Proceedings of the 16th International Conference on Industrial Engineering and Operations Management
Ankara, Türkiye, June 23-25, 2026

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/AN16.20260075
Published: June 23, 2026

A Time-Expanded Carbon and Congestion-Aware Electric Vehicle Routing Problem with Station Capacity and a Heterogeneous Fleet

Ecem Yavuzdogan and Aysenur Bal
Department of Industrial Engineering
Ozyegin University
Istanbul,Türkiye
ecem.yavuzdogan@ozu.edu.tr, aysenur.bal@ozu.edu.tr

Milad Elyasi
Assistant Professor
Department of Industrial Engineering
Ozyegin University
Istanbul,Türkiye
milad.elyasi@ozyegin.edu.tr

Abstract

We study an electric vehicle routing problem where routing and charging decisions must be coordinated with traffic conditions and electricity grid dynamics. We propose a time-expanded, carbon- and congestion-aware electric vehicle routing problem with station capacity (TE-CC-EVRP) that extends classical EVRP models in several directions. The planning horizon is discretized into time periods and represented by a time-expanded network that captures time-dependent travel times due to congestion, as well as time-varying electricity prices and emission factors. Charging stations have limited resources in each period, modeled through plug availability and aggregate power limits, which may create competition among vehicles for charging. The vehicle fleet is heterogeneous with respect to battery capacity, energy consumption, maximum charging power, and load capacity constraints. The objective is to minimize total operating cost, including travel and charging costs, while satisfying a global carbon budget that limits indirect emissions from electricity use. We formulate the problem as a mixed-integer linear program on a time-expanded network and develop a parallel memetic genetic matheuristic. The algorithm evolves vehicle routes through genetic operators while delegating charging scheduling and feasibility verification to an exact mixed-integer optimization subproblem. Computational experiments on benchmark instances derived from realistic traffic and grid profiles show the impact of congestion and station capacity on routing and charging strategies, highlight cost–emission trade-offs under different carbon budgets, and demonstrate that the proposed matheuristic efficiently produces high-quality solutions for large-scale instances.

Keywords
Electric Vehicle Routing Problem, Time-expanded Network, Carbon-aware Routing, Congestion-aware Routing, Memetic Matheuristic

1. Introduction
E-commerce growth and urbanization have increased demand for sustainable logistics, accelerating interest in Electric Vehicles (EVs) for last-mile delivery (Kunnapapdeelert et al., 2022). While EVs suit urban routes that

typically fall within battery limits, adoption remains constrained by range limitations, travel-time variability, charging infrastructure, and dynamic electricity pricing and carbon intensity.

Addressing these spatial and temporal interactions requires EV routing models that integrate routing, charging behavior, traffic dynamics, and infrastructure constraints. Yet most EVRP studies rely on simplifying assumptions such as static travel times, fixed electricity prices, simplified charging models, or unconstrained charging networks, and often neglect the grid's environmental profile (Huang, 2023). These limitations are particularly problematic in last-mile logistics, which must coordinate time windows, congestion, heterogeneous fleets, and en-route charging. Empirical analyses have documented notable hour-to-hour variations in electricity prices and carbon intensity, revealing consistent differences between peak and off-peak periods across regions (Hintermann, 2016). Likewise, traffic data reveal substantial variability in link travel times and speeds across the day, especially in dense urban environments. These variations create opportunities for logistics operators to synchronize routing and charging decisions with favorable periods strategically, for example, by avoiding heavily congested arcs during rush hours or by shifting charging to low-carbon, off-peak grid periods, thereby reducing operational costs and emissions without sacrificing service quality (Lin et al., 2021). However, capturing these opportunities requires optimization approaches that can jointly model (i) routing and customer sequencing in a time-dependent network, (ii) the evolution of battery state-of-charge and load over time, (iii) the timing and magnitude of charging actions under dynamic electricity attributes, and (iv) competition for limited charging capacity at stations. Such integration further tightens the coupling between time, space, and energy-related decisions. It demands more expressive modeling frameworks than those typically used for classical vehicle routing problems (VRPs), including standard time-indexed or continuous-time formulations.

In this study, we investigate a Time-Expanded Carbon- and Congestion-Aware Electric Vehicle Routing Problem with station capacity and a heterogeneous fleet (TE-CC-EVRP). A fleet of EVs with different battery capacities, charging technologies, and energy efficiencies must serve customer demands and may recharge at charging stations across the network. The planning horizon is discretized into time periods, creating a time-expanded network in which each location appears across multiple time layers. Arcs capture time-dependent travel times due to congestion, as well as time-varying electricity prices and carbon intensities. Charging stations are modeled as capacitated resources with limited charging plugs and aggregate power availability in each period. Charging decisions therefore generate both monetary costs and emissions depending on the station, time period, vehicle charging rate, and grid conditions. The objective is to minimize total operating cost, including travel and charging, while satisfying a carbon budget, or, equivalently, to minimize a weighted combination of cost and emissions. The TE-CC-EVRP is particularly challenging because it integrates time-dependent routing, heterogeneous fleets, vehicle energy dynamics, and time-varying charging constraints. Although the time-expanded formulation captures these interactions, it produces large-scale mixed-integer models with many time-layered nodes and arcs, making large instances difficult to solve with straightforward optimization or local search methods.

The contributions of this paper are threefold. First, we introduce a carbon- and congestion-aware EVRP that models time-dependent travel times, electricity prices, carbon intensities, and charging-station capacity for a heterogeneous fleet using a time-expanded network representation. Second, we develop a Mixed-Integer Programming (MIP) formulation that jointly captures routing, load evolution, battery state-of-charge dynamics, and time-dependent charging decisions, providing a high-fidelity benchmark for evaluating heuristic solutions. Third, to address the scale and complexity of realistic instances, we propose a parallel memetic genetic matheuristic for the TE-CC-EVRP. Candidate solutions represent vehicle routes and potential charging locations and are evolved through genetic operators and local search. For each routing solution, an exact mixed-integer subproblem determines optimal charging schedules while respecting station capacity and time-dependent congestion, electricity prices, and carbon intensities. Computational experiments on synthetic urban delivery instances show that the proposed approach efficiently produces high-quality solutions for large-scale networks and highlight the operational impact of congestion, station capacity, and carbon-aware charging decisions.

2. Literature Review

Urbanization and the rapid rise of e-commerce are driving the shift to sustainable last-mile logistics. Consequently, electrifying delivery fleets is both an operational trend for logistics providers and an environmental priority, as the transportation sector accounts for a significant share of global CO_2 emissions. Schneider, Stenger, and Goeke (2014) developed the E-VRPTW to combine battery limits and station visits into classical VRPTW models, laying the mathematical foundation for electric vehicle routing with time windows and recharging stations. SoC (State-of-

Charge) constraints became an explicit operational decision in routing models thanks to their formulation and computational analysis. The assumptions made by Schneider et al. were expanded and relaxed in various ways by later literature: Keskin and Çatay (2016) allowed partial recharging, eliminating unnecessary deadhead energy/time and modeling shorter, more useful charging stops. Montoya et al. (2017) proposed nonlinear charging functions (piecewise or nonlinear charge-time relations). They demonstrate that disregarding nonlinear charging can result in plans that are either infeasible or excessively costly. Goeke & Schneider (2015) and related work improved on basic linear distance-based energy assumptions by extending energy models to account for consumption that is dependent on speed and operation. Time-varying traffic is another challenge that urban routing must address. When stop-and-go dynamics and congestion are introduced, static travel-time assumptions can yield infeasible or suboptimal plans (Ichoua et al., 2003).

Additionally, empirical research shows that traffic dynamics change EV energy use by tens of percent depending on driving mode and speed profiles (Grigorev et al., 2021). For realistic routing, modeling frameworks that explicitly capture time-dependent travel times and congestion-influenced drainage are crucial. Researchers use time-expanded or time-dependent network formulations, which compute travel time and energy consumption as functions of departure time, to show such temporal dynamics at the routing level (Ichoua et al., 2003; Lu et al., 2020). The use of time-variant electricity price modeling and carbon-aware objective functions is a key extension in recent EVRP research. Research on combined routing and charging/discharging scheduling under time-varying electricity prices shows that fleet operators can co-optimize routing and charging/discharging scheduling to minimize costs and shift charging to periods of lower carbon electricity (or to times of lower prices) (Zhang et al., 2013). When pricing and grid-intensity signals are used, these approaches show potential to reduce emissions at a limited cost penalty. Abbasi, Turki, and Dellagi (2025) also point to multi-period supply-chain models that impose carbon limitations or include carbon policy tools to align transport and production decisions with climate goals.

The capacity of charging stations cannot be neglected. Waiting times have a significant impact on cost and feasibility, according to models that treat chargers as capacitated resources, where service times differ, and several vehicles compete for a limited number of plugs. Simulation-based and capacitated-station formulations have been created to account for stochastic queues and station-level power limits (Sadeghi-Barzani et al., 2014). Disregarding these constraints can result in irrational schedules and cost increases in practice (Sweda and Klabjan, 2012).

In large-scale instances, metaheuristics remain the dominant solution approach due to the NP-hardness of the TE-CC-EVRP and related time-dependent, station-capacitated EVRP variants (Schneider et al., 2014; Felipe et al., 2014; Keskin & Çatay, 2016). Population-based evolutionary metaheuristics are particularly well-suited to complex routing problems due to their ability to explore large solution spaces efficiently. In this context, memetic algorithms, which combine genetic search with local improvement procedures, provide an effective framework for balancing exploration and intensification (Labadi et al., 2008; Cattaruzza et al., 2014). Their hybrid structure enables the integration of problem-specific operators to handle routing decisions, charging station assignments, and time-dependent constraints in EV routing problems (Peng et al., 2019).

To conclude, a multi-layer modeling approach for urban EV logistics is supported by recent literature: (1) realistic energy models (nonlinear charging, speed-dependent consumption), (2) time dependent travel times and traffic-aware energy accounting, (3) station capacity and queuing representation, and (4) multi-objective/carbon aware formulations that integrate cost, energy, and emission metrics. For large, realistic instances, hybrid metaheuristic frameworks such as memetic algorithms, which combine evolutionary search with local improvement procedures, offer a practical and effective solution strategy.

3. MILP Model: Time-Expanded EVRP with Tiered Charging Emissions

In this section, we present the Mixed-Integer Linear Programming (MILP) formulation for the proposed Time-Expanded Carbon- and Congestion-Aware Electric Vehicle Routing Problem (TE-CC-EVRP). The model captures the interactions among vehicle routing, battery dynamics, charging decisions, station capacity limitations, and time-dependent electricity attributes within a unified time-expanded network representation.

The planning horizon is discretized into equal time intervals, with each physical node replicated across time layers to form a time-expanded network in which vehicles traverse arcs representing travel, waiting, or charging activities. Travel arcs incorporate congestion-driven travel times and energy consumption, while charging arcs capture energy replenishment decisions subject to station capacity limits and vehicle-specific charging power. The charging infrastructure is modeled using a two-tier electricity structure, in which each station and time period has limited clean-energy capacity, representing a lower-emission electricity supply. At the same time, additional charging demand is met by a higher-emission tier that represents carbon-intensive generation sources. This structure captures the operational reality of power systems, where marginal electricity generation becomes increasingly carbon-intensive as demand rises. Vehicles are heterogeneous in battery capacity, charging capability, and energy consumption characteristics, and must serve each customer exactly once while maintaining feasible battery levels, cargo capacity, and station capacity constraints. Charging decisions influence both operational costs and carbon emissions, and the model minimizes total operational costs comprising charging costs, carbon-emission penalties, and travel-time penalties while respecting a carbon budget constraint. The formulation jointly determines vehicle routes, charging schedules, and energy allocations across the two electricity tiers. Table 1 summarizes the sets, parameters, and decision variables used in the formulation.

Table 1: TE-CC-EVRP Nomenclature

Sets and Indices	
$\mathcal{K}$	set of vehicles ($k \in \mathcal{K}$)
$\mathcal{C}$	set of customers ($c \in \mathcal{C}$)
$\mathcal{S}$	set of charging stations ($s \in \mathcal{S}$)
$\mathcal{N}$	set of all nodes, $\mathcal{N} = \mathcal{C} \cup \mathcal{S} \cup \{0\}$ (depot 0)
$\mathcal{H}$	set of discrete time points (e.g., every Δ minutes), ($h \in \mathcal{H}$)
$\mathcal{A}$	set of time-expanded arcs (i, h, j, m) with $i, j \in \mathcal{N}, h, m \in \mathcal{H}$

Parameters	
Δ	time step size (minutes)
$T_{\max}$	time horizon (minutes)
d_{ij}	distance from node i to node j (km)
$E_{i,h,j,m}$	energy consumption (kWh) on arc $(i, h, j, m) \in \mathcal{A}$
B_k	initial battery energy (kWh) for vehicle k
$\bar{b}$	battery upper bound (kWh)
D_c	demand at customer c
$\bar{q}$	cargo upper bound
P_k	max charging power (kW) for vehicle k
G	clean-tier capacity (kW) at each station and time
U	max simultaneous vehicles charging/waiting at a station per time step
$\Psi_{s,h}$	electricity price ($/kWh) at station s and time h
$\Omega^{\text{cln}}_{s,h}$	clean-tier carbon intensity (kg CO_2/kWh)
$\Omega^{\text{drt}}_{s,h}$	dirty-tier carbon intensity (kg CO_2/kWh)
ω	weight for carbon cost
ϵ	weight for travel time penalty
X_k	carbon budget limit for vehicle k (kg CO_2)
M_b, M_q	big-M constants
Decision Variables	
$x^k_{i,h,j,m}$	1 if vehicle k uses arc $(i, h, j, m) \in \mathcal{A}$
$g^{\text{cln}}_{k,s,h}$	clean energy charged by vehicle k at station s starting at time h (kWh)
$g^{\text{drt}}_{k,s,h}$	dirty energy charged by vehicle k at station s starting at time h (kWh)
$b_{k,i,h}$	battery energy of vehicle k at node i and time h (kWh)
$q_{k,i,h}$	cargo remaining on vehicle k at node i and time h
$L^{\text{cln}}_{s,h}$	total clean-tier load at station s and time h
$L^{\text{drt}}_{s,h}$	total dirty-tier load at station s and time h

$$\min \sum_{k\in\mathcal{K}} \sum_{s\in\mathcal{S}} \sum_{h\in\mathcal{H}} \left(\Psi_{s,h}\, g^{\text{cln}}_{k,s,h} + 1.2\,\Psi_{s,h}\, g^{\text{drt}}_{k,s,h}\right) \tag{1}$$

$$+ \, \omega \sum_{k\in\mathcal{K}} \sum_{s\in\mathcal{S}} \sum_{h\in\mathcal{H}} \left(\Omega^{\text{cln}}_{s,h}\, g^{\text{cln}}_{k,s,h} + \Omega^{\text{drt}}_{s,h}\, g^{\text{drt}}_{k,s,h}\right) \tag{2}$$

$$+ \, \epsilon \sum_{k\in\mathcal{K}} \sum_{(i,h,j,m)\in\mathcal{A}: i\neq j} (m-h)\, x^k_{i,h,j,m}. \tag{3}$$

$$\sum_{(u,t):(u,t,i,h)\in\mathcal{A}} x^k_{u,t,i,h} - \sum_{(j,m):(i,h,j,m)\in\mathcal{A}} x^k_{i,h,j,m} = -1, \qquad \forall k \in \mathcal{K},\ (i,h) = (0,0) \tag{4}$$

$$\sum_{(u,t):(u,t,0,T_{\max})\in\mathcal{A}} x^k_{u,t,0,T_{\max}} - \sum_{(j,m):(0,T_{\max},j,m)\in\mathcal{A}} x^k_{0,T_{\max},j,m} = 1, \qquad \forall k \in \mathcal{K} \tag{5}$$

$$\sum_{(u,t):(u,t,i,h)\in\mathcal{A}} x^k_{u,t,i,h} = \sum_{(j,m):(i,h,j,m)\in\mathcal{A}} x^k_{i,h,j,m}, \qquad \forall k \in \mathcal{K},\ \forall i \in \mathcal{N},\ \forall h \in \mathcal{H} \setminus \{0, T_{\max}\} \tag{6}$$

$$\sum_{k\in\mathcal{K}} \sum_{(i,h,m):(i,h,c,m)\in\mathcal{A},\, i\neq c} x^k_{i,h,c,m} = 1, \qquad \forall c \in \mathcal{C} \tag{7}$$

$$\sum_{(i,t,j,m)\in\mathcal{A}:t\le h,\ m\ge h+\Delta} x^k_{i,t,j,m} = 1, \qquad \forall k\in\mathcal{K},\ \forall h\in\mathcal{H}\setminus\{T_{\max}\} \tag{8}$$

$$b_{k,0,0} = B_k, \qquad \forall k\in\mathcal{K} \tag{9}$$

$$q_{k,0,0} = \sum_{(i,h,j,m)\in\mathcal{A}:j\in C} D_j\, x^k_{i,h,j,m}, \qquad \forall k\in\mathcal{K} \tag{10}$$

$$0 \le b_{k,i,h} \le \bar{b}, \qquad \forall k\in\mathcal{K},\ \forall i\in\mathcal{N},\ \forall h\in\mathcal{H} \tag{11}$$

$$0 \le q_{k,i,h} \le \bar{q}, \qquad \forall k\in\mathcal{K},\ \forall i\in\mathcal{N},\ \forall h\in\mathcal{H} \tag{12}$$

$$b_{k,j,m} \le b_{k,i,h} - E_{k,i,h,j,m} + \begin{cases} g^{\text{cln}}_{k,i,h} + g^{\text{drt}}_{k,i,h}, & \text{if } i=j \text{ and } i\in S,\\ 0, & \text{otherwise,}\end{cases} + M_b\left(1 - x_{k,i,h,j,m}\right), \qquad \forall k\in\mathcal{K},\ \forall(i,h,j,m)\in\mathcal{A} \tag{13}$$

$$b_{k,j,m} \ge b_{k,i,h} - E_{k,i,h,j,m} + \begin{cases} g^{\text{cln}}_{k,i,h} + g^{\text{drt}}_{k,i,h}, & \text{if } i=j \text{ and } i\in S,\\ 0, & \text{otherwise,}\end{cases} - M_b\left(1 - x_{k,i,h,j,m}\right), \qquad \forall k\in\mathcal{K},\ \forall(i,h,j,m)\in\mathcal{A} \tag{14}$$

$$q_{k,j,m} \le q_{k,i,h} - D_j + M_q\left(1 - x^k_{i,h,j,m}\right), \qquad \forall k\in\mathcal{K},\ \forall(i,h,j,m)\in\mathcal{A} \tag{15}$$

$$q_{k,j,m} \ge q_{k,i,h} - D_j - M_q\left(1 - x^k_{i,h,j,m}\right), \qquad \forall k\in\mathcal{K},\ \forall(i,h,j,m)\in\mathcal{A} \tag{16}$$

$$\sum_{k\in\mathcal{K}} x^k_{s,h,s,h+\Delta} \le U, \qquad \forall s\in\mathcal{S},\ \forall h\in\mathcal{H}: h+\Delta\in\mathcal{H} \tag{17}$$

$$g^{\text{cln}}_{k,s,h} + g^{\text{drt}}_{k,s,h} \le \left(P_k\,\frac{\Delta}{60}\right) x^k_{s,h,s,h+\Delta}, \qquad \forall k\in\mathcal{K},\ \forall s\in\mathcal{S},\ \forall h\in\mathcal{H}: h+\Delta\in\mathcal{H} \tag{18}$$

$$L^{\text{cln}}_{s,h} = \sum_{k\in\mathcal{K}} g^{\text{cln}}_{k,s,h}, \qquad \forall s\in\mathcal{S},\ \forall h\in\mathcal{H} \tag{19}$$

$$L^{\text{drt}}_{s,h} = \sum_{k\in\mathcal{K}} g^{\text{drt}}_{k,s,h}, \qquad \forall s\in\mathcal{S},\ \forall h\in\mathcal{H} \tag{20}$$

$$0 \le L^{\text{cln}}_{s,h} \le G, \qquad \forall s\in\mathcal{S},\ \forall h\in\mathcal{H} \tag{21}$$

$$\sum_{s\in\mathcal{S}}\sum_{h\in\mathcal{H}} \left(\Omega^{\text{cln}}_{s,h}\, g^{\text{cln}}_{k,s,h} + \Omega^{\text{drt}}_{s,h}\, g^{\text{drt}}_{k,s,h}\right) \le X_k, \qquad \forall k\in\mathcal{K} \tag{22}$$

$$x^k_{i,h,j,m} \in \{0,1\}, \qquad \forall k\in\mathcal{K},\ \forall(i,h,j,m)\in\mathcal{A} \tag{23}$$

$$g^{\text{cln}}_{k,s,h} \ge 0, \quad g^{\text{drt}}_{k,s,h} \ge 0, \qquad \forall k\in\mathcal{K},\ \forall s\in\mathcal{S},\ \forall h\in\mathcal{H} \tag{24}$$

$$L^{\text{cln}}_{s,h} \in [0,G], \quad L^{\text{drt}}_{s,h} \ge 0, \qquad \forall s\in\mathcal{S},\ \forall h\in\mathcal{H} \tag{25}$$

$$b_{k,i,h} \in [0,\bar{b}], \quad q_{k,i,h} \in [0,\bar{q}], \qquad \forall k\in\mathcal{K},\ \forall i\in\mathcal{N},\ \forall h\in\mathcal{H} \tag{26}$$

The objective includes three components. The first term (1) represents the electricity cost associated with charging vehicles at stations, with dirty-tier electricity penalized by a higher price multiplier. The second term (2) accounts for the carbon emissions resulting from electricity consumption, weighted by parameter ω. The third term (3) represents a travel-time penalty reflecting network congestion. Constraint (4) ensures that each vehicle departs from the depot at the beginning of the planning horizon. Constraint (5) guarantees that every vehicle returns to the depot at the end of the planning horizon. Constraint (6) enforces flow conservation at every node and time layer of the time-expanded network. Constraint (7) ensures that each customer is visited exactly once by the fleet. Constraint (8) guarantees that each vehicle occupies exactly one arc during every time interval. Constraint (9) sets the initial battery level for each vehicle at the depot. Constraint (10) initializes the vehicle load based on the customers assigned to the route. Constraints (11) and (12) enforce feasible limits for battery energy and cargo capacity throughout the route. Constraints (13) and (14) link battery levels between consecutive nodes, accounting for energy consumption during travel and energy gained from charging at stations. Constraints (15) and (16) track cargo evolution along the route and ensure customer demands are met. Constraint (17) limits the number of vehicles that can simultaneously charge or wait at a station during a given time interval. Constraint (18) ensures that charging power does not exceed the maximum charging rate of the vehicle. Constraints (19) and (20) define the total charging load supplied by the clean and dirty electricity tiers at each station and time step. Constraint (21) limits the amount of clean electricity available at each station. Constraint (22) enforces a carbon budget for each vehicle by limiting total emissions associated with electricity consumption. Finally, constraints (23)–(26) define the dvariableecision decision domains.

4. Solution Methodology

This section presents the solution techniques developed for the TE-CC-EVRP. Our primary approach is a *parallel memetic genetic matheuristic* that combines (i) population-based evolutionary search over routing structures with (ii) an exact mixed-integer optimization subproblem that optimally schedules charging and enforces time-expanded feasibility for selected candidate routing patterns. The resulting method is a *matheuristic*: the genetic algorithm explores the combinatorial space of customer orderings, vehicle partitioning, and candidate charging-station visits, while a MIP-based "oracle" computes the best feasible charging and timing decisions for fixed route structures under congestion, electricity dynamics, and station capacity constraints. Since exact optimization of the full TE-CC-EVRP quickly becomes intractable as the number of time layers and charging interactions grows, the proposed parallel memetic matheuristic constitutes the core methodological contribution of this study.

The final algorithm is designed to (a) preserve the global "each customer served exactly once" structure through a chromosome-and-repair design, (b) incorporate embedded local improvement (memetic search) to intensify promising solutions, and (c) leverage surrogate screening, parallel exact evaluation, and memorization to reduce the number and cost of computationally expensive exact subproblem calls. In what follows, we describe the chromosome representation, the exact and surrogate evaluation scheme, the genetic operators, the memetic local search, and the overall algorithmic framework.

4.1 Parallel Memetic Genetic Matheuristic

The TE-CC-EVRP couples routing, time-dependent travel, and charging under station capacity. A key design choice of our method is to separate the *route structure* decision (which customers and charging-station visits each vehicle performs and in which order) from the *continuous and time-indexed* decisions (when to depart, when to charge, how much to charge per time step, and how to respect station capacity constraints). For any candidate routing structure, we delegate detailed scheduling and charging decisions to an exact subproblem, solved via the MILP formulation of Section **3**. This creates an evaluation "oracle" that returns a rigorous cost (or infeasibility signal) for a decoded route structure. The evolutionary layer then uses these evaluations, together with a fast surrogate screening mechanism, to guide selection, crossover, mutation, and memetic intensification.

4.1.1 Solution Representation

In the final implementation, a solution is not represented directly as a dictionary of routes. Instead, each solution is encoded as a chromosome composed of two parts:

$$X = (\pi, \mathbf{z}),$$

where π is a giant-tour permutation and $\mathbf{z}$ is a binary activation vector associated with charging-station copies. The permutation is defined over the set of customers, station-copy tokens, and vehicle-separator tokens.

More precisely, if each physical charging station $s \in \mathsf{S}$ is replicated into L station copies and if $|\mathsf{K}|$ vehicles are available, then

$$\pi \in \mathrm{Perm}\big(\mathcal{C} \cup \mathcal{S}^{\text{copy}} \cup \Sigma\big),$$

where S^{copy} denotes the set of station copies, and Σ contains $|\mathsf{K}| - 1$ separator markers.

This representation is motivated by the structure of the exact evaluator. In the fixed-route MIP, a charging station may be visited more than once by the same vehicle or by different vehicles at different times. A representation containing only one copy of each physical station would therefore be too restrictive. By introducing station copies, the chromosome can express repeated visits to the same physical charging station while remaining permutation-based. The binary vector $\mathbf{z}$ determines which station copies are active in the decoded solution; inactive station copies are ignored during decoding.

A decoded solution is represented as a dictionary of routes, one route per vehicle:

$$R(X) = \{R_k : k \in \mathcal{K}\}, \qquad R_k = (n_{k,1}, n_{k,2}, \ldots, n_{k,|R_k|}),$$

where each $n_{k,r} \in \mathsf{C} \cup \mathsf{S}$ denotes either a customer or a physical charging station. The depot is implicit: each route is assumed to start and end at the main depot, i.e., the traveled sequence is $0 \rightarrow R_k \rightarrow 0$.

[1] The exact evaluator exploits a reduced variant of the full model in which routing decisions are fixed to candidate routes, while preserving the time-expanded charging and station-capacity constraints.

The decoding procedure scans the giant-tour permutation from left to right. Separator tokens split the giant tour into consecutive vehicle segments. Customer tokens are always inserted into the current vehicle route. Station-copy tokens are inserted only if their corresponding activation bit equals one, in which case the copy is mapped back to its underlying physical station. Because customers appear as unique genes in the permutation, the global customer-service structure is embedded in the chromosome itself. At the same time, the station-copy mechanism preserves flexibility for repeated charging visits.

4.1.2 Exact Evaluation and Fitness Function

Given a decoded candidate route set $R = R(X)$, we evaluate its quality using a two-stage mechanism The first stage computes a lightweight surrogate score for all uncached decoded route sets. The second stage calls the exact time-expanded scheduling and charging subproblem only for a selected subset of promising uncached candidates.

For exact evaluation, the subproblem uses the TE-CC-EVRP formulation from Section **3,** with routing decisions fixed to reflect the decoded route sequence for each vehicle. The evaluator returns the minimum objective value (or declares infeasibility) under: (i) congestion-dependent travel times and energy consumption,
(ii) time-varying electricity prices and emission factors, (iii) station plug and aggregate capacity constraints, and (iv) vehicle battery, charging power, and cargo constraints. Let Exact Cost(R) denote the objective value returned by the exact evaluator.

Before invoking the exact evaluator, we compute lightweight structural violation counts:

$$\mathrm{miss}(R) = \text{number of customers not appearing in any } R_k,$$

$$\mathrm{dup}(R) = \text{number of duplicate customer occurrences beyond one.}$$

These checks provide a fast rejection mechanism and avoid unnecessary exact solves.

Because exact evaluation dominates runtime, all uncached decoded routes are first ranked by a cheap surrogate score based on routing structure. In particular, the surrogate combines the total route distance, the number of charging-station visits, a route-balance term, and a penalty for excessively long legs. This proxy does not replace exact evaluation; rather, it prioritizes which chromosomes deserve an exact solution to the time-expanded subproblem. Let $\hat{F}(R)$ denote the resulting surrogate score.

The fitness minimized by the evolutionary algorithm is therefore defined as

$$F(R) = \begin{cases} \mathrm{ExactCost}(R), & \text{if } \mathrm{miss}(R)=0,\ \mathrm{dup}(R)=0, \text{ and } R \text{ is exactly evaluated as feasible,} \\ \mathrm{ExactCost}(R) + \alpha_{\mathrm{inf}}, & \text{if } R \text{ is exactly evaluated and the subproblem is infeasible,} \\ M + \alpha_{\mathrm{miss}}\,\mathrm{miss}(R) + \alpha_{\mathrm{dup}}\,\mathrm{dup}(R), & \text{if } \mathrm{miss}(R) > 0 \text{ or } \mathrm{dup}(R) > 0, \\ \widehat{F}(R), & \text{if } R \text{ is structurally valid but not selected for exact evaluation in the current generation.} \end{cases} \tag{27}$$

Here **M** is a large constant and $\alpha_{\mathrm{miss}}, \alpha_{\mathrm{dup}}, \alpha_{\mathrm{inf}} > 0$ are penalty coefficients. This definition ensures that structurally invalid chromosomes are strongly dominated. At the same time, infeasibilities detected by the exact evaluator are penalized but still allow the search to learn which routing patterns tend to violate charging or timing constraints. At the same time, surrogate scoring enables screening the entire population without solving an exact MIP for each candidate.

The exact evaluation budget in each generation is bounded. Among the uncached decoded route sets, only a limited fraction of the best surrogate-ranked candidates is sent to the exact evaluator, subject to lower and upper bounds on the number of exact calls per generation. Additionally, route-level exact MIP evaluations use restricted time and gap limits to prevent excessive effort in proving optimality for a single route structure memorization (fitness caching). Because exact evaluation is computationally expensive, we store previously evaluated decoded route sets in a hash-based cache. Each route dictionary is converted into a unique string key via a canonical ordering of vehicles and their node sequences. If a decoded-route hash exists in the cache, we reuse its stored fitness and diagnostics, avoiding redundant exact solves. Importantly, caching is performed at the level of decoded route sets rather than raw chromosomes, since different chromosomes may decode into the same vehicle routes.

4.1.3 Greedy Randomized Initial Population

The algorithm initializes a population of size P directly in chromosome space. For each individual, the full set of customer genes, station-copy genes, and separator genes is randomly permuted to create the giant-tour component. Then, the station-copy activation vector is generated by assigning each station copy a binary usage value. This randomized construction provides broad diversification in customer orders, vehicle segmentation, and candidate charging patterns. To avoid an initial population with a systematically inactive charging structure, the procedure may activate at least one station copy when all activation bits are initially zero.

Although this initialization differs from the earlier route-dictionary construction, its purpose remains the same: to create a diverse population while preserving broad customer coverage and a wide variety of charging patterns. Each initial chromosome undergoes a repair procedure to ensure consistency before evaluation.

4.1.4 Feasibility Repair Mechanism

To maintain solution validity, we apply a repair operator after initialization and after every genetic modification. In the final chromosome design, repair operates on the giant-tour representation rather than on route dictionaries directly. Specifically, it enforces that the permutation contains each required gene exactly once and that the station-copy activation vector has the correct length and binary values. Extraneous genes are removed, and any missing genes are reinserted.

This mechanism plays the same structural role as in the earlier route-based version: it preserves the global "visit each customer exactly once" condition while allowing charging-station visits to remain flexible. Because customer genes are unique elements of the permutation, the repair step guarantees customer completeness at the chromosome level, and decoded routes inherit this property automatically.

4.1.5 Selection

We use tournament selection to choose parents for reproduction. In each tournament, k_{tour} individuals are sampled uniformly at random, and the one with the lowest fitness value $F(\cdot)$ is selected. Tournament selection balances selection pressure and diversity and performs well when candidate quality may be a mixture of exact scores, surrogate scores, and feasibility penalties.

4.1.6 Crossover Operator

The final crossover operator acts directly on chromosome space. For the permutation component, we use an order crossover mechanism that preserves a contiguous subsequence from one parent and fills the remaining positions with the relative order induced by the second parent. This is well-suited to the giant-tour representation because all genes in the permutation are unique. For the station-copy activation vector, we use a uniform recombination rule in which each bit is independently inherited from one of the two parents.

This crossover preserves meaningful customer subsequences, vehicle-partition structure, and candidate charging patterns, while still allowing substantial recombination between parents. After the crossover, the repair mechanism of Section 4.1.4 is applied to restore chromosome consistency if needed.

4.1.7 Mutation Operators

We employ multiple mutation moves selected at random to diversify the search. These moves operate on the chromosome representation and include: swapping two positions in the giant-tour permutation, relocating a gene from one position to another, reversing a subsequence of the permutation, and flipping one or more station-copy activation bits. Through these operators, the search can simultaneously alter customer ordering, implicit vehicle partitioning, and charging-station usage patterns. Conceptually, these moves play the same role as the earlier route-level mutations: they diversify the route structure and charging pattern. The difference is that they now act as a representation compatible with repeated station visits and can be repaired more systematically. After mutation, we apply the repair mechanism to restore the integrity of the chromosome.

4.1.8 Memetic Local Search

To intensify the search around promising individuals, we apply a memetic local search to a subset of offspring with probability p_{ls}. Starting from a candidate chromosome, local search iteratively proposes small-neighborhood moves, such as permutation swaps, relocations, segment reversals, and station-copy activation flips, for a fixed number of iterations. Each neighbor is repaired and decoded before being scored.

A key refinement in the final implementation is that local search does not invoke the exact evaluator at every neighborhood step. Instead, it first uses the cached exact score if the decoded route structure has already been evaluated; otherwise, it relies on the surrogate score described in Section 4.1.2. If an improving neighbor is found under this fast-scoring rule, it becomes the new incumbent in the local search. This design preserves the memetic component's role in intensification while avoiding an excessive number of route-level exact MIP solves.

4.1.9 Parallel Population Evaluation

Fitness evaluation dominates runtime because each exact call to the subproblem may require solving a time-expanded MILP. We therefore evaluate the selected exact candidates in parallel using a process pool of size **W** (i.e., the number of workers). Each worker constructs and solves its own exact model instance for a decoded route set, ensuring solver isolation and avoiding shared-state issues. Parallel evaluation is combined with memorization: decoded routes with cached hashes are excluded from parallel jobs and assigned their stored fitness immediately.

In the final implementation, the worker pool persists across generations, eliminating the need to create and destroy parallel processes repeatedly. Moreover, each worker solves its exact subproblem using a restricted solver-thread setting, preventing excessive core oversubscription when several route-level evaluations are carried out simultaneously. This combination of bounded exact evaluation, persistent parallelism, and memorization yields substantial speedups relative to a naive "evaluate every individual exactly" strategy.

4.2 Overall Framework

Algorithm 1 summarizes the overall parallel memetic genetic matheuristic. The procedure starts by constructing and repairing an initial population of chromosomes, then repeatedly applies elitism, tournament- based reproduction, crossover, mutation, optional memetic local search, surrogate screening, and bounded parallel exact evaluation. The best exactly verified feasible solution observed during the run is maintained throughout the search. Before termination, a final exact-polishing step is applied to the best-ranked candidates so that the returned incumbent is exactly evaluated.

Algorithm 1 Parallel Memetic Genetic Matheuristic for the TE-CC-EVRP

Input: Instance data, time limit T, population size P, generations $G_{\max}$, crossover rate p_c, mutation rate p_m, local-search rate p_{ls}, workers W
Output: Best route set $R^\star$ and its objective value
Construct an initial population $\mathcal{P}$ of P chromosomes (Section 4.1.3)
Repair all individuals to enforce chromosome consistency (Section 4.1.4)
Decode and evaluate $\mathcal{P}$ using surrogate screening, caching, and bounded exact evaluation (Sections 4.1.2–4.1.9)
Set $R^\star \leftarrow$ best exactly evaluated feasible solution in $\mathcal{P}$
for $g = 1$ to $G_{\max}$ **or until** wall-clock time exceeds T **do**
    Sort $\mathcal{P}$ by fitness; copy the top $\lceil \eta P \rceil$ elites into $\mathcal{P}_{\text{new}}$
    while $|\mathcal{P}_{\text{new}}| < P$ **and** time remains **do**
        Select parents via tournament selection (Section 4.1.5)
        Create offspring by crossover with probability p_c (Section 4.1.6)
        if rand $< p_m$ **then**
            Apply a random mutation operator (Section 4.1.7)
        end if
        Repair the offspring (Section 4.1.4)
        if rand $< p_{\text{ls}}$ **then**
            Apply memetic local search using cached or surrogate scores (Section 4.1.8)
        end if
        Add the offspring to $\mathcal{P}_{\text{new}}$
    end while
    $\mathcal{P} \leftarrow \mathcal{P}_{\text{new}}$
    Decode and evaluate $\mathcal{P}$ using caching, surrogate screening, and bounded parallel exact evaluation
    Update $R^\star$ if a better exactly evaluated feasible solution is found
end for
Perform a final exact evaluation of the best-ranked candidates
return $R^\star$ and $F(R^\star)$

Implementation notes. The algorithm is implemented with both a wall-clock time limit and a generation limit. Elitism is used to preserve high-quality individuals across generations. To reduce redundant computation, all evaluated decoded route sets are stored in a global cache keyed by a canonical hash of the route dictionary. Exact evaluation is applied only to a limited promising subset of uncached candidates in each generation, and these exact solves are executed in parallel via independent worker processes. The memetic local search relies on cached exact values or surrogate scores rather than invoking exact optimization at every neighborhood step. The final output is the best, exactly verified solution encountered during the run, together with diagnostic information recorded during evaluation.

5. Numerical Results

We conducted extensive computational experiments to validate the proposed Time-Expanded Carbon- and Congestion-Aware Electric Vehicle Routing Problem (TE-CC-EVRP) formulation and to evaluate the performance of the proposed parallel memetic genetic matheuristic. The experiments are designed to assess both the quality of solutions produced by the algorithm and its ability to handle the complex interactions among routing decisions, time-dependent congestion, charging infrastructure capacity, and dynamic electricity attributes. The computational study focuses on three main objectives. First, we examine the scalability of the proposed approach across instances of different sizes and network structures. Second, we evaluate the algorithm's ability to generate feasible routing and charging strategies under time-dependent electricity prices and carbon intensities. Third, we analyze the impact of dynamic grid conditions and charging infrastructure limitations on routing behavior and charging patterns. The proposed algorithms were implemented in Python 3.11. All experiments were conducted on a 64-bit Windows Server equipped with two 2.4 GHz Intel Xeon CPUs and 24 GB of RAM. The evolutionary search component and the exact charging-scheduling evaluator were executed within the same computational environment. Parallel fitness evaluation was performed using multiple worker processes, enabling simultaneous evaluation of candidate routing solutions.

5.1 Instance Generation

To evaluate the proposed methodology, we generated a comprehensive set of synthetic instances designed to capture the key structural characteristics of the TE-CC-EVRP. All instances were produced using a dedicated instance generator implemented in Python. The generator creates problem instances by sampling spatial locations, vehicle characteristics, demand profiles, and dynamic electricity attributes, ensuring full reproducibility via controlled random seeds.

Each instance consists of a depot, customer locations, charging stations, and a fleet of electric delivery vehicles. The depot is fixed at the origin of the coordinate system. Customers are generated using a polar-coordinate scheme, where each is assigned a random angle and a distance sampled from a predefined radius interval, resulting in spatially dispersed delivery points around the depot. Charging stations are distributed across the service region using a predefined spatial template to ensure network accessibility.

Vehicles are heterogeneous in battery capacity and maximum charging power, drawn from predefined sets. Battery capacities range from 350 to 500 energy units, depending on the instance group, and charging powers correspond to typical EV charging levels. Customer demands are randomly generated within predefined ranges, and each instance includes a vehicle-level carbon constraint that limits charging-related emissions along a route. The planning horizon is discretized into 10-minute intervals ($\Delta = 10$) to build the time-expanded network used in the MILP formulation (Section 3). For each station and time step, electricity prices and carbon intensities are assigned. Prices follow a time-of-use structure (off-peak, mid-peak, peak), while carbon intensities vary over time to reflect changes in the electricity generation mix.

Electricity supply is modeled with two charging tiers: clean (low-carbon) and dirty (higher-emission). Emission factors for each station and time step are derived from a base carbon-intensity profile using deterministic scaling factors. Travel times and energy consumption are computed using a physics-inspired traffic model based on Euclidean distances and stochastic traffic density for each arc and departure time.

This mechanism produces realistic variations in travel speeds and travel times across the network. Travel durations are discretized according to the time-step size of the time-expanded network. Energy consumption depends on both travel distance and speed, capturing the effect of congestion on vehicle energy usage. The experimental dataset consists of six instance groups with increasing network size and operational complexity. These groups differ in the number of customers, charging stations, vehicles, time-horizon length, customer spatial dispersion, battery capacities, and carbon limits. In total, 100 instances were generated using different random seeds, ensuring variability in spatial layouts, demand realizations, and dynamic electricity conditions (Table 2).

Table 2. Summary of TE-CC-EVRP instance generation parameters

Instance Group	# Instances	# Customers	# Stations	# Vehicles	T_{max}
sA	20	4	3	3	700
sB	20	5	3	4	900
sC	20	6	3	4	1000
mA	15	7	4	4	1100
mB	15	8	4	5	1200
mC	10	9	4	5	1300

All instances were generated using fixed random seeds to ensure complete reproducibility of the computational experiments. The resulting datasets exhibit diverse spatial configurations, heterogeneous demand distributions, and time-dependent electricity price and emission patterns. This diversity enables a thorough evaluation of the proposed solution methodology across a wide range of operational conditions.

5.2 Performance Evaluation

Table 3 summarizes the computational results obtained for all generated instances. The experiments include 100 instances distributed across the six instance groups described in Table 2. For every instance, both the MILP formulation and the proposed parallel memetic genetic matheuristic were executed, allowing a direct comparison of solution quality and computational performance across different problem sizes.

Table 3: Aggregate performance comparison of MILP and the proposed genetic matheuristic by instance group (corrected solver labels)

Group	# Inst.	GA better	MILP better	Tie	Avg GA Obj.	Avg MILP Obj.	Avg MILP Time (s)	Avg GA Time (s)
sA	20	8	5	7	22.75	22.87	482.77	902.45
sB	20	12	6	2	27.50	28.51	1787.45	905.37
sC	20	10	5	5	34.70	35.12	1874.73	906.66
mA	15	5	7	2	40.57	39.19	1991.81	920.25
mB	15	6	5	0	99.24	120.26	2332.39	939.80
mC	10	5	1	1	55.00	68.22	2592.01	990.24
Total / Avg.	100	46	29	17	40.69	35.69[a]	1734.77	920.69

Both methods produce feasible routing and charging plans across all benchmark instances but differ in solution quality and runtime. Overall, the genetic matheuristic achieves better objective values in 46 instances, compared to 29 for the MILP, while both obtain identical results in 17 cases. These results show that the matheuristic frequently improves upon MILP solutions within the given computational limits. Performance differences also depend on instance size. For the smallest instances (group sA), both methods perform similarly, with nearly identical average objective values (22.75 for GA and 22.87 for MILP) and a balanced number of best solutions, indicating that both approaches can effectively explore the solution space when the problem size is small.

For the intermediate instance groups sB and sC, the genetic matheuristic shows a clearer advantage in terms of solution quality. In group sB, the matheuristic produces better solutions in 12 out of 20 instances, while the MILP performs better in six instances. A similar pattern appears in group sC, where the matheuristic again outperforms the MILP in ten instances and ties occur in several cases. These results suggest that the evolutionary search mechanism becomes increasingly beneficial as the interactions between routing and charging become more complex.

The medium-sized instance groups mA, mB, and mC further illustrate the complementary strengths of the two approaches. In group mA, the MILP slightly outperforms the matheuristic, achieving better solutions in seven instances compared to five for the GA. In contrast, the matheuristic shows stronger performance in groups mB and mC. In group mB, the average objective value obtained by the matheuristic (99.24) is substantially lower than that of the MILP (120.26), indicating a clear improvement in solution quality. A similar pattern is observed in group mC, where the matheuristic achieves better solutions in 5 instances than the MILP does in 1. These findings suggest that heuristic search becomes more effective as the problem size and structural complexity increase. The computational time comparison also reveals a consistent difference between the two approaches. On average, the MILP requires 1734.77 seconds per instance, whereas the genetic matheuristic requires approximately 920.69 seconds. The MILP runtime increases considerably as instance size grows, reaching more than 2500 seconds on average for the largest instance group. The matheuristic runtime remains stable across all instance groups, staying close to the 900-second limit due to the evolutionary algorithm's fixed-generation design. Results show that the proposed parallel memetic genetic matheuristic is an effective alternative to the exact MILP formulation. While MILP performs competitively on small instances, its runtime grows substantially with problem size. In contrast, the matheuristic maintains stable runtimes and produces high-quality solutions, particularly for medium and large instances with complex routing, congestion, and charging interactions.

6. Conclusion

This study introduced the Time-Expanded Carbon- and Congestion-Aware Electric Vehicle Routing Problem with station capacity and a heterogeneous fleet (TE-CC-EVRP), which extends classical EVRP models by jointly capturing routing, charging, congestion, time-dependent travel times, time-varying electricity prices, and carbon intensities, station capacity limits, heterogeneous fleets, and a carbon budget. These elements are embedded in a unified time-expanded network representing electric last-mile delivery operations.

We developed two complementary approaches: a mixed-integer linear programming (MILP) formulation to provide exact benchmarks for small and medium instances, and a parallel memetic genetic matheuristic for larger problems. The method efficiently explores the complex routing-energy space by combining an evolutionary route search with an exact charging subproblem, enhanced by surrogate screening, local improvement, memorization, and parallel evaluation. The computational results demonstrate that the proposed matheuristic constitutes a strong alternative to the direct solution of the full MILP. Computational results show that the matheuristic frequently outperforms the MILP under the same computational limits while maintaining lower and more stable runtimes. Its advantage increases with instance size and complexity, where the MILP runtime grows rapidly but the matheuristic continues to produce high-quality solutions.

Beyond computational performance, the study also highlights the managerial relevance of jointly modeling transportation and energy-system dynamics. Accounting for traffic congestion, station limits, and fluctuating electricity attributes yields solutions that differ significantly from static assumptions. The proposed framework, therefore, provides a useful decision-support tool for logistics operators seeking to balance cost efficiency, charging feasibility, and environmental performance in electric fleet operations.

Several directions remain open for future research: (i) incorporating stochastic models for traffic, pricing, and demand; (ii) integrating richer charging models that account for nonlinear functions and battery degradation; (iii) developing multi-objective variants to analyze trade-offs between cost, emissions, and service; and (iv) testing the methodology on real-world urban freight datasets.

References

Abassi, B., Turki, S., and Dellagi, S., Optimal multi-period manufacturing–remanufacturing–transport planning in carbon conscious supply chain: An approach based on prediction and optimization, Sustainability, vol. 17, no. 11, 5218, 2025.

Cattaruzza, D., Absi, N., Feillet, D., and Vidal, T., A memetic algorithm for the multi-trip vehicle routing problem, European Journal of Operational Research, vol. 236, no. 3, pp. 833–848, 2014.

Felipe, Á., Ortuño, M. T., Righini, G., and Tirado, G., A heuristic approach for the green vehicle routing problem with multiple technologies and partial recharges, Transportation Research Part E: Logistics and Transportation Review, vol. 71, pp. 111–128, 2014.

Goeke, D., and Schneider, M., Routing a mixed fleet of electric and conventional vehicles, European Journal of Operational Research, vol. 245, no. 1, pp. 81–99, 2015.

Grigorev, A., Mao, T., Berry, A., Tan, J., Purushothaman, L., and Mihaita, A. S., How will electric vehicles affect traffic congestion and energy consumption: An integrated modelling approach, Proceedings of 2021 IEEE International Intelligent Transportation Systems Conference (ITSC), pp. 1635–1642, IEEE, 2021.

Hintermann, B., Pass-through of CO2-emission costs into hourly electricity prices in Germany, Journal of the Association of Environmental and Resource Economists, vol. 3, no. 4, pp. 857–891, 2016.

Huang, Y., A review of the electric vehicle routing problem and its variations, Frontiers in Social Science and Technology, vol. 5, no. 5, pp. 85–90, 2023.

Ichoua, S., Gendreau, M., and Potvin, J. Y., Vehicle dispatching with time-dependent travel times, European Journal of Operational Research, vol. 144, no. 2, pp. 379–396, 2003.

Keskin, M., and Çatay, B., Partial recharge strategies for the electric vehicle routing problem with time windows, Transportation Research Part C: Emerging Technologies, vol. 65, pp. 111–127, 2016.

Kunnapapdeelert, S., Johnson, J. V., and Phalitnonkiat, P., Green last-mile route planning for efficient e-commerce distribution, Engineering Management in Production and Services, vol. 14, no. 1, pp. 1–12, 2022.

Lu, J., Chen, Y., Hao, J. K., and He, R., The time-dependent electric vehicle routing problem: Model and solution, Expert Systems with Applications, vol. 161, 113593, 2020.

Labadi, N., Prins, C., and Reghioui, M., A memetic algorithm for the vehicle routing problem with time windows, RAIRO-Operations Research, vol. 42, no. 3, pp. 415–431, 2008.

Lin, B., Ghaddar, B., and Nathwani, J., Electric vehicle routing with charging/discharging under time-variant electricity prices, Transportation Research Part C: Emerging Technologies, vol. 130, 103285, 2021.

Montoya, A., Guéret, C., Mendoza, J. E., and Villegas, J. G., The electric vehicle routing problem with a nonlinear charging function, Transportation Research Part B: Methodological, vol. 103, pp. 87–110, 2017.

Peng, B., Zhang, Y., Gajpal, Y., and Chen, X., A memetic algorithm for the green vehicle routing problem, Sustainability, vol. 11, no. 21, 6055, 2019.

Sadeghi-Barzani, P., Rajabi-Ghahnavieh, A., and Kazemi-Karegar, H., Optimal placement and sizing of fast charging stations, Applied Energy, vol. 125, pp. 289–299, 2014.

Schneider, M., Stenger, A., and Goeke, D., The electric vehicle routing problem with time windows and charging stations, Transportation Science, vol. 48, no. 4, pp. 500–520, 2014.

Sweda, T. M., and Klabjan, D., Finding minimum-cost paths for electric vehicles, Proceedings of 2012 IEEE International Electric Vehicle Conference, pp. 1–4, IEEE, 2012.

Zhang, Y., Gatsis, N., and Giannakis, G. B., Robust energy management for microgrids with high-penetration renewables, IEEE Transactions on Sustainable Energy, vol. 4, no. 4, pp. 944–953, 2013.

Biographies

Ecem Yavuzdoğan is a senior Industrial Engineering student at Özyeğin University. Her academic work centers on routing problems and data science, with a particular interest in using analytical methods to solve large-scale industrial challenges. She is working on several research projects following a successful research internship in optimization.

Ayşenur Bal is a senior student in Industrial Engineering at Özyeğin University. Her academic interests include operations research, mathematical optimization, and data-driven decision making. She is involved in research focusing on the development and application of optimization models for complex industrial and logistical systems. She aims to pursue graduate studies in industrial engineering and contribute to research in this area.

Milad Elyasi is an Assistant Professor in the Industrial Engineering Department at Özyeğin University, Istanbul, Turkey. He received his Ph.D. in Industrial Engineering. His research interests include operations research, mathematical optimization, transportation and energy systems, and decision analytics. His work focuses on optimization models and algorithms for industrial and energy problems, including sustainable manufacturing, energy-efficient system design, and large-scale optimization under uncertainty. He has published in international journals and conference proceedings and supervises graduate research