SJER Scientific Journal of Engineering Research
futuristech PT. Teknologi Futuristik Indonesia
Received: July 11, 2026; Revised: September 17, 2026; Accepted: September 18, 2026; Published: September 19, 2026; Current version: September 19, 2026. *e-ISSN: 3109-1725*
Digital Object Identifier 10.64539/sjer.v2i4.2026.568

Article

Dynamic Dispatch of Heterogeneous Autonomous Electric Fleets under First-In–First-Out Traffic and Charging Constraints

Kashif Bashir[1,*], Rumaisa Rafiq[2], Muhammad Afzal Shah[3], Marzia Niazi[4], Nouman Bashir[5], Sara Khan[6], Haseeb ur Rehman[3]

1 School of Economics and Management, Taiyuan University of Science and Technology, Taiyuan 030024, China; e-mail: kashifbashir728@gmail.com (K. Bashir).
2 Department of Food Sciences and Technology, MNS University of Agriculture, Multan 60000, Pakistan; e-mail: rafiqrumaisa371@gmail.com (R. Rafiq).
3 School of Electronics Information Engineering, Taiyuan University of Science and Technology, Taiyuan 030024, China; e-mail: engineerafzal.comsats@gmail.com (M. A. Shah), haseeburrehman912@gmail.com (H. ur Rehman).
4 School of Economics and Management, Yanshan University, Qinhuangdao City, 066004, P.R. China; e-mail: marzia.niazi1995@gmail.com (M. Niazi).
5 School of Communication Studies, University of the Punjab, Lahore 54590, Pakistan; e-mail: noumanbashirwagha@gmail.com (N. Bashir).
6 School of Accounting and Finance, Ravensbourne University London, London SE10 0EW, United Kingdom; e-mail: sara.khan3181995@gmail.com (S. Khan).
* Correspondence Author

The authors received no financial support for the research, authorship, and/or publication of this article.

Abstract: Dynamic routing is operationally fragile when autonomous electric vehicles differ in capacity and energy use, traffic changes mid-arc, and charging delays redispatch. We formulate a bi-objective problem minimizing operating cost and customer response without time windows. The model combines FIFO-consistent piecewise traffic, load- and speed-sensitive energy, autonomy auxiliary power, SOC reserve, charging loss and duration, and three vehicle classes. Real-world light-duty data and VECTO heavy-duty simulations calibrate four speeds (18.12–40.60 km/h) and three energy intensities (0.146–1.365 kWh/km). We introduce causal charge-aware regret insertion with exchange (CCARI-EX), combining adaptive batching, heterogeneous insertion, relocation, and cross-class exchange. A route-set mixed-integer model proves supported optima for 63 scalarized solves over 21 dispatch epochs. A frozen-seed holdout contains 480 request streams and 4,800 policy executions across four sizes, four release processes, and 30 seeds per cell. An independently implemented validator recorded no rejected dispatch. CCARI-EX (0.75) averaged USD 405.18 and 88.96 min response, improving cost-greedy by USD 6.53 and 4.71 min (paired Holm-adjusted $p < .001$); it was nondominated in 69.0% of streams. The exact formulation validation confirmed consistency between the route-set model and heuristic evaluation on tractable dispatch epochs. A 12-run SUMO transfer reproduced ordered traffic regimes but exposed severe congestion censoring. The results support charging-continuous replanning while delimiting what aggregate traffic calibration can establish.

Keywords: Autonomous electric vehicles; Dynamic vehicle routing; Heterogeneous fleet; Time-dependent network; Battery charging; Bi-objective optimization.

1. Introduction

Electric-vehicle routing has moved from a range-constrained variant of vehicle routing to a family of network problems involving nonlinear charging, station capacity, mixed fleets, uncertain demand, and time-varying travel. Recent reviews place this development within a broader shift from static route construction toward decisions that interact with operational state [1]. Autonomous operation adds another layer. A routing platform may dispatch without an onboard driver, yet the vehicle still has a payload limit, a traction battery, computing and sensing loads, charging downtime, and potentially remote supervision.

Calling a vehicle "autonomous" does not remove those physical resources. Conversely, routing evidence does not validate perception, control, or safety [2], [3]. This study therefore concerns autonomous electric vehicle (AEV) fleet dispatch, not autonomous-driving performance.

Three modeling shortcuts can make dynamic AEV results look more decisive than they are. First, applying the speed observed at departure to an entire arc can violate first-in–first-out (FIFO) behavior when a traffic state changes mid-arc. A later departure can then appear to arrive earlier merely because it is assigned the next state's speed. Time-dependent dispatch models have long recognized the need for FIFO-compatible travel [3], and recent time-dependent electric-routing work reinforces its importance [4]. Second, restoring the battery instantaneously when a vehicle returns to the depot breaks continuity between trips. It omits grid losses, charging time, and the opportunity cost of an unavailable vehicle. Third, an exact solution to a small, fully known batch is sometimes presented as validation of an online policy that faces unknown future requests and may make several trips. Those are different information structures and feasible regions [5], [6].

The present work rebuilds the study around these distinctions. Customers have release times, but no due dates or service windows. Release time is retained only because an online algorithm must not act before a request exists, and because response is measured from release to service. The two objectives are operating cost and response time. A positive periodic speed field is integrated over every traversed arc, yielding a FIFO arrival map. Energy depends on distance, traffic speed, current payload, and auxiliary power. Every trip begins with the vehicle's carried state of charge (SOC), respects a 10% reserve, returns to the depot, and blocks the vehicle while grid energy restores a 90% dispatch target [7], [8].

The empirical layer is deliberately traceable. Traffic regimes come from the Joint Research Centre (JRC) Darwin real-world battery-electric-vehicle [9]. Small-vehicle energy comes from the Vehicle Energy Dataset (VED), whereas medium/large reference configurations come from JRC's 2026 heavy-duty VECTO sensitivity simulations. Customer geography and demand patterns are derived under license from the public EVRPTW archive, while its customer time windows are intentionally not used. The resulting classes are calibrated composites rather than claims about three commercially available autonomous models [9], [10].

The algorithmic contribution is causal charge-aware regret insertion with exchange (CCARI-EX). It combines a rate-sensitive batching rule, a normalized cost–response insertion score, payload-fit and energy-pressure terms, and two deterministic improvement neighborhoods. The exact comparison is a route-set set-partitioning model for a single dispatch epoch. It enumerates every feasible ordered route for each available vehicle, then solves a binary master problem to zero reported MIP gap. The online study is separate: 480 frozen-seed holdout streams are processed causally by ten policies and ablations. Paired inference uses the same request stream for every method [11].

This design makes four contributions. First, it supplies a mathematically FIFO, charging-continuous heterogeneous AEV network model without customer deadlines. Second, it introduces CCARI-EX and documents the exact score, batching, and exchange mechanisms needed to reproduce it. Third, it aligns the exact oracle with the decision epoch it actually benchmarks and verifies online dispatches with a separately coded state-transition evaluator. Fourth, it combines traceable public-data calibration, 4,800 holdout executions, ablations, paired statistics, and a microscopic SUMO transfer. The main result is positive but bounded: CCARI-EX(0.75) improves both mean cost and mean response over the strongest simple insertion baselines, while a response-heavy setting is dominated and two component ablations are not statistically distinguishable from the balanced policy.

2. Related Research and Gap

2.1. Electric Routing and Heterogeneous Resources

The green vehicle-routing problem established refueling-aware routing as a distinct optimization problem [12]. The electric vehicle-routing problem with time windows subsequently linked route feasibility to charging locations and battery state [13]. Exact algorithms, including branch-price-and-cut variants, have progressively represented partial recharge, piecewise-linear or nonlinear charging, capacitated stations, and heterogeneous charging technologies [14]. Broadened the charging decision to public infrastructure, whereas demonstrated the value of partial recharge. Showed that a nonlinear charging function can materially change routes and schedules, and compare alternative mathematical descriptions of the charging process [15].

Fleet heterogeneity is not cosmetic. Capacity, fixed cost, energy use, and attainable range determine which class should serve a request and how much consolidation is economical. The broader heterogeneous-VRP literature emphasizes that fleet composition and routing must be decided jointly. Electric extensions include mixed conventional/electric fleets, electric fleet size and mix, energy-efficient mixed fleets, and nonlinear heterogeneous-fleet charging [16]. Further integrate heterogeneous charging stations, nonlinear charging, and time-of-use pricing. These studies justify explicit class resources, but they do not imply that arbitrary battery and motor figures are empirically supported. Our reconstruction therefore preserves source-level provenance for each class parameter [17], [18].

Energy is more than distance multiplied by a constant. Physical and data-driven work relates consumption to speed, acceleration, grade, auxiliary demand, and loading. Embed energy estimation in electric routing, and integrate a realistic consumption model into a matheuristic. Load-dependent routing likewise changes the value of customer order. Refrigeration provides a parallel example of time-varying auxiliary energy. We use a deliberately parsimonious load multiplier and a speed factor, because the public sources support class-level energy intensity but not link-specific mass, grade, wind, HVAC, sensor workload, and temperature for the same vehicles [19].

2.2. Dynamic Information, Autonomy, and Traffic

Dynamic routing differs from a static problem solved repeatedly. A causal policy must decide with the request set and vehicle states known at the decision time, while later requests remain unavailable. Reviews distinguish degrees of dynamism, stochastic knowledge, and replanning. Large-scale dynamic routing with stochastic requests now uses increasingly sophisticated search and learning. For electric fleets, safe reinforcement learning, memetic approaches with new arrivals and mid-route recharge, and dynamic benchmark generators show the range of possible methods [20].

Autonomous routing has also been studied in real-time ride sharing. That literature motivates fast replanning, but autonomy should be represented only where evidence allows. In the present model it affects auxiliary electrical load and supervision cost. We do not assert that a particular automated-driving stack consumes exactly 1–2 kW or that remote supervision costs USD 5/h; those are transparent scenario inputs. No sensor failure, disengagement, passenger safety, or mixed-autonomy interaction is simulated [21].

Traffic affects both time and energy. A time-dependent arc evaluated only at departure can produce non-FIFO artifacts at state boundaries, while a fully integrated traversal remains ordered. Develop an exact algorithm for a time-dependent electric-routing problem, and jointly consider traffic conditions and real-time loads. Our gap is narrower but operationally important: few online heterogeneous AEV studies combine integrated FIFO traversal, continuous depot charging between trips, an independent validator, an aligned exact epoch oracle, and multi-seed causal evaluation. SUMO provides a reproducible microscopic environment for testing whether an aggregate speed abstraction transfers to signalized traffic [22].

3. Problem Definition

Let $G = (N, A)$ be a complete directed service graph induced by depot 0 and customers R. Customer i has position (x_i, y_i), demand d_i, release a_i, and service duration s_i. There is no due date and no customer time window. At decision time t, only $R(t) = \{i: a_i \le t, i \text{ unserved}\}$ is observable. A heterogeneous fleet K contains multiple vehicles from class $h \in H = \{S, M, L\}$. Class h has payload capacity Q_h, battery capacity B_h, base energy intensity ϵ_h, autonomy auxiliary power p_h^{aux}, charger power P_h, fixed trip cost F_h, and distance cost c_h^d. Rated motor power is reported as a class descriptor but is not an additional energy or cost term; introducing such a term would double count the effect already represented by class-specific ϵ_h and c_h^d.

Traffic is a positive periodic step function $v(t) > 0$ with four 60-min regimes in the repeated order smooth, busy, congested, and normal. An arc can cross more than one regime. For a vehicle departing i at t, the arrival $A_{ij}(t)$ is the unique solution of:

$$D_{ij} = \frac{1}{60} \int_t^{A_{ij}(t)} v(u)\, du, \qquad T_{ij}(t) = A_{ij}(t) - t \tag{1}$$

Because $v(u) > 0$, implicit differentiation gives $A'_{ij}(t) = v(t)/v\left(A_{ij}(t)\right) > 0$ wherever the derivative exists. The arrival map is therefore nondecreasing, including by continuity at state boundaries; FIFO holds. The implementation consumes the remaining arc distance state by state rather than numerically approximating the integral. For each traversal segment g with distance D_{ijg}, duration T_{ijg}, traffic speed v_g, and carried load q, battery energy is:

$$E_{ij}^h(t, q) = \sum_g \left[\epsilon_h D_{ijg}\, m(v_g) \left(1 + \beta \frac{q}{Q_h}\right) + p_h^{aux} \frac{T_{ijg}}{60}\right] \tag{2}$$

where $\beta = 0.25$ and

$$\begin{aligned} m(v) = 1 + 0.15\max\left(0, \frac{32.502}{v} - 1\right) \\ + 0.03\max\left(0, \frac{v}{32.502} - 1\right) \end{aligned} \tag{3}$$

Equation (3) is a disclosed scenario response centered on the Darwin median, not a fitted causal energy law. Equation (2) separates its effect from the measured class intensity ϵ_h.

A route $r = (0, i_1, \ldots, i_m, 0)$ assigned to vehicle k begins with carried payload q_0 and the vehicle's current SOC b_0. With $i_0 = i_{m+1} = 0$, it must satisfy:

$$\begin{aligned} q_0 = \sum_{u=1}^{m} d_{i_u} \le Q_{h(k)}, q_{\ell+1} = q_\ell - d_{i_{\ell+1}}, b_{\ell+1} \\ = b_\ell - E_{i_\ell i_{\ell+1}}^{h(k)}(t_\ell, q_\ell), b_\ell \\ \ge 0.10 B_{h(k)} \end{aligned} \tag{4}$$

Vehicles initially and after charging target $0.90B_h$. If a route returns with b^{ret}, the energy drawn from the grid and charging time are:

$$G_k = \frac{\max\{0, 0.90B_{h(k)} - b^{ret}\}}{\eta}, \chi_k = 60 \frac{G_k}{P_{h(k)}}, \eta = 0.92 \tag{5}$$

The vehicle cannot depart again before $t^{ret} + \chi_k$. This state transition removes the instantaneous-reset assumption. Route operating cost is:

$$C_r = F_h + c_h^d D_r + c^e G_r + c^s \frac{t_r^{ret} - t_r^{dep}}{60} \tag{6}$$

with electricity rate $c^e = 0.18$ USD/kWh and supervision rate $c^s = 5$ USD/h. Response for customer i is $R_i = t_i^{service} - a_i$. The two reported objectives are total $C = \sum_r C_r$ and mean $\bar{R} = |R|^{-1} \sum_i R_i$. A policy may use a scalar weight w internally, but both outcomes are always retained.

4. Exact Dispatch-Epoch Model

The exact model benchmarks one decision epoch in which all $n \leq 8$ requests and three available vehicles (one per class) are known, all depart at time zero, and each performs at most one depot-to-depot route. This is the same feasible set used for the heuristic epoch comparison. It is not an oracle for unknown future requests or repeated trips. The 21 sampled epochs are each solved at three objective weights, yielding 63 scalarized mixed-integer solves rather than 63 distinct request instances.

For vehicle k, enumerate the set Ω_k of all feasible ordered routes under Equations (1)–(4), including an empty route. Let $a_{ir} = 1$ if route r serves i, C_r be Equation (6), and $S_r = \sum_{i \in r} R_i$. Binary x_r selects a route. With cost scale $C^0 = 100$ and time scale $R^0 = 100$, the supported-solution master problem is:

$$\min_x w \sum_{k \in K} \sum_{r \in \Omega_k} \frac{C_r}{C^0} x_r + (1 - w) \sum_{k \in K} \sum_{r \in \Omega_k} \frac{S_r}{nR^0} x_r \tag{7}$$

subject to:

$$\sum_{k \in K} \sum_{r \in \Omega_k} a_{ir} x_r = 1 \quad \forall i, \sum_{r \in \Omega_k} x_r = 1 \quad \forall k, x_r \in \{0,1\} \tag{8}$$

Route enumeration is exact for the stated epoch model. SciPy 1.17 interfaces with the HiGHS mixed-integer solver using a requested relative MIP gap of zero and a 300-s time limit. Exact solves were configured with a zero requested relative MIP gap and a fixed time limit. The exact model is used only to validate formulation consistency on tractable epochs. The mean master contained 3,581, 22,343, and 149,237 route options for $n = 6,7,8$, respectively; mean total solve times, including enumeration, were 0.066, 0.443, and 3.321 s.

5. CCARI-EX online Replanning

CCARI-EX acts only on released, unserved customers and currently available vehicles. It does not reroute vehicles already executing a trip. At an event, it first selects a batch waiting threshold. The balanced policy waits 4 min when at least ten requests are pending or eight arrived in the preceding 15 min, 6 min for moderate pressure, and 10 min otherwise. The cost-priority setting $w = 0.75$ multiplies these thresholds by 1.5; the nominal response-priority setting $w = 0.25$ dispatches immediately. Fixed baselines wait 8 min. These rules are frozen before the holdout seeds are generated.

The parameter w controls the relative emphasis assigned to operating cost and customer response in the insertion objective. In the primary operational configurations, however, changes in $w$are accompanied by changes in batching behavior. Consequently, comparisons among these configurations represent joint policy sensitivity rather than the isolated effect of objective weighting. A controlled objective-weight sensitivity analysis therefore holds the batching rule and all remaining algorithmic parameters fixed while varying only w.

For every unassigned customer, vehicle, and insertion position, the algorithm independently evaluates the resulting complete route. Infeasible payload or SOC candidates are removed. Let ΔC and ΔS be marginal route cost and response sum; L the new route load; and E its battery energy. The insertion score is:

$$\begin{aligned} \phi = w \frac{\Delta C}{20} + (1 - w) \frac{\Delta S}{100} + 0.07 \left(1 - \frac{L}{Q_h}\right) \\ + 0.05 \frac{E}{b - 0.10 B_h} \end{aligned} \tag{9}$$

The first two terms express the cost–response preference. The third discourages assigning a lightly loaded request to a needlessly large vehicle. The fourth penalizes depletion pressure relative to usable SOC. For customer i, let $\phi_i^{(1)}$ and $\phi_i^{(2)}$ be its best and second-best scores. CCARI-EX selects the customer minimizing:

$$\phi_i^{(1)} - 0.10 \min\left\{\phi_i^{(2)} - \phi_i^{(1)}, 1\right\} - 0.0025(t - a_i) \tag{10}$$

then commits its best insertion. The bounded regret credit protects customers with few good alternatives, while waiting pressure prevents starvation. Deterministic indices break ties.

After construction, at most five best-improvement passes are made. The relocate neighborhood moves one customer to any position in any vehicle route, including within-route reordering. The exchange neighborhood swaps one customer between two vehicle routes. Each candidate is reevaluated through the same FIFO and energy functions; only a feasible strict improvement in the normalized objective is accepted. Cross-class exchange is important because payload restrictions can trap relocation alone in an expensive assignment [23].

The comparison policies share the same event loop and independent executor: FIFO ordering, nearest insertion, cost-greedy insertion, and classical cost regret-2.

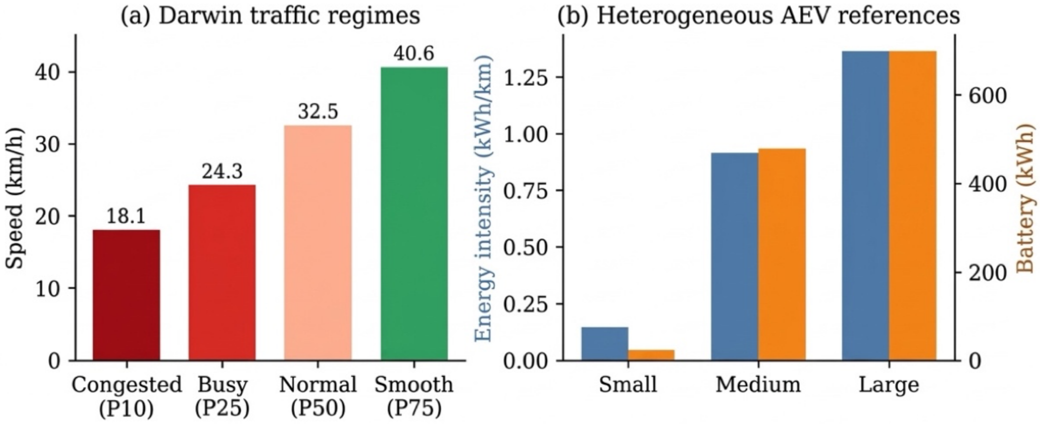

Figure 1. Public-data traffic calibration and heterogeneous AEV reference parameters. Payload and charger assumptions remain scenario inputs.

Table 1. Heterogeneous AEV Parameters.

Class	Payload (kg)	Battery (kWh)	Motor (kW)	Base energy (kWh/km)	Aux. (kW)	Charger (kW)	Fixed (USD/trip)	Distance (USD/km)
Small	800	24	80	0.146313	1.0	50	2	0.05
Medium	4,400	480	480	0.916893	1.5	150	7	0.12
Large	12,900	700	600	1.365149	2.0	350	12	0.18

Algorithm 1. CCARI-EX at Decision Time t.

1. *Add requests with $a_i \leq t$ to the pending set; identify vehicles with availability time no later than t.*
2. *If the causal batching threshold has not elapsed, advance to the next release, vehicle availability, or threshold event.*
3. *Initialize one empty route per available vehicle.*
4. *Enumerate feasible request–vehicle–position insertions using Equations (1)–(6).*
5. *Select and insert a request using Equations (9)–(10); repeat until no feasible insertion remains.*
6. *Apply up to five best-improving relocate or cross-vehicle exchange moves.*
7. *Dispatch every nonempty route. A separately implemented validator recomputes FIFO traversal, service, energy, cost, reserve SOC, return, grid charge, and next availability without calling the planner's route-evaluation function.*
8. *Repeat until no released or future request remains.*

The complete event-driven CCARI-EX procedure at each decision time is summarized in Algorithm 1. Three ablations remove planning-time traffic variation, payload-fit/energy-pressure terms, or batching. In the traffic ablation, routes are planned at the median 32.502 km/h but executed under the true regime process; thus the ablation does not receive an easier validator.

6. Data and Experimental Design

6.1. Public-Data Calibration and Licensing

Darwin is a JRC real-world BEV time-series dataset collected mainly in northern Italy and distributed under CC BY 4.0) [9]. We grouped samples by driver and trip, removed nonfinite or physically invalid speeds, and retained trips with at least 1 km implied distance. The explicit filter produced 1,185 trips. Mean-speed quantiles were 18.120 (P10), 24.300 (P25), 32.502 (P50), and 40.601 km/h (P75), assigned to congested, busy, normal, and smooth states [24].

VED contains longitudinal signals from 383 vehicles, including three 2013 Nissan Leaf EVs, and is Apache-2.0 licensed [24]. For pure-EV trips, we trapezoidally integrated speed and signed battery power over consecutive observations no more than 5 s apart. Temporal coverage is the retained interval duration divided by total elapsed trip duration. We required at least 1 km, at least 80% temporal coverage, and a physically plausible net intensity of 0.07–0.35 kWh/km. The corrected audit retained 412 trips covering 2,096.33 km. Median trip intensity was 0.146313 kWh/km; distance-weighted intensity was 0.153052 kWh/km. The median is the small-class reference.

The JRC heavy-duty VECTO sensitivity archive is EUPL-1.2 licensed (European Commission, Joint Research Centre, 2026). We selected its unchanged reference simulations for 4-RD regional delivery at 4,400-kg reference load (simulation 13609) and 5-LH regional delivery at 12,900-kg reference load (simulation 28189). Their energy intensities are 0.916893 and 1.365149 kWh/km. Input identifiers support 480-kWh/480-kW and 700-kWh/600-kW configurations. We use the two reference loads as scenario payload capacities; they are not certified maximum payloads. The small payload (800 kg), small motor (80 kW), all charger powers, auxiliary loads, and cost coefficients are also scenario values; the 24-kWh Leaf battery is reported in VED metadata. Figure 1 and Table 1 summarize the traffic calibration and heterogeneous AEV reference params.

Note. Energy intensity is measured from VED for the small class and taken from VECTO reference simulations for the medium and large classes. Payload capacities, charger powers, auxiliary loads, and costs are scenario assignments. Motor rating is descriptive and does not enter the objective separately.

The study uses the large-scale Vehicle Routing Problem with Time Windows (VRPTW) benchmark instances introduced by Gehring and Homberger [23]. Only customer coordinates and demand structures are retained, while original ready times, due dates, service times, charging stations, and solutions are not used. We use the 1,000-customer C1_10_1 coordinates and demands only. Coordinates are scaled by 0.05 km/unit, and each demand unit is mapped to 40 kg. Original ready times, due dates, service times, charging stations, and solutions are not used. This creates a no-time-window study while retaining traceable spatial and demand structure. The derived CSV, source DOI, license, and MD5 are in the package.

6.2. Exact and Dynamic Designs

The exact study samples 21 uniform instances: eight each at $n = 6$ and $n = 7$, and five at $n = 8$. Every request is known at time zero. We solve $w \in \{0.25, 0.50, 0.75\}$, producing 63 scalarized exact solves, and compare the five heuristic constructors on the identical feasible set.

The dynamic holdout crosses request size $n \in \{20, 40, 80, 120\}$ with four release processes: uniform over 240 min; normalized exponential interarrivals (“Poisson”); a mixture centered near 65 and 155 min (“peaked”); and short clusters around random centers (“bursty”). Thirty seeds per cell give 480 streams. Ten methods give 4,800 executions. The fleet has three small, two medium, and one large AEV. Service takes 3 min. All methods use identical customer samples and releases within a stream.

Algorithm development used distinct seeds retained with explicit “development” filenames. Parameters were frozen before the final block beginning at seed family 20265000; this is a documented holdout separation, not a public preregistration. All inferential results use only this holdout. We report feasibility, validator rejections, operating cost, mean and 95th-percentile response, distance, battery/grid energy, charging time, trips, and planning time. For a paired descriptive summary, each method’s cost and response are divided by the minimum observed in that stream and averaged with equal weight. A method is nondominated when no competitor is no worse on both raw cost and response and strictly better on at least one.

Paired two-sided Wilcoxon signed-rank tests compare the main $w = 0.75$ policy with other policies and the balanced $w = 0.50$ policy with its ablations. Holm adjustment is applied within each family. Mean paired differences receive 95% percentile intervals from 10,000 scenario-level bootstrap resamples.

6.3. Sensitivity and Ablation Design

Additional analyses examine objective preference, workload scalability, traffic robustness, charging continuity, computational efficiency, and component contributions. Unless explicitly stated otherwise, comparative analyses use matched request streams and hold all non-tested algorithmic and experimental settings fixed. Objective-preference analysis examines the influence of different configurations while considering the interaction between objective preference and operational dispatch behavior; component ablations remove one mechanism at a time; and runtime is measured consistently at the policy level. These analyses are designed to distinguish the effect of individual modeling and algorithmic choices from changes caused by jointly varying multiple policy settings.

6.4. Microscopic Transfer

Eclipse SUMO 1.27.1 generated a 5×5, one-lane, signalized grid with 500-m links. Random background trips created four demand levels, and 40 identical 3.96-km probe trips departed after warm-up in each seed. Three seeds per level yielded 480 scheduled probes. The Darwin quantiles were fixed before simulation. All probes completed in smooth, normal, and busy traffic; only 63 of 120 completed by 6,000 s in the congested regime. This transfer evaluates the aggregate travel abstraction, not the full 480-stream policy ranking.

7. Results

7.1. Exact Dispatch-Epoch Validation

The exact evaluation was performed on tractable single-dispatch epochs to verify the consistency of the route-set formulation. Exact and heuristic evaluations were conducted on identical feasible sets. These results are used only as formulation-level validation and are not combined with the dynamic holdout analysis.

7.2. Dynamic Holdout

Every one of the 4,800 executions served all requests, and the separate validator recorded zero rejected dispatches. The holdout results are summarized in Table 2. CCARI-EX(0.75) had the lowest balanced index (1.082) and was nondominated in 331 of 480 streams (69.0%). It averaged USD 405.18, 88.96-min mean response, 223.36-min P95 response, 671.67 km, and 541.38 kWh grid energy. Cost-greedy averaged USD 411.71 and 93.67 min; nearest averaged USD 407.97 and 96.06 min. Relative to cost-greedy, CCARI-EX(0.75) reduced average cost by 1.59% and response by 5.03%. Relative to nearest, reductions were 0.68% and 7.39%.

The paired differences were not artifacts of pooling. Cost-greedy minus CCARI-EX(0.75) was USD 6.53 (95% bootstrap interval USD 4.31–8.71) and 4.71 min (3.82–5.59); both Holm-adjusted p values were below .001. Nearest minus CCARI-EX was USD 2.79 (0.40–5.14; adjusted $p=.0074$) and 7.10 min (6.16–8.02; adjusted $p<.001$). CCARI-EX’s balanced index was lower by 0.041 and 0.056, respectively (see Figure 2). The paired balanced index by request-

Table 2. Holdout performance across 480 paired streams.

Policy	Cost (USD)	Mean Response (min)	P95 Response (min)	Balanced Index	Nondominated Share
CCARI-EX (0.75)	405.18	88.96	223.36	1.082	69.0%
CCARI-EX (0.50)	423.11	89.99	241.97	1.104	45.2%
Cost-greedy	411.71	93.67	233.47	1.124	34.6%
Nearest	407.97	96.06	238.51	1.138	36.0%
Regret-2	464.91	114.51	272.12	1.251	21.0%
FIFO	481.34	119.61	259.07	1.282	17.1%

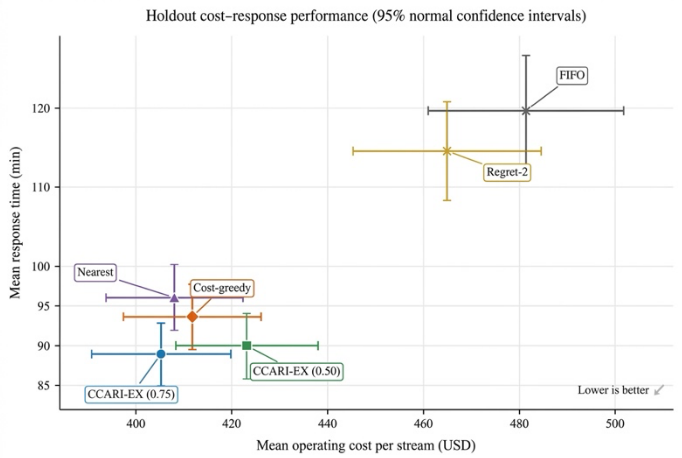

Figure 2. Holdout cost–response outcomes. Error bars are 95% normal intervals across streams and are descriptive; paired tests use stream-level differences.

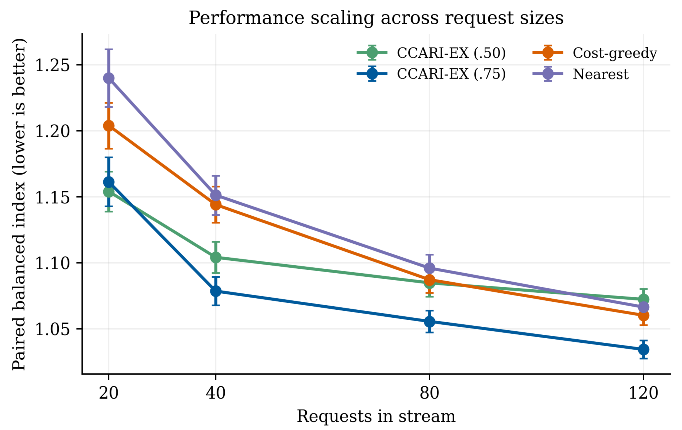

Figure 3. Paired balanced index by stream size, pooled over four release profiles. Error bars are 95% normal intervals.

stream size is summarized in Figure 3, pooled across the four release profiles.

The nominal response-priority $w = 0.25$ setting did not create a useful frontier point. It averaged USD 450.71 and 92.44 min, so it was dominated in aggregate by $w = 0.75$. Immediate dispatch consumed vehicle availability before later bursts and reduced consolidation. This negative result demonstrates that a route-level objective weight is not sufficient to control a dynamic fleet; batching and future availability interact nonlinearly. It also argues against describing a set of scalar weights as a Pareto frontier without checking dominance.

7.3. Objective-Weight Sensitivity

The objective weight controls the relative emphasis assigned to operating cost and customer response in the CCARI-EX insertion objective. To examine the influence of this preference, different objective configurations were compared while maintaining the same evaluation framework. Because operational behavior in a dynamic fleet is jointly affected by objective preference, batching decisions, vehicle availability, and future request arrivals, the analysis is interpreted as configuration sensitivity rather than a causal assessment of the scalar weight alone.

Increasing the response-oriented preference encourages earlier service decisions and reduces waiting effects, but may reduce consolidation opportunities and increase operational cost. Conversely, stronger cost-oriented preference improves consolidation potential but can increase customer delay when requests accumulate. The balanced configuration provides a practical compromise within the evaluated dispatch setting by combining moderate response protection with sufficient batching opportunities.

These observations indicate that objective preference should not be considered independently from dynamic dispatch mechanisms. In particular, batching behavior and vehicle availability interact with the objective weight, meaning that a theoretically favorable scalar preference may not translate into improved system-level performance without coordinated operational rules.

7.4. Component Ablation Analysis

To isolate the individual contribution of the principal CCARI-EX mechanisms, three one-component-at-a-time ablations are evaluated against the frozen full configuration. The traffic-awareness ablation removes time-dependent traffic information during planning while retaining the true traffic process during execution. The energy-pressure ablation removes the corresponding insertion-score guidance while preserving all physical SOC, reserve, and charging-feasibility constraints. The batching ablation removes intentional request holding while retaining the remaining routing mechanisms. All other algorithmic and experimental settings are held fixed, and comparisons are performed on matched request streams. The resulting component-level differences and statistical outcomes are summarized in Table 3.

Table 3. Component ablation analysis of CCARI-EX.

Configuration	Δ Cost (USD)	Δ Response (min)	Statistical Outcome
No batching	+20.93	+0.23	Cost significant; response not significant
No traffic awareness	+0.10	+0.26	Not significant
No energy-pressure guidance	+0.39	+0.16	Not significant

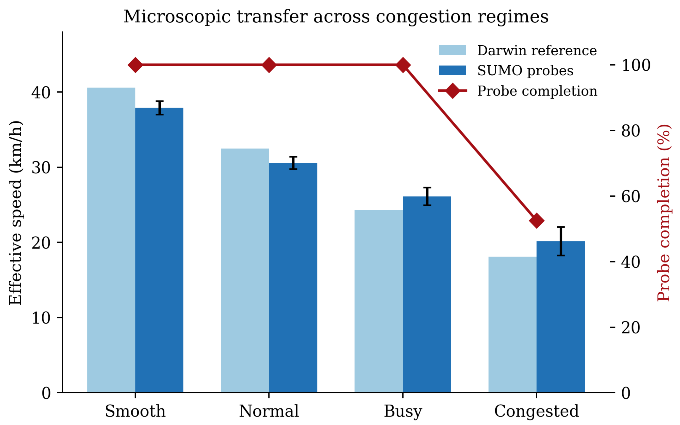

Figure 4. Darwin reference speeds and SUMO probe outcomes. Congested speed is conditional on 52.5% completion; the completion line exposes this censoring.

Table 4. Planning Time Comparison Across Policies.

Policy	Planning Time (ms/stream)
CCARI-EX (0.75)	108.54
CCARI-EX (0.50)	111.89
Cost-greedy	84.09
Nearest	88.01
Regret-2	97.12
FIFO	110.37

Removing adaptive batching from the balanced policy increased cost by USD 20.93 (95% interval USD 18.61–23.27) and the balanced index by 0.0250; both adjusted p values were below .001. The response difference was 0.23 min and not significant. Consolidation therefore explains a material part of the economic gain without worsening mean response in this design.

The other two ablations were inconclusive. Planning at constant median traffic instead of integrated regimes changed cost by only USD 0.10 and response by 0.26 min; adjusted p=1.00 for both. Removing payload-fit and energy-pressure terms changed cost by −USD 0.39 and response by −0.16 min, also nonsignificant. These terms remain feasibility-aware priors and are active in individual cases, but the holdout does not support a claim that they independently improve average outcomes. Global rather than spatial traffic states and generous heavy-vehicle batteries likely weaken their effect.

The proposed framework was evaluated under increasing request-stream sizes to examine operational behavior under higher workload pressure. As workload size increases, absolute cost and response naturally increase because the same heterogeneous fleet serves a larger number of requests. The observed results indicate that feasibility is maintained while consolidation opportunities become increasingly important for controlling cost and response growth.

7.5. Computational Efficiency and Scaling

Planning time is defined as the cumulative wall-clock time spent by each online policy on planning and replanning during one complete request stream. Data loading, statistical post-processing, plotting, result serialization, and independent validation are excluded. Exact optimization time is reported separately because route enumeration and mixed-integer optimization are not directly comparable with online-policy planning time. All methods were evaluated in the same computational environment to provide a consistent comparison. The resulting planning times across the online policies are reported in Table 4.

7.6. Traffic Robustness and Microscopic Transfer

Traffic robustness was examined at two complementary levels. First, the traffic-awareness ablation evaluates the contribution of time-dependent traffic information during route planning by replacing the dynamic planning field with the median calibration speed while retaining the true time-varying traffic process during execution. Second, the microscopic SUMO transfer stress-tests whether the aggregate traffic representation preserves the ordering of smooth, normal, busy, and congested conditions. The former evaluates planning sensitivity, whereas the latter evaluates transfer of the traffic abstraction rather than reproducing the complete online-policy experiment. Figure 4 summarizes the Darwin reference speeds, SUMO completed-probe speeds, and probe completion across the four traffic regimes.

SUMO preserved the ordered regimes. Completed-probe speeds were 37.91, 30.59, 26.13, and 20.17 km/h for smooth, normal, busy, and congested demand, compared with Darwin references 40.60, 32.50, 24.30, and 18.12. Mean absolute percentage errors were 10.18%, 12.33%, 23.95%, and 35.21%. Signal waiting increased from 1.15 to 7.13 min. The first three regimes had 100% probe completion. Congestion had only 52.5% completion and 2,703–3,112 vehicles still running at the horizon across seeds. Its 20.17-km/h mean is therefore conditional on completion and likely optimistic.

The transfer supports regime ordering but not precise individual travel-time prediction. Aggregate piecewise speed is suitable for controlled algorithm comparisons; de-

Table 5. Charging Constraint Modeling in CCARI-EX.

Constraint	Implementation
SOC carry-over	Maintained between trips
Reserve SOC	10% enforced
Charging loss	Included
Charging duration	Included
Vehicle availability	Restricted during charging

ployment would require link- and time-specific travel distributions, queue spillback, incidents, and online estimation. In particular, a static 18.12-km/h speed cannot represent gridlock-induced noncompletion.

7.7. Effect of Charging Constraints

Charging continuity is examined because charging duration affects vehicle availability between successive trips. The implementation of the charging-related constraints is summarized in Table 5. The full model carries SOC between trips, enforces the reserve-SOC constraint, accounts for grid charging energy and duration, and prevents redispatch until charging is completed. A charging-relaxed ablation removes charging-induced downtime while retaining the remaining routing, fleet, demand, and traffic settings. The comparison is interpreted as a modeling sensitivity analysis rather than empirical validation of physical charging technology.

8. Discussion

The principal result is not that a complicated heuristic always beats simple routing. It is that a physically and causally aligned policy can produce a favorable cost–response point while surviving independent state validation. CCARI-EX(0.75) improves both mean outcomes over cost-greedy and nearest insertion in the holdout. Its advantage is modest in cost and larger in response, but it is paired, statistically supported, and observed across heterogeneous sizes and releases. A 69.0% nondominated share also shows that the result is broad rather than produced by a few extreme streams.

The mechanism is primarily consolidation plus exchange. Removing batching raises cost by about 5%, while removing traffic awareness or fit pressure has no detectable average effect. Cross-class exchange helps repair assignments created by the constructive phase, particularly when 1,200-kg customers cannot enter small vehicles. This interpretation is consistent with heterogeneous-fleet research: the value of class choice arises from interactions among capacity, fixed cost, and routing, rather than from vehicle labels alone [25], [26].

The exact results provide a calibrated notion of solution quality. CCARI-EX is close to supported optima for response-heavy $n = 8$ epochs but can be more than 20% away in individual cost-heavy cases. The maximum balanced gap of 41.0% at $n = 6$ is especially important: good average online outcomes do not imply consistently near-optimal local construction. More powerful large-neighborhood search, route-pool reuse, or a rolling restricted master could target those cases. Exact branch-price-and-cut work suggests how route generation could scale beyond exhaustive enumeration [27], [28].

The dominated $w = 0.25$ setting adds a second caution. In static bicriteria optimization, reducing the cost weight usually moves toward faster solutions. In a dynamic system, immediate dispatch can occupy the wrong vehicle when a burst arrives. A myopic response weight may worsen later response. A future policy should couple weight selection to an estimate of near-term arrival pressure or learn a state-dependent batching decision, while retaining hard feasibility checks. Safe reinforcement learning is relevant here, but any learned policy would need out-of-distribution and constraint validation [29].

Operationally, the model demonstrates why charging cannot be treated as a reset. A returned vehicle's charging duration determines whether it is available for the next release cluster. Grid energy exceeds battery energy by the specified efficiency loss and is the quantity priced. For the holdout, CCARI-EX (0.75) averaged 541.38 kWh from the grid. Changing electricity tariff, charge efficiency, or charger capacity could alter both class selection and response. Time-of-use prices and nonlinear charging would be natural extensions [30], [31].

The study also clarifies the scope of "autonomous." Auxiliary power and supervision enter the operating model, but no autonomy-specific operational dataset was found that jointly identifies computing draw, remote interventions, freight payload, and route context. We therefore avoid claiming that the selected power and supervision values are measured. An AEV deployment study should replace them with platform logs and include fallback events, geofenced operational design domains, and safety-driven service interruptions.

9. Limitations

First, the service graph uses scaled Euclidean distances from one Homberger base instance. Although the dynamic samples span 1,000 spatial points, they do not reproduce turn restrictions, grades, asymmetric links, or a real city topology. The SUMO grid is a separate transfer layer, not a one-to-one replay of the routing streams.

Second, Darwin speed quantiles are trip-level summaries applied as global 60-min regimes. They capture smooth-to-congested variation but not spatial congestion. This explains why the traffic-awareness ablation is weak and why the SUMO congested regime produces censoring that a speed step cannot express. Link-level historical and live data should replace the global field.

Third, vehicle classes combine measured, simulated, and assumed parameters. The small payload and motor, the medium/large capacity assignments anchored to reference-load cases, all charging powers, autonomy auxiliary loads, maintenance costs, fixed costs, electricity tariff, supervision rate, reserve, and target SOC are scenarios. JRC heavy-duty outputs are simulations generated by VECTO rather than road measurements, and motor rating is not an independently active model term. Sensitivity analysis in the workbook enables re-evaluation, but conclusions are conditional on these inputs.

Fourth, the exact oracle allows one route per available vehicle and at most eight requests. It validates an epoch constructor, not repeated-trip online optimality. The dynamic policy does not reroute vehicles already underway and has no public charging, depot-plug capacity, stochastic energy, battery degradation, or charger queue. Literature on capacitated and nonlinear charging shows that these decisions can matter [32], [33].

Fifth, 30 seeds per design cell support paired comparisons but do not establish performance on other cities, demand distributions, fleet mixes, tariffs, or weather. The balanced index uses the best observed method in each stream and is descriptive; raw cost and response remain primary. Wilcoxon inference assumes independent streams, and bootstrap intervals quantify sampling variability under the generator, not structural model uncertainty.

Finally, no claim of journal acceptance or deployment readiness follows from the reported results. The code and processed data make the computational claims auditable. Authors and reviewers must still evaluate application validity, authorship, institutional declarations, and any journal-specific reporting requirements.

10. Conclusion

This study presents a reproducible bi-objective network model for dynamic routing of heterogeneous autonomous electric vehicles without customer time windows. It replaces departure-state travel with FIFO-integrated traversal, carries SOC across trips, prices grid-side energy, and blocks vehicles during charging. Public sources calibrate traffic and class energy, while unsupported fleet values are labeled as scenarios.

CCARI-EX combines causal batching, charge-aware heterogeneous insertion, relocate search, and cross-class exchange. On 480 held-out streams and 4,800 executions, its cost-priority setting served every request, incurred zero independently validated dispatch rejections, reduced both cost and mean response relative to cost-greedy and nearest insertion, and was nondominated in 69.0% of streams. Exact route-set models define what this result means: the heuristic is competitive but not exact, with gaps that vary materially by objective weight. Ablations identify batching as the strongest supported component, while traffic and fit penalties show no significant average gain. SUMO preserves regime ordering but exposes congested noncompletion beyond an aggregate speed abstraction.

The publishable contribution is therefore an aligned modeling-and-validation framework plus a competitive online policy, not a universal optimum. The next scientific step is link-specific stochastic traffic, capacitated and nonlinear charging, real autonomous-freight logs, and a rolling route-pool master evaluated on multiple real networks.

11. Declarations

11.1. Author Contributions

Kashif Bashir: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing – Original Draft; **Rumaisa Rafiq:** Investigation, Data Curation, Writing – Review & Editing; **Muhammad Afzal Shah:** Methodology, Formal analysis, Visualization, Writing – Review & Editing; **Marzia Niazi:** Literature analysis, Validation, Writing – Review & Editing; **Nouman Bashir:** Data analysis, Investigation, Writing – Review & Editing; **Sara Khan:** Formal analysis, Writing – Review & Editing; **Haseeb ur Rehman:** Software, Validation, Visualization.

11.2. Institutional Review Board Statement

Not applicable.

11.3. Informed Consent Statement

Not applicable.

11.4. Data Availability Statement

The data used in this study is available on request from the corresponding author.

11.5. Acknowledgment

Not applicable.

11.6. Conflicts of Interest

The authors declare no conflict of interest.

11.7. Generative AI Use Disclosure

During the preparation of this manuscript, the authors used generative AI tools only for language improvement and proofreading assistance. The authors reviewed and edited all content and take full responsibility for the final published manuscript.

12. References

[1] C. Archetti, L. C. Coelho, M. G. Speranza, and P. Vansteenwegen, "Beyond fifty years of vehicle routing: Insights into the history and the future," *European Journal of Operational Research*, vol. 330, no. 2, pp. 355–372, 2026. https://doi.org/10.1016/j.ejor.2025.06.014.

[2] K. Rajashekara, "Energy impacts of autonomous vehicles—Present and the future," *iEnergy*, vol. 3, no. 2, pp. 73–74, 2024, https://doi.org/10.23919/IEN.2024.0012.

[3] G. Lera-Romero, J. J. Miranda Bront, and F. J. Soulignac, "A branch-cut-and-price algorithm for the time-dependent electric vehicle routing problem with time windows," *European Journal of Operational Research*, vol. 312, no. 3, pp. 978–995, 2024. https://doi.org/10.1016/j.ejor.2023.06.037.

[4] S. Zhou, D. Zhang, B. Ji, S. Zhou, S. Li, and L. Zhou, "A MILP model and heuristic method for the time-dependent electric vehicle routing and scheduling problem with time windows," *Journal of Cleaner Production*, vol. 434, Art. no. 140188, 2024, https://doi.org/10.1016/j.jclepro.2023.140188.

[5] N. Soeffker, M. W. Ulmer, and D. C. Mattfeld, "Stochastic dynamic vehicle routing in the light of prescriptive analytics: A review," *European Journal of Operational Research*, vol. 298, no. 3, pp. 801–820, 2022. https://doi.org/10.1016/j.ejor.2021.07.014.

[6] J. Zhang, K. Luo, A. M. Florio, and T. Van Woensel, "Solving large-scale dynamic vehicle routing problems with stochastic requests," *European Journal of Operational Research*, vol. 306, no. 2, pp. 596–614, 2023. https://doi.org/10.1016/j.ejor.2022.07.015.

[7] W. Wu, Y. Lin, R. Liu, and W. Jin, "The multi-depot electric vehicle scheduling problem with power grid characteristics," *Transportation Research Part B: Methodological*, vol. 155, pp. 322–347. https://doi.org/10.1016/j.trb.2021.11.007.

[8] M. H. de Vos, R. N. van Lieshout, and T. Dollevoet, "Electric vehicle scheduling in public transit with capacitated charging stations," *Transportation Science*, vol. 58, no. 2, pp. 279–294, 2024. https://doi.org/10.1287/trsc.2022.0253.

[9] A. Tansini, J. Suarez Corujo, A. Laverde Marin, M. Otura Garcia, and G. Fontaras, Real-world driving data—BEV (Darwin) [Dataset]. *European Commission, Joint Research Centre*, 2025. https://doi.org/10.2905/JRC.A754JZ8.

[10] A. Hanif, M. Abid, M. Tabaa, H. Bensag, and M. Youssfi, "Electric vehicle routing with time windows and heterogeneous charging-station attribute dataset," *Data*, vol. 11, no. 4, Art. no. 83, 2026. https://doi.org/10.3390/data11040083.

[11] O. Rainio, J. Teuho, and R. Klén, "Evaluation metrics and statistical tests for machine learning," *Scientific Reports*, vol. 14, Art. no. 6086, 2024. https://doi.org/10.1038/s41598-024-56706-x.

[12] İ. Küçükoğlu, R. Dewil, and D. Cattrysse, "The electric vehicle routing problem and its variations: A literature review," *Computers & Industrial Engineering*, vol. 161, Art. no. 107650, 2021. https://doi.org/10.1016/j.cie.2021.107650.

[13] Y. Xiao, Y. Zhang, I. Kaku, R. Kang, and X. Pan, "Electric vehicle routing problem: A systematic review and a new comprehensive model with nonlinear energy recharging and consumption," *Renewable and Sustainable Energy Reviews*, vol. 151, Art. no. 111567, 2021, doi: 10.1016/j.rser.2021.111567. https://doi.org/10.1016/j.rser.2021.111567.

[14] E. Lam, G. Desaulniers, and P. J. Stuckey, "Branch-and-cut-and-price for the electric vehicle routing problem with time windows, piecewise-linear recharging and capacitated recharging stations," *Computers & Operations Research*, vol. 145, Art. no. 105870, 2022. https://doi.org/10.1016/j.cor.2022.105870.

[15] R. Raeesi and K. G. Zografos, "Coordinated routing of electric commercial vehicles with intra-route recharging and en-route battery swapping," *European Journal of Operational Research*, vol. 301, no. 1, pp. 82–109, 2022. https://doi.org/10.1016/j.ejor.2021.09.037.

[16] M. Seyfi, M. Alinaghian, E. Ghorbani, B. Çatay, and M. S. Sabbagh, "Multi-mode hybrid electric vehicle routing problem," *Transportation Research Part E: Logistics and Transportation Review*, vol. 166, Art. no. 102882, 2022. https://doi.org/10.1016/j.tre.2022.102882.

[17] S. Dönmez, Ç. Koç, and F. Altıparmak, "The mixed fleet vehicle routing problem with partial recharging by multiple chargers: Mathematical model and adaptive large neighborhood search," *Transportation Research Part E: Logistics and Transportation Review*, vol. 167, Art. no. 102917, 2022. https://doi.org/10.1016/j.tre.2022.102917.

[18] W. Wang, Y. Adulyasak, and J.-F. Cordeau, "Electric vehicle routing with heterogeneous charging stations," *Computers & Operations Research*, vol. 189, Art. no. 107374, 2026. https://doi.org/10.1016/j.cor.2025.107374.

[19] M. Bruglieri, M. Paolucci, and O. Pisacane, "A matheuristic for the electric vehicle routing problem with time windows and a realistic energy consumption model," *Computers & Operations Research*, vol. 157, Art. no. 106261, 2023. https://doi.org/10.1016/j.cor.2023.106261.

[20] J. Dong, H. Wang, and S. Zhang, "Dynamic electric vehicle routing problem considering mid-route recharging and new demand arrival using an improved memetic algorithm," *Sustainable Energy Technologies and Assessments*, vol. 58, Art. no. 103366, 2023. https://doi.org/10.1016/j.seta.2023.103366.

[21] M. Bruglieri, R. Peruzzini, and O. Pisacane, "Efficiently routing a fleet of autonomous vehicles in a real-time ride-sharing system," *Computers & Operations Research*, vol. 168, Art. no. 106668, 2024. https://doi.org/10.1016/j.cor.2024.106668.

[22] J. Zhao, Z. Zeng, and Y. Liu, "Electric vehicle routing problem considering traffic conditions and real-time loads," *Transportation Research Part C: Emerging Technologies*, vol. 176, Art. no. 105150, 2025. https://doi.org/10.1016/j.trc.2025.105150.

[23] H. Gehring and J. Homberger, "Two evolutionary metaheuristics for the vehicle routing problem with time windows," *INFOR: Information Systems and Operational Research*, vol. 37, no. 3, pp. 297–318, 1999. https://doi.org/10.1080/03155986.1999.11732386.

[24] G. Oh, D. J. LeBlanc, and H. Peng, "Vehicle Energy Dataset (VED), a large-scale dataset for vehicle energy consumption research," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 4, pp. 3302–3312, 2022. https://doi.org/10.1109/TITS.2020.3035596.

[25] W. Wang and J. Zhao, "Partial linear recharging strategy for the electric fleet size and mix vehicle routing problem with time windows and recharging stations," *European Journal of Operational Research*, vol. 308, no. 2, pp. 929–948, 2023. https://doi.org/10.1016/j.ejor.2022.12.011.

[26] Y. Jiang, W. Hu, W. Gu, Y. Yu, and M. Xu, "A multi-mode hybrid electric vehicle routing problem with time windows," *Transportation Research Part E: Logistics and Transportation Review*, vol. 195, Art. no. 103976, 2025. https://doi.org/10.1016/j.tre.2025.103976.

[27] A. Parmentier, R. Martinelli, and T. Vidal, "Electric vehicle fleets: Scalable route and recharge scheduling through column generation," *Transportation Science*, vol. 57, no. 3, pp. 631–646, 2023. https://doi.org/10.1287/trsc.2023.1199.

[28] Y. Zang, M. Wang, and M. Qi, "A column generation tailored to electric vehicle routing problem with nonlinear battery depreciation," *Computers & Operations Research*, vol. 137, Art. no. 105527, 2022. https://doi.org/10.1016/j.cor.2021.105527.

[29] R. Basso, B. Kulcsár, I. Sánchez-Díaz, and X. Qu, "Dynamic stochastic electric vehicle routing with safe reinforcement learning," *Transportation Research Part E: Logistics and Transportation Review*, vol. 157, Art. no. 102496, 2022. https://doi.org/10.1016/j.tre.2021.102496.

[30] B. Lin, B. Ghaddar, and J. Nathwani, "Electric vehicle routing with charging/discharging under time-variant electricity prices," *Transportation Research Part C: Emerging Technologies*, vol. 130, Art. no. 103285, 2021. https://doi.org/10.1016/j.trc.2021.103285.

[31] A. Froger, O. Jabali, J. E. Mendoza, and G. Laporte, "The electric vehicle routing problem with capacitated charging stations," *Transportation Science*, vol. 56, no. 2, pp. 460–482, 2022. https://doi.org/10.1287/trsc.2021.1111.

[32] M. Dastpak, F. Errico, O. Jabali, and F. Malucelli, "Dynamic routing for the electric vehicle shortest path problem with charging station occupancy information," *Transportation Research Part C: Emerging Technologies*, vol. 158, Art. no. 104411, 2024. https://doi.org/10.1016/j.trc.2023.104411.

[33] G. M. Nafstad, G. Desaulniers, and M. Stålhane, "Branch-price-and-cut for the electric vehicle routing problem with heterogeneous recharging technologies and nonlinear recharging functions," *Transportation Science*, vol. 59, no. 3, pp. 628–646, 2025. https://doi.org/10.1287/trsc.2024.0725.