Schrotenboer, A.H. The two-echelon vehicle routing problem with covering options: City logistics with cargo bikes and parcel lockers. *Comput. Oper. Res.* **2020**, *118*, 104919. [\[CrossRef\]](#)
13. Li, H.; Wang, H.; Chen, J.; Bai, M. Two-echelon vehicle routing problem with satellite bi-synchronization. *Eur. J. Oper. Res.* **2021**, *288*, 775–793. [\[CrossRef\]](#)
14. Dos Santos, A.G.; Viana, A.; Pedroso, J.P. 2-echelon lastmile delivery with lockers and occasional couriers. *Transp. Res. Part E Logist. Transp. Rev.* **2022**, *162*, 102714. [\[CrossRef\]](#)
15. Dahimi, A.; Lurkin, V.; Mohammadi, M.; Van Woensel, T. A two-echelon vehicle routing problem with mobile satellites and multiple commodities. *Eur. J. Oper. Res.* **2025**, *326*, 124–140. [\[CrossRef\]](#)
16. Tadaros, M.; Kyriakakis, N.A. A Hybrid Clustered Ant Colony Optimization Approach for the Hierarchical Multi-Switch Multi-Echelon Vehicle Routing Problem with Service Times. *Comput. Ind. Eng.* **2024**, *190*, 110040. [\[CrossRef\]](#)
17. Gutierrez, A.; Labadie, N.; Prins, C. A Two-echelon Vehicle Routing Problem with time-dependent travel times in the city logistics context. *EURO J. Transp. Logist.* **2024**, *13*, 100133. [\[CrossRef\]](#)
18. Zamal, M.A.; Schrotenboer, A.H.; Van Woensel, T. The two-echelon vehicle routing problem with pickups, deliveries, and deadlines. *Comput. Oper. Res.* **2025**, *179*, 107016. [\[CrossRef\]](#)
19. Şahin, M.K.; Yaman, H. A Branch and Price Algorithm for the Heterogeneous Fleet Multi-Depot Multi-Trip Vehicle Routing Problem with Time Windows. *Transp. Sci.* **2022**, *56*, 1636–1657. [\[CrossRef\]](#)
20. Guo, S.; Hu, H.; Xue, H. A Two-Echelon Multi-Trip Capacitated Vehicle Routing Problem with Time Windows for Fresh E-Commerce Logistics under Front Warehouse Mode. *Systems* **2024**, *12*, 205. [\[CrossRef\]](#)
21. Crainic, T.G.; Gobbato, L.; Perboli, G.; Rei, W. Logistics capacity planning: A stochastic bin packing formulation and a progressive hedging meta-heuristic. *Eur. J. Oper. Res.* **2016**, *253*, 404–417. [\[CrossRef\]](#)
22. Ghilas, V.; Demir, E.; Van Woensel, T. A scenario-based planning for the pickup and delivery problem with time windows, scheduled lines and stochastic demands. *Transp. Res. Part B Methodol.* **2016**, *91*, 34–51. [\[CrossRef\]](#)
23. Klapp, M.A.; Erera, A.L.; Toriello, A. The Dynamic Dispatch Waves Problem for same-day delivery. *Eur. J. Oper. Res.* **2018**, *271*, 519–534. [\[CrossRef\]](#)
24. Van Heeswijk, W.J.A.; Mes, M.R.K.; Schutten, J.M.J. The Delivery Dispatching Problem with Time Windows for Urban Consolidation Centers. *Transp. Sci.* **2019**, *53*, 203–221. [\[CrossRef\]](#)
25. He, Y.; Hewitt, M.; Lehuédé, F.; Medina, J.; Péton, O. A continuous-time service network design and vehicle routing problem. *Networks* **2023**, *83*, 300–323. [\[CrossRef\]](#)
26. Medina, J.; Hewitt, M.; Lehuédé, F.; Péton, O. Integrating long-haul and local transportation planning: The Service Network Design and Routing Problem. *EURO J. Transp. Logist.* **2019**, *8*, 119–145. [\[CrossRef\]](#)
27. Zamal, M.A.; Schrotenboer, A.H.; Van Woensel, T. End-to-end logistics in metropolitan areas: A stochastic dynamic order-assignment and dispatching problem. *Transp. Res. Part B Methodol.* **2025**, *199*, 103249. [\[CrossRef\]](#)
28. Sluijk, N.; Florio, A.M.; Kinable, J.; Dellaert, N.; Van Woensel, T. A Chance-Constrained Two-Echelon Vehicle Routing Problem with Stochastic Demands. *Transp. Sci.* **2023**, *57*, 252–272. [\[CrossRef\]](#)
29. Cai, D.; Chen, K.; Lin, Z.; Li, D.; Zhou, T.; Leung, M.-F. JointSTNet: Joint Pre-Training for Spatial-Temporal Traffic Forecasting. *IEEE Trans. Consum. Electron.* **2025**, *71*, 6239–6252. [\[CrossRef\]](#)

30. Beni Prathiba, S.; Ram Krishnamoorthy, S.; Soundari Kannan, K.; Selvaraj, A.K.; Ranganayakulu, D.; Fang, K.; Gadekallu, T.R. Digital Twin-Enabled Real-Time Optimization System for Traffic and Power Grid Management in 6G-Driven Smart Cities. *IEEE Internet Things J.* **2025**, *12*, 29164–29175. [[CrossRef](#)]
31. Garvey, M.D.; Carnovale, S.; Yenyurt, S. An analytical framework for supply network risk propagation: A Bayesian network approach. *Eur. J. Oper. Res.* **2015**, *243*, 618–627. [[CrossRef](#)]
32. Hosseini, S.; Ivanov, D.; Dolgui, A. Ripple effect modelling of supplier disruption: Integrated Markov chain and dynamic Bayesian network approach. *Int. J. Prod. Res.* **2020**, *58*, 3284–3303. [[CrossRef](#)]
33. Liu, M.; Liu, Z.; Chu, F.; Chu, C. A new robust dynamic Bayesian network approach for disruption risk assessment under the supply chain ripple effect. *Int. J. Prod. Res.* **2021**, *59*, 265–285. [[CrossRef](#)]
34. Liu, M.; Tang, H.; Chu, F.; Ding, Y.; Zheng, F.; Chu, C. A signomial programming-based approach for multi-echelon supply chain disruption risk assessment with robust dynamic Bayesian network. *Comput. Oper. Res.* **2024**, *161*, 106422. [[CrossRef](#)]
35. Hosseini, S.; Ivanov, D. A multi-layer Bayesian network method for supply chain disruption modelling in the wake of the COVID-19 pandemic. *Int. J. Prod. Res.* **2022**, *60*, 5258–5276. [[CrossRef](#)]
36. Liu, M.; Liu, Z.; Chu, F.; Zheng, F.; Dolgui, A. Dynamic structural adaptation for building viable supply chains under super disruption events. *Transp. Res. Part B Methodol.* **2025**, *195*, 103190. [[CrossRef](#)]
37. Attar, S.F.; Mohammadi, M.; Pasandideh, S.H.R. A Bayesian network approach to production decisions by incorporating complex causal factors. *J. Manag. Sci. Eng.* **2025**, *in press*. [[CrossRef](#)]
38. Harikrishnakumar, R.; Nannapaneni, S. Forecasting bike sharing demand using quantum Bayesian network. *Expert Syst. Appl.* **2023**, *221*, 119749. [[CrossRef](#)]
39. Liu, M.; Liu, Z.; Chu, F.; Zheng, F.; Chu, C. Risk-averse dynamic Bayesian network for supply chain risk evaluation under ripple effects and risk dispersion. *IFAC-PapersOnLine* **2025**, *59*, 733–738. [[CrossRef](#)]
40. Liang, Y.; Dong, H.; Li, D.; Song, Z. Adaptive eco-cruising control for connected electric vehicles considering a dynamic preceding vehicle. *eTransportation* **2024**, *19*, 100299. [[CrossRef](#)]
41. Powell, W.B. Reinforcement learning and stochastic optimization: A unified framework for sequential decisions. *Quant. Financ.* **2022**, *22*, 2151–2154. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.