Fajardo-Calderin et al. European Transport Research Review (2026) 18:22
https://doi.org/10.1186/s12544-026-00782-9

European Transport Research Review

ORIGINAL PAPER

Intelligent vehicle routing problem solver for simulating urban last-mile logistics

Jenny Fajardo-Calderin[1*], Antonio D. Masegosa[1,2], Xabier Cantero[1], Asier Moreno[1] and Pilar Elejoste[1]

Abstract
The exponential growth of e-commerce and increasing urbanization have significantly complicated last-mile logistics, making more efficient and sustainable delivery solutions necessary. This study presents a solver for simulating urban last-mile logistics, developed within the framework of the European SENATOR project, which aims to optimize urban logistics. The proposed solution integrates key concepts such as multimodal transportation, time-dependent travel conditions, and driving constraints to enhance delivery efficiency and minimize environmental impact. To evaluate its effectiveness, a pilot study in Zaragoza, Spain, was conducted, simulating three experimental scenarios: 1) Baseline Scenario – Reflecting the current postal delivery operations, 2) Zero Emission Zone (ZEZ) – Assessing delivery performance when restricting access to fuel-powered vehicles in the city center, 3) Urban Consolidation Centre (UCC) – Evaluating the impact of centralized parcel distribution. Each scenario was analyzed under three fleet electrification levels (current composition, 50% electrified fleet, and 100% electrified fleet). The results demonstrate how the solver can simulate and quantify the operational impact of different policies on an urban logistics distributor, such as the implementation of zero-emission zones (leading to fewer packages delivered and an increase in total distance traveled). Furthermore, it shows how this impact varies with different levels of fleet electrification (higher electrification results in a lower negative impact on operations). Furthermore, the developed tool highlights how an urban consolidation center can help reduce the distance traveled by logistics vehicles and, consequently, their emissions.

Keywords Collaborative delivery, Dynamic planning, Last-mile logistics, Multi-modal, Optimization, Vehicle routing problem

1 Introduction

In recent years, last-mile logistics has gained significant attention due to the exponential growth of e-commerce, particularly during the COVID-19 pandemic. This surge in demand for urban order deliveries has substantially increased the complexity of urban logistics operations. Consequently, the efficient design and management of last-mile logistics are now considered crucial for enhancing societal welfare. The term “last mile” refers to the final segment of the delivery process, which concludes at the recipient’s predefined destination. Various terminologies are used to describe this process, including last-mile logistics, last-mile delivery, and last-mile parcel distribution, among others.

To mitigate the environmental externalities associated with last-mile logistics, companies must strike a balance between fast delivery times, adherence to customers’ time windows, and environmental sustainability. In this context, green transportation has become a

*Correspondence:
Jenny Fajardo-Calderin
fajardo.jenny@deusto.es
[1]Faculty of Engineering, University of Deusto, Avda. de las Universidades 24, Bilbao, Basque Country 48007, Spain
[2]Ikerbasque, Basque Foundation for Science, Bilbao, Basque Country 48009, Spain

widely adopted alternative among logistics operators [1]. Given that most last-mile deliveries occur in urban areas, the promotion of environmentally friendly transport modes—such as bicycles, scooters, electric cars, motorcycles, and vans—has gained significant importance. Additionally, there is a growing trend toward collaborative last-mile logistics models, where multiple logistics service providers manage deliveries through a shared platform. This collaborative approach enhances resource utilization and maximizes operational efficiency [2, 3].

This study is developed within the framework of SENATOR[1], an European research project focused on optimizing collaborative urban logistics by implementing innovative planning and routing strategies. As part of the project, a solver was developed to simulate real-world last-mile logistics environments. This solver enables the analysis and optimization of delivery operations while accounting for variables such as operational constraints, fleet electrification, and environmental impact. To validate the effectiveness of this approach, a pilot study was conducted in the city of Zaragoza, Spain, where different logistics scenarios of the logistics company were simulated, to evaluate the impact of various urban logistics strategies.

Our contribution focuses on (i) a unified input model for pickup/delivery jobs, vehicles and vehicle types, and urban access rules; (ii) time-dependent routing by replacing static travel-time assumptions with time-sliced distance/time profiles calibrated from historical traffic, enabling more realistic shift planning; (iii) operational robustness through explicit, configurable per-shipment service times (iv) heterogeneous fleets and access restrictions (e.g., low/zero-emission zones) embedded into the feasibility of routes; and (v) a configurable multi-objective optimization that extends the classic distance/time criterion to include operational costs and per-vehicle emissions via a weighted-sum aggregation. Importantly, our solver addresses these elements in an offline planning setting: uncertainty is handled through data-driven profiles and configurable parameters rather than an explicit stochastic vehicle routing problem (SVRP) formulation, and real-time re-optimization for dynamic request arrivals (DVRP) is left for future work. Likewise, "green" aspects are captured through fleet composition and emission factors; modeling energy-routing decisions such as battery/range constraints, charging operations, or speed choices is outside the present scope.

With these challenges in mind, the SENATOR project aims to create a new urban logistic model for enhancing the sustainability of cities. For this purpose, the project will develop a smart network operator, as a control tower supported on an ICT Platform that will work as a support tool for decision-making, integration and planning of all logistics operations.

To evaluate the proposed solution, three experimental scenarios were designed, each addressing a specific last-mile logistics challenge:

- Scenario 1 (Baseline Scenario): Represents the current postal operations in Zaragoza, with nine Delivery Units (DUs) and two Special Service Units (SSUs) operating in two daily shifts. This scenario serves as a benchmark for comparing the effects of different logistics strategies.
- Scenario 2 (Zero Emission Zone - ZEZ): Simulates the implementation of a Zero Emission Zone in the historic city center, restricting access to combustion vehicles during daytime hours. This scenario assesses the impact of fleet electrification on delivery efficiency and emissions reduction.
- Scenario 3 (Urban Consolidation Centre - UCC): Examines the effect of establishing an Urban Consolidation Centre at the San Vicente de Paúl Market, centralizing parcel distribution to reduce vehicle congestion and optimize last-mile logistics.

Each scenario was analyzed under three fleet compositions: current fleet composition, 50% electrification, and 100% electrification, allowing for a comprehensive assessment of the environmental and operational impact of different electrification strategies. This research contributes to the field of sustainable last-mile logistics by proposing an optimized route planning model that integrates environmental, operational, and economic considerations.

The remainder of this paper is structured as follows: Sect. 2 reviews the most relevant related work in the field. Section 3 presents the integrated modeling and solver architecture: 3.1 describes the architecture and components; 3.2 details intelligent route management (robustness to stochastic travel/service times, driving/access restrictions, time dependence, and portability to different urban contexts); 3.3 introduces the Large Neighborhood Search algorithm and the destroy/repair operators; and 3.4 summarizes implementation details. Section 4 describes the experimental set-up and the three scenarios: a baseline case, the deployment of a low emission zone, and the implementation of an urban consolidation center. Section 7 reports the experiments and results obtained under three different scenarios. Finally, Sect. 6 provides the main conclusions drawn from this work.

2 Related work

Among routing optimization problems, one of the most extensively studied in the fields of computer science and operations research is the Travelling Salesman Problem (TSP). The TSP, along with the Vehicle Routing Problem

[1] https://www.senatorproject.eu/

(VRP), represents a cornerstone in combinatorial optimization, given its complexity and wide range of real-world applications [4]. The VRP aims to determine an optimal set of routes that minimizes the total cost while satisfying the following constraints: (i) each route must start and end at a designated depot, (ii) each customer must be visited exactly once, and (iii) the total demand of customers assigned to a route must not exceed the vehicle's capacity [5, 6]. These problems are classified as NP-hard, meaning that obtaining an optimal solution becomes computationally infeasible as the problem size increases [7].

In real-world applications, uncertainty is an inherent characteristic of routing problems. This uncertainty may arise due to stochastic variations in demand, travel times, unexpected disruptions (e.g., vehicle breakdowns), or other dynamic factors. The Stochastic Vehicle Routing Problem (SVRP) [8] is a variant of the VRP in which one or more problem parameters are represented as random variables following known probability distributions. Beyond uncertainty, many real-world logistics problems exhibit dynamic elements that evolve over time. The Dynamic Vehicle Routing Problem (DVRP) [9], also known as real-time or online VRP, involves scenarios where input data is partially disclosed or subject to modification during the execution of distribution operations. Common dynamic events include the arrival of new customer pick-up/delivery requests or fluctuations in service and travel times.

Furthermore, solving VRP in real-world logistics, particularly in last-mile delivery scenarios, requires more than just optimizing travel distance or operational time. Multiple objectives must be considered, such as cost minimization, service level improvements, and environmental impact reductions (e.g., emissions) [10]. At the core of last-mile logistics challenges lies the VRP. However, real-world last-mile logistics introduce additional complexities, requiring the consideration of multiple constraints and objectives.

Several VRP variants are particularly relevant to last-mile delivery operations:

- Capacitated VRP (CVRP) – Considers vehicle capacity limitations when designing routes.
- VRP with Time Windows (VRPTW) – Ensures that deliveries occur within predefined time slots.
- Stochastic VRP (SVRP) – Incorporates uncertainties such as fluctuating demand and variable travel times.
- Multi-Objective VRP (MOVRP) – Simultaneously optimizes conflicting objectives, such as cost minimization, service quality, and environmental impact.
- Green VRP (GVRP) – Focuses on optimizing routes while considering fuel consumption and carbon emissions.

Given the complex and dynamic nature of last-mile logistics, solving VRP efficiently requires the use of advanced optimization techniques, including metaheuristics (e.g., Genetic Algorithms, Tabu Search, and Ant Colony Optimization) and hybrid approaches that integrate artificial intelligence and machine learning [11].

An influential contribution within ruin-and-recreate/Large Neighborhood Search (LNS) is the Slack Induction by String Removals (SISR) operator introduced by Christiaens and Vanden Berghe [12]. The key idea is to remove contiguous sequences of customers ("strings") using time-window and load slack measures to induce slack in incumbent routes, then reconstruct them with informed insertion strategies (e.g., regret- or greedy-based) under explicit constraint handling. By enabling large-scale, structure-aware moves that are hard to reach with classical local search, SISR has demonstrated competitive performance on VRPTW and related variants. This principle—destroying in a guided way to create slack that makes better reconstructions feasible—directly motivates our choice of a ruin-and-recreate scheme and the design of operators tailored to urban access restrictions, time dependence, and electric-vehicle constraints.

One of the most recently studied objectives in the literature, and one that has taken special attention in last-mile logistics, is the reduction of emissions. Freight transport in cities can be responsible for up to 40% of direct carbon emissions. Two variants of the VRP problem exist to address this objective: the Green VRP (GVRP) [13], which considers carbon emissions in routing, and the Pollution Routing Problem (PRP), which considers fuel consumption [14].

A wide variety of tools have been developed to simulate and analyze different scenarios in last-mile logistics, leveraging diverse optimization strategies ranging from exact methods to heuristic, metaheuristic, and hybrid techniques. Additionally, numerous computational tools exist for modeling and solving logistics problems, some of which share conceptual similarities with the solver presented in this work. In this section, we review a selection of these existing solutions and provide a comparative analysis to position our solver within the current state of the art.

- JSprit[2] is an open-source, Java-based vehicle route optimization engine that employs a Ruin & Recreate metaheuristic. It is highly flexible and supports multiple VRP variants, including capacitated VRP, multi-depot VRP, time-dependent VRP, and heterogeneous fleet VRP. Its modular design allows for easy modification and extension, making it suitable for customization. The tool enables the

[2] https://github.com/graphhopper/jsprit/tree/master/docs

definition of custom constraints, optimization configurations, and solution visualization. JSprit employs Java Reflection for constraint management, dividing them into hard and soft constraints related to routes and activities. Its adaptability and robust testing framework make it a reliable option for VRP optimization.

- OR-Tools[3] developed by Google is an open-source optimization suite designed to tackle complex routing, flow, and constraint programming problems. It is implemented in C++ with support for Python, C#, and Java wrappers. The tool includes a specialized routing library capable of solving various VRP variants, such as VRPTW, VRPPD, and CVRP. OR-Tools integrates multiple solving strategies, including Greedy Descent, Guided Local Search, Simulated Annealing, and Tabu Search. Its high computational efficiency and configurable parameters make it a widely used tool for solving large-scale VRP instances.
- VROOM[4] is an open-source, C++-based optimization tool designed to address VRP-related problems in logistics and geographically distributed task management. It integrates with open-source routing engines such as OSRM, OpenRouteService, and Valhalla. VROOM utilizes heuristics and metaheuristics to provide high-quality solutions efficiently, making it suitable for large-scale problem instances. It supports various VRP types, including TSP, CVRP, VRPTW, and multi-depot heterogeneous fleet VRP. Its solving approach emphasizes speed and scalability, allowing it to handle dynamic and complex routing problems effectively.
- ArcGIS VRP[5] is a commercial routing service developed by ESRI that provides a customizable solution for solving a broad range of VRP variants. It supports multiple depots, heterogeneous fleets, time windows, and pickup-and-delivery constraints. A key advantage of ArcGIS VRP is its integration with ESRI's geographic information system (GIS), allowing users to incorporate geospatial factors such as restricted areas and varying speed zones. However, it is a pay-per-use service, limiting accessibility for large-scale applications. It can be accessed in different ways such as JavaScript Application Programming Interface (API) and SDKs in different programming languages.
- Circuit[6] is a commercial route optimization tool available as a web service and mobile application (Android and iOS). It optimizes delivery routes by considering real-time traffic conditions and the latest map data. The tool supports up to 1000 stops per route, time windows, and priority-based stop ordering. Circuit is commonly used for driver tracking, courier management, and local delivery planning. It integrates with third-party navigation applications such as Google Maps but has limited configurability and extension options due to licensing constraints.
- LOCUS[7] is a route planning and vehicle assignment tool designed for logistics operations, particularly last-mile delivery. It supports various optimization problems, including VRP with capacity constraints, time windows, and pickups/deliveries. LOCUS employs exact, heuristic, and hybrid algorithms, though the specific methods are not publicly disclosed. The platform provides real-time tracking, dynamic route planning, and API integration for seamless connectivity with external systems. However, it is a commercial tool, limiting accessibility for potential users.
- OptaPlanner[8] is an open-source Java-based optimization tool under the Apache License, designed for solving complex combinatorial problems, including VRP and its variants. It implements multiple metaheuristics such as Tabu Search, Simulated Annealing, and Late Acceptance. The tool supports custom heuristic construction and parallelization for performance improvement. OptaPlanner allows integration with mapping services like Google Maps and OpenStreetMap, making it a versatile choice for route optimization. It also provides benchmarking tools for performance comparison, making it valuable for research and industry applications.
- HERE[9] is a commercial tool that provides a route planning API to solve the VRP, implementing the Capacitated VRP, VRP with Time Windows, Multi-Depot VRP, Open Vehicle Routing, Heterogeneous or Mixed Fleet VRP, and Pickup and Delivery VRP. It allows the calculation of routes using real-time and historical traffic information and the re-planning of routes in real time if new orders appear.
- GraphHopper[10] is an open-source software tool developed in Java that uses JSprit as the route

[3] https://github.com/google/or-tools

[4] https://github.com/VROOM-Project/vroom

[5] https://desktop.arcgis.com/es/arcmap/latest/extensions/network-analyst/vehicle-routing-problem.htm

[6] https://getcircuit.com/

[7] https://locus.sh/

[8] https://www.optaplanner.org/

[9] https://developer.here.com/products/tour-planning

[10] https://github.com/graphhopper/graphhopper

optimization engine. It provides an API to solve a variety of vehicle routing problems, including the Traveling Salesman optimization problem, and all the VRP variants implemented in JSprit. Its main advantages are the possibility of designing vehicle types and defining time windows and service times for drivers.

Despite the wide range of functionalities offered by the tools analyzed, none of them simultaneously integrate key features essential for comprehensive last-mile logistics optimization, as can be seen in Table 1. Specifically, none of these solutions provide a holistic approach that combines multi-modal fleet optimization, inter-modal and transfer route optimization, multi-objective optimization, and the consideration of driving constraints within a single framework. While some tools excel in specific areas—such as dynamic routing, vehicle capacity management, or GIS integration—none effectively address all these challenges in an unified manner. This limitation highlights the need for an advanced, flexible, and adaptive solution capable of handling the complex and evolving demands of modern urban logistics.

Positioning with respect to JSprit. Our solver uses JSprit as the underlying metaheuristic engine (ruin-and-recreate / LNS), and extends it towards realistic last-mile planning in three main ways. First, we provide a data-driven integration layer (JSON schema + REST-oriented workflow) that translates operational inputs (pickup/delivery jobs, heterogeneous fleets, and access rules) into the JSprit problem model. Second, we enhance realism by operationalizing time dependence and robustness: we support multi-period travel-time/distance matrices and adapt the cost interfaces and time-propagation routines so that evaluation and insertion use the appropriate time slice, while allowing per-shipment service times (with configurable defaults). Third, we extend the optimization criteria beyond JSprit's default distance/time objective by incorporating operational cost terms and per-vehicle emissions through a configurable weighted-sum multi-objective function. These additions keep compatibility with JSprit's search operators while enabling scenario-based policy analysis (e.g., ZEZ/UCC) in an offline planning setting.

Furthermore, it is important to highlight that, despite the extensive development of optimization techniques and the proliferation of variants within the field of operations research, many of these approaches have yet to be applied to the specific context of urban last-mile logistics. Techniques such as advanced metaheuristics, hybrid models integrating machine learning with combinatorial optimization, and decomposition methods designed for large-scale problems remain largely unexplored in addressing the unique constraints and dynamic conditions of urban delivery environments. This underutilization of existing optimization strategies represents a significant opportunity for future research aimed at enhancing the efficiency, adaptability, and sustainability of last-mile logistics through innovative algorithmic solutions.

3 Integrated modeling and solver architecture

This section describes the overall system and methodology: we first outline the solver architecture, then detail the operational data and uncertainty models (driving restrictions and time-dependent travel), and finally present the optimization algorithm and its implementation.

3.1 Architecture and components

The core of the architecture is the optimization engine, which comprises several sub-modules, each dedicated to a specific function essential for the model's efficiency and adaptability. These sub-modules include:

- Data Processing: This module handles the preprocessing and transformation of all input data

Table 1 Feature comparison of VRP tools and platforms according to the requested criteria

Tool / Platform	Open-source	Multimodal fleets	Multi-objective	Dynamic model	EV/Energy Emissions	Driving restrictions	Time-dependent
OR-Tools	Yes	Partial	Partial	Partial	Partial	Partial	Partial
GraphHopper	Yes	Partial	Partial	Partial	Partial	Partial	Yes
VROOM	Yes	Partial	Partial	Partial	No	Partial	Partial
OptaPlanner	Yes	Partial	Yes	Yes	Partial	Partial	Partial
ArcGIS VRP	No	Partial	Partial	Partial	Partial	Yes	Yes
HERE	No	Partial	Partial	Yes	Partial	Yes	Yes
Circuit	No	No	Partial	Yes	No	No	Yes
LOCUS	No	Partial	Partial	Yes	Partial	Yes	Yes
JSprit	Yes	Partial	Partial	Partial	Partial	Partial	Yes
Solver (JSprit-based)	Yes	Yes	Yes	Partial	Partial	Yes	Yes

Notes: Yes = native support; Partial = indirect/limited support via custom modeling, callbacks, or external integration; No = not natively supported into the data structures required by the optimization algorithm. It ensures that the data is correctly formatted and validated for seamless integration into the optimization process.

- Living Labs Services: This module manages the diverse constraints and intricacies associated with the optimization model. It ensures that the model can be adapted to different use cases by incorporating specific operational and regulatory constraints relevant to each scenario.
- Output Solution Processing: This module is responsible for generating the final output of the optimization process. It computes and organizes key performance indicators (KPIs) to evaluate the quality of the solution, enabling a comprehensive assessment of its effectiveness in meeting the desired objectives.

The general solver that encompasses the optimization engine is primarily responsible for four key tasks: (1) integration with the Representational State Transfer (REST) API and related services, (2) execution of the optimization engine that loads all input data and formulates the problem instance, (3) integration with the JSprit framework, and (4) processing the algorithm's output to generate the solution, which is then provided through the output REST API. In Fig. 1, the different components and their relationships are depicted. Certain components serve specific purposes and have been designed to strengthen the optimization model.

Additionally, a dedicated JSprit module has been developed with the following functionalities: (1) it implements the Vehicle Routing Problem (VRP) framework, which allows for adapting a dynamic model to delivery planning in the future; and (2) the main modifications focus on constraint modeling, operator strategies, and fitness function computation to align with the dynamic nature of the problem.

A specialized optimization subsystem has also been developed to meet the specific requirements of each use case defined in the proposed model. This subsystem is responsible for processing input data and generating optimized solutions, which are stored in a database that can be accessed through REST API. The planning model enables real-time last-mile route planning for heterogeneous fleets of vehicles, optimizing resource utilization, and ensuring efficient delivery operations.

3.2 Intelligent route management

The developed solver is designed to simulate realistic last-mile logistics scenarios by addressing complex vehicle routing problems that capture key aspects of real-world operations. Built upon the well-established Rich Vehicle Routing Problem (RVRP) [5], the solver enables the analysis of diverse operational conditions, providing a flexible framework to evaluate and optimize delivery strategies under varying constraints and requirements.

The primary objective was to achieve more realistic optimization functions by incorporating, along with a wide range of real-world constraints related to time, distance, and the use of heterogeneous fleets. The integration of these advanced route planning concepts is essential for developing intelligent route management solutions capable of meeting the evolving demands of urban logistics. This section defines and analyzes these concepts as they have been integrated into our proposed solver.

3.2.1 Robustness

Uncertainty is a fundamental characteristic of real-world problems, and in vehicle routing, certain parameters may be unknown or subject to variation. In contrast to deterministic optimization, where all data is known in advance, stochastic optimization is used to handle uncertainty. Stochastic models rely on historical data to build

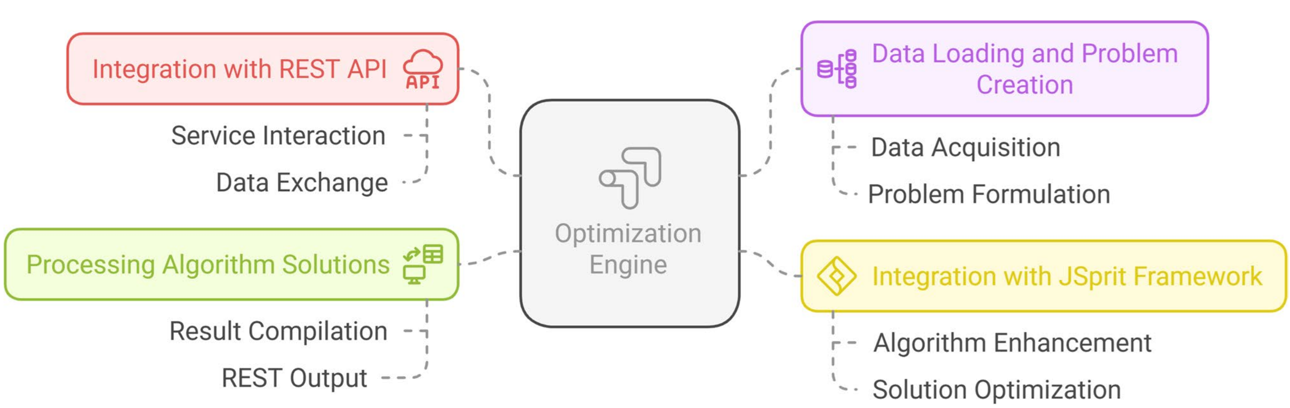

Fig. 1 Architecture of the proposed solver. The optimization engine sits at the core and exchanges information with four modules: (i) data loading & problem creation (data acquisition and formulation), (ii) integration with JSprit (algorithm enhancements and route optimization), (iii) integration with REST API (service interaction and data exchange), and (iv) solution processing (result compilation and REST outputs). Dashed connectors indicate data/control interfaces

probabilistic distributions, enabling the generation of more robust solutions that can withstand variations in travel times, service durations, or demand [15].

The Stochastic Vehicle Routing Problem (SVRP) addresses uncertainty through several variants, such as VRPSD (stochastic demand), VRPSDC (stochastic demand and customers), and VRPSTS (stochastic travel times and service). By considering these stochastic components, the robustness of the solutions can be improved, especially in terms of adapting to travel time and service duration fluctuations [16, 17].

To improve solution quality and robustness, the input model allows service times to be specified per shipment (for both pickup and delivery activities). If a shipment specific value is not provided, a configurable default is applied at data loading time; this default can be adjusted through the input data settings. We modified existing and introduced new in JSprit data structures, to accommodate variable service times.

The solver proposed in this work does not implement an explicit SVRP formulation (e.g., two-stage stochastic programming with recourse decisions, chance constraints, or sampling-based optimization embedded in the search). Instead, we solve a deterministic VRP instance parameterized with data driven, travel time profiles and configurable service time assumptions. In this sense, our approach aligns with a common operational practice: uncertainty is handled by using expected (or otherwise conservative) values and by testing the resulting plan under alternative assumptions.

3.2.2 Driving restrictions

Operational constraints are a critical component of vehicle routing models, especially in urban environments, where factors such as vehicle capacity, driving regulations, time windows, and access restrictions play a significant role. In many cities, specific areas are accessible only to certain types of vehicles, such as Low Emission Zones (LEZ), or are subject to time-based driving restrictions.

Managing last-mile deliveries in such environments is a highly complex process, as many historical city centers have narrow streets, one-way traffic, and limited loading and unloading spaces. Geographic Information Systems (GIS) are valuable tools for identifying these constraints and guiding optimization engines to generate feasible routes. Urban consolidation centers can further enhance delivery efficiency by organizing distribution according to these constraints and reducing congestion in central areas.

3.2.3 Time dependence

In urban logistics, travel times are highly dependent on the time of day due to variations in traffic conditions, weather, and road closures. Traditional vehicle routing models often assume static travel times, which can reduce the accuracy of the proposed solutions. The Time-Dependent Vehicle Routing Problem (TDVRP) addresses this issue by incorporating travel time variability into the optimization process [18, 19].

In TDVRP, travel times are represented as time-dependent functions, reflecting different traffic conditions throughout the day. For example, a driver's 8-hour shift can be divided into four 2-hour intervals, each with different travel time costs due to varying traffic patterns. This enables the model to generate more realistic solutions and improve the quality of route planning [18, 20].

While several tools are available for measuring travel times and collecting historical traffic data, consolidating this information remains a challenge due to incomplete datasets. Experts recommend analyzing traffic patterns for each geographical area separately, as travel time variability is highly dependent on local conditions and traffic flow. Geographic and road conditions may also require the use of different types of vehicles, each with specific operational requirements.

One of the developments incorporated into the proposed solver is the explicit modeling of time dependence. Initially, travel times and distances between locations were handled using a single distance matrix and a single time matrix. We introduced a new data structure that supports multiple time periods, each with its own distance and time matrices. In addition, we modified JSprit's cost interfaces and time-propagation routines so that, both during solution construction and during objective evaluation, the solver selects the appropriate matrix based on the departure time and propagates arrival times accordingly.

Our time-dependent travel times/distances are represented as a piecewise-constant function through a set of time-sliced OD matrices. During construction and move evaluation, the solver selects the matrix corresponding to the departure time and propagates arrival times accordingly.

This modeling choice implicitly follows the common FIFO (First-In-First-Out) assumption used in time-dependent routing: for a given OD pair, departing later should not lead to an earlier arrival. Under FIFO, there is no benefit in deliberately delaying departure purely to obtain a shorter travel time, and therefore our solver does not include an explicit decision to "wait in order to arrive sooner"; vehicles depart as soon as service is completed, except for waiting that may arise from time-window feasibility (early arrival).

When the departure time is close to a period boundary, the travel time is evaluated using the slice corresponding to that departure time. As with any discretization, abrupt changes between consecutive slices may create small boundary effects. In practice, these effects can be mitigated by using finer time slices and/or by calibrating the time profiles so that they are consistent with FIFO at the chosen granularity. Modeling non-FIFO effects and explicit waiting decisions to exploit traffic regime changes is left for future work.

3.2.4 Portability to different urban contexts

Our solver is data-driven and city-agnostic: adapting it to a new environment requires configuration and data calibration, not code changes. City-specific characteristics are mapped to inputs as follows:

- Topography and network: time-of-day distance and travel-time matrices.
- Service-time distributions by zone/time band that incorporate parking/search time and curb-access frictions.
- Local regulation and policies: low-emission zones, access bans, and delivery windows, with rules by vehicle class and time of day.
- Logistics infrastructure: depots, UCCs, and micro-hubs as facility sets with opening hours and capacities (or candidate sites for scenario testing).
- Fleet and objectives: vehicle catalog (capacities, fixed/variable costs, emission factors, driving/access restrictions) and objective weights (e.g., cost vs. emissions).

3.3 Optimization algorithm

The Large Neighborhood Search (LNS) algorithm was employed in the resolution of the proposed solver. This algorithm is a metaheuristic in which the neighborhood of a solution is implicitly defined through the use of destroy and repair operators. The destroy operator removes a portion of the current solution, while the repair operator reconstructs the destroyed portion, aiming to find an improved solution. The destroy process is often implemented with a degree of randomness, allowing different aspects of the current solution to be modified in order to explore a broader search space. Compared to classical local search metaheuristics, LNS employs a significantly larger neighborhood exploration strategy, which enhances its ability to escape local optima and explore more promising regions of the solution space [21, 22].

The LNS algorithm is a hybrid metaheuristic that combines multiple destroy and repair operators also referred to as ruin and recreate operators along with solution acceptance strategies to guide the search process. This hybrid nature allows it to adapt to various problem structures and makes it suitable for solving complex instances of the vehicle routing problem. Consequently, the LNS algorithm is frequently integrated into many libraries related to vehicle routing optimization [23].

3.3.1 Ruin and recreate operators

The ruin and recreate strategy plays a crucial role in the performance of the Large Neighborhood Search (LNS) algorithm. We follow the LNS terminology and use destroy and repair to denote the removal and reinsertion phases. The JSprit library we build upon adopts the synonymous terms ruin and recreate; when we refer to specific JSprit operators/classes we keep that naming. In what follows, $destroy \equiv ruin$ and $repair \equiv recreate$; there is no methodological difference implied both describe the same pattern of removing a subset of visits and reinserting them with insertion heuristics.

The optimization core follows LNS scheme, where a current solution is iteratively partially destroyed and then repaired to explore a large neighborhood efficiently. In each iteration, a subset of assigned jobs is removed using destroy operators (e.g., random or related removal), and the solution is reconstructed using repair/insertion heuristics (including regret-based insertion), after which standard acceptance and stopping criteria are applied.

3.4 Implementation details

The architecture of the proposed solver is structured around the distinct responsibilities and functionalities of its components, each designed to perform specialized tasks that collectively contribute to the consolidation of the optimization model. This architecture defines the interaction between the various components and their integration within the system.

To enhance realism in routing solutions and accommodate more complex constraints, we introduced several modifications to JSprit. Specifically, we structured the problem model to take a JSON-formatted input file containing key data on pickups/deliveries, vehicles, vehicle types, and driving access restrictions. Based on this input, we constructed the model using JSprit's data structures, incorporating modifications in different components of the model and algorithm, as detailed below.

The first major enhancement introduced to our model was the incorporation of robustness into the solution by considering stochastic travel and service times. To achieve this, the problem formulation was extended to include service time parameters for both pickups and deliveries, allowing end users to adjust these values based on their operational experience. Additionally, travel times were estimated from historical data, contributing to a more realistic and resilient optimization process.

Stochastic travel times are closely related to time-dependent functions, which model varying traffic conditions throughout the day. To accommodate this, we modified JSprit's internal data structures to manage distance and time matrices segmented into different time intervals. This modification accounts for fluctuating traffic patterns, improving the accuracy of route planning.

Additionally, our input data defines multiple vehicle types, each associated not only with cost and capacity parameters but also with distinct distance and time matrices tailored to different periods of the day.

Our model also includes driving constraints, recognizing that certain delivery or pickup locations may have temporary or permanent access restrictions. These restrictions are defined within the input data using geospatial polygons that outline restricted areas. These constraints are then incorporated into the generation of distance and time matrices, ensuring that the routing algorithm respects access limitations and delivers feasible solutions.

One of the most critical components of the algorithm is the objective function. In JSprit, this function is originally defined to minimize the total distance traveled and the total time required, combining both components into the final solution. To further reduce urban traffic congestion and greenhouse gas emissions, we modified this function to incorporate operational costs and the emissions generated by each vehicle. For instance, minimizing the number of vehicles used not only reduces investment costs but also decreases pollutant emissions.

Balancing multiple objectives is often complex, as some may conflict with one another. To address the simultaneous optimization of different objectives, our model employs the weighted sum method. In this approach, each objective function is assigned a weighting coefficient that reflects its priority or relevance, and the goal is to minimize the weighted sum of all objectives.

The process involves assigning specific weights to each objective in order to construct a weighted sum for a multi-objective function. These weights reflect the relative importance of each objective, ensuring a balanced optimization process that aligns with predefined priorities. Initially, we assigned weights to each objective ensuring that their total sum equals 1 while defining the relative importance of each objective. This approach allows for a balanced optimization process that aligns with predefined priorities. By adjusting these coefficients, the model can emphasize certain criteria over others, facilitating trade-off analysis and achieving a more tailored solution.

All data used to calculate operational costs and emissions were integrated into the problem model, forming part of the input data for each vehicle, taking into account the specific vehicle type.

The following points summarize the main implementation contributions of the proposed solver. They cover the data model and interfaces, the operationalization of time dependence and uncertainty, the encoding of constraints and fleet heterogeneity, and the multi-objective optimization with its integration into JSprit.

- **Data ingestion and modeling:** JSON input schema unifying pickup/delivery jobs, vehicles, vehicle types, and access rules; the schema is translated into the internal data structures of the solver and JSprit.
- **Time-dependent routing:** replacement of static matrices with time-slice profiles (distance and time matrices per period) and selection of the appropriate slice based on the departure time during insertion and move evaluation.
- **Operational robustness:** explicit per-shipment service-time parameters (with configurable defaults) and travel-time estimation from historical speed profiles to capture typical congestion patterns.
- **Fleet heterogeneity:** vehicle types with specific costs/capacities and, optionally, per-period time and distance matrices for each type.
- **Access and driving restrictions:** restricted areas represented as geospatial polygons; constraints embedded in matrix generation to ensure feasible access in the resulting routes.
- **Extended objective function:** augmentation of the classic distance/time criterion to include operational costs and per-vehicle emissions/energy, encouraging fewer vehicles and lower environmental impact.
- **Configurable multi-objective aggregation:** weighted-sum with scenario-specific weights to prioritize criteria and enable trade-off analysis without altering the algorithmic core.
- **JSprit integration:** adapted cost interfaces and time-propagation routines to handle time-dependent functions and additional feasibility checks, while remaining compatible with the search operators.
- **Reproducible configuration:** all key parameters (weights, defaults, temporal profiles) managed via data-loading configuration, facilitating traceability across scenarios.

4 Experiments and set-up

For the validation and analysis of the proposed solver, a set of specific indicators was defined. These indicators serve as a crucial tool for structuring the evaluation framework, enabling both the analysis of existing phenomena in the case studies and the assessment of their impact.

The set of 19 indicators is structured as follows:

- Environmental and Emission Indicators:
  CO concentration, SOx concentration, NOx concentration, NH3 concentration, PM10 concentration, CO2 emissions, CH4 emissions, NO2

emissions, and social costs associated with air quality and greenhouse gas (GHG) emissions.
- Logistics and Mobility Indicators: Number of shipments, number of routes, total kilometers covered (including walking), total kilometers covered by green modes (including walking), total vehicle-kilometers traveled by freight vehicles, total vehicle-kilometers traveled by green freight vehicles, vehicle utilization factor.
- Economic Indicators: Fixed costs, running costs, and capital costs.

The classification matrix of these indicators, as presented in Table 2, considers the following impact areas and evaluation criteria:

- Environment & Society: Air quality, GHG emissions, and social costs.
- Transport & Mobility: Accessibility, urban freight transport (UFT) vehicles, and operational costs.

This structured classification provides a comprehensive framework for evaluating the effects of different logistics strategies, allowing for a multi-dimensional assessment of both environmental and operational impacts.

Three scenarios were designed with the primary objective of validating the proposed route planning approach and gaining a deeper understanding of the impact that different policies may have on enhancing the sustainability of last-mile logistics. Specifically, the following three scenarios were considered:

- Scenario 1: Represents the current operations of a logistics operator in the city of Zaragoza, serving as a baseline for comparing the results of the other scenarios.
- Scenario 2: Corresponds to the implementation of a Zero Emission Zone (ZEZ) in the city center of Zaragoza.
- Scenario 3: Simulates the establishment of an Urban Consolidation Center (UCC) in the city center of Zaragoza.

Furthermore, for each of the defined scenarios, three different fleet compositions were simulated, varying in terms of the vehicle electrification ratio. This approach aims to assess the impact of different fleet electrification levels on each scenario. The following sections provide a detailed description of the scenarios and the fleet compositions evaluated.

4.1 Scenario 1: baseline Scenario

Scenario 1 represents the current state of postal operations in the urban area of Zaragoza, providing a baseline for evaluating the impact of different logistics strategies. This scenario reflects the existing infrastructure, operational framework, and delivery patterns of the postal operator, which are structured as follows:

Infrastructure and Distribution Network:

- Nine Delivery Units (DUs): These units are strategically distributed throughout the city and are responsible for the delivery of postal items and small parcels. The predominant delivery method from

Table 2 Impact areas and indicators

Impact Area	Criteria	Indicator	Data/unit
Environment & Society	Air quality	CO concentration	g/day
		SOx concentration	g/day
		NOx concentration	g/day
		NH3 concentration	g/day
		PM10 concentration	g/day
	GHG emissions	CO2	g/day
		CH4	g/day
		NO2	g/day
	Social costs	Social costs of air quality and GHG emissions	€/day
Transport & Mobility	Accessibility	Number of shipments	n./day
		Number of routes	n./day
		Total km covered (including walking)	km/day
		Total km covered by green modes (including walking)	km/day
		Total veh-km covered by freight vehicles	Veh-km/day
		Total veh-km covered by green freight vehicles	Veh-km/day
	UFT vehicles	Vehicle utilisation factor	%/day
Operative costs	Fixed costs		€/day
	Running costs		€/day
	Capital costs		€/day

Table 3 Shipments per shift on different days

Day	Working Shift	Shipments	% of ship. in the ZEZ
Sept 13, 22	Morning	11723	15%
Sept 13, 22	Afternoon	3794	9%
Sept 14, 22	Morning	11429	15%
Sept 14, 22	Afternoon	4296	13%

Table 4 Vehicle composition under different electrification scenarios

Vehicle Type	Technology	Current Composition	Elect. 50%	Elect. 100%
Large Van	Combustion	9	9	–
Large Van	Electric	–	–	9
Motorcycle	Electric	12	27	27
Small Van	Combustion	37	15	–
Small Van	Electric	2	24	39

these units is walking routes, with postal workers covering assigned areas on foot.

- Two Special Service Units (SSUs): Located in the northern and southern areas of Zaragoza, these specialized units handle the distribution of larger parcels, which require motorized transport for efficient delivery.

The postal operator follows a two-shift system to manage deliveries efficiently:

- Morning shift: Runs from 7:00 AM to 3:00 PM.
- Afternoon shift: Runs from 3:00 PM to 10:00 PM.

However, due to operational requirements, postal workers dedicate time at the beginning of each shift to sort and prepare items for delivery and at the end of the shift to process undelivered items. As a result, the actual delivery time slots are: 8:00 AM to 2:00 PM for the morning shift, and 4:00 PM to 9:00 PM for the afternoon shift.

With regard to the demand data, for the purpose of analysis, we have selected the delivery data from 13 and 14 September 2022, which are presented in Table 3 distributed by day and by shift.

For the simulation of this scenario, as well as the subsequent ones, three different fleet compositions were defined to assess their environmental and operational impact based on fleet electrification levels. The compositions considered are: (1) Current fleet composition, preserving the existing configuration of the logistics operator's vehicle fleet; (2) 50% Electrification, where approximately half of the fleet is electrified while preserving the same vehicle typology (e.g., motorbikes, small vans, large vans); and (3) 100% Electrification, in which all vehicles are fully electric while maintaining the current fleet structure for comparability. The corresponding vehicle distribution by type and technology for each composition is detailed in the following Table 4.

This study is energy agnostic and does not explicitly model Electric Vehicle (EV) constraints (battery range, state of charge (SOC), charging times). Results should be read as an upper bound relative to Electric Vehicle constrained operations. Future work will extend the model with: (i) per-route range and minimum SOC reserve, (ii) depot/opportunity charging activities with power limits and charging times, and (iii) stochastic evaluation of consumption and ambient effects.

4.2 Scenario 2: deployment of a low emission zone

The second scenario analyzed in this study considers the establishment of a Zero Emission Zone (ZEZ) in the historic center of Zaragoza, as outlined in Fig. 2.

The implementation of this ZEZ imposes access restrictions on polluting vehicles between 7:00 AM and 11:00 PM, significantly impacting postal operations within the designated area. Under these conditions, combustion-engine vehicles would no longer be permitted to enter the zone for deliveries, requiring an adaptation of logistics strategies. Consequently, only postal workers on foot or using electric vehicles would be able to carry out deliveries within this restricted zone.

To assess the impact of these restrictions, Table 3 presents the distribution of orders affected by the ZEZ. The data reveal that the percentage of deliveries falling within the restricted area fluctuates between 9 and 15%, depending on the shift and the day. This scenario serves as a basis for evaluating the feasibility of alternative last-mile delivery strategies, such as increased fleet electrification or micro-consolidation centers, to ensure operational efficiency while complying with urban sustainability policies.

Similar to the previous scenarios, three different fleet compositions were defined based on fleet electrification levels, following the same criteria as in Scenario 1: Current Fleet Composition, 50% Electrification, and 100% Electrification.

4.3 Scenario 3: deployment of an urban consolidation center

The third scenario analyzed involves the establishment of an Urban Consolidation Centre (UCC) in the city center of Zaragoza, specifically at the current location of the San Vicente de Paúl Market. This site has been identified as a key location for the SENATOR project. Figure 3 illustrates the location of the proposed Urban Consolidation Centre along with its designated area of influence, which encompasses postal codes 50,001 to 50,005, 50,008, 50,010, and 50,013.

The Urban Consolidation Centre is designed to streamline the distribution of larger parcels, functioning

Fig. 2 Delimitation of the area considered for the deployment of the ZEZ in Zaragoza

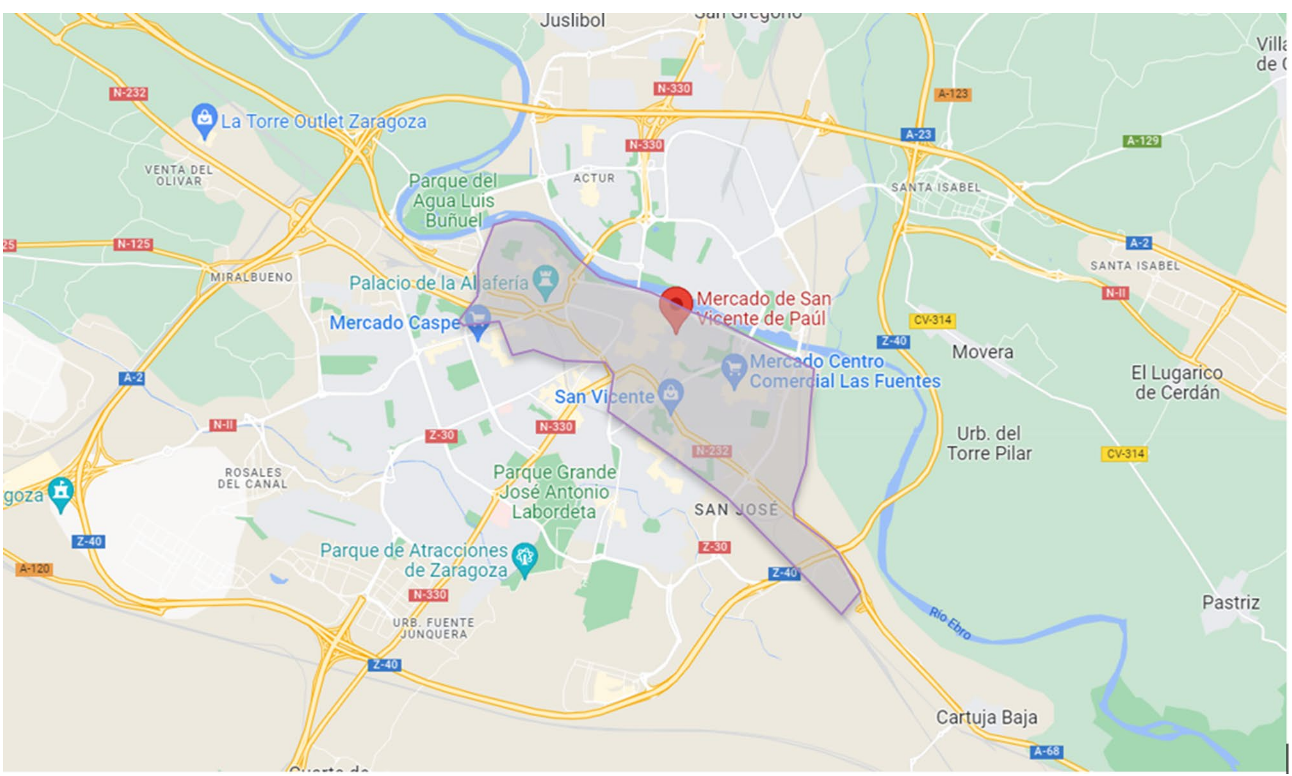

Fig. 3 Location of the UCC and area of influence of the UCC

Table 5 Number of shipments per day and shift to be delivered at the area of influence of the UCC

Day	Shift	Shipments	% over all shipments	% SSE shipments
Sept 13, 22	Morning	459	4%	37%
Sept 13, 22	Afternoon	903	24%	33%
Sept 14, 22	Morning	400	3%	34%
Sept 14, 22	Afternoon	1167	27%	37%

similarly to the existing SSUs. To simulate the impact of this consolidation strategy, we redirected all parcels originally handled by SSUs within the designated area to be processed and distributed from the new UCC.

Table 5 presents the volume of shipments that would be processed through the Urban Consolidation Centre under this scenario. The data reveal that, when considering both postal items and parcels, the proportion of deliveries routed through the UCC is relatively low in the morning shift but significantly higher in the afternoon shift. However, when focusing specifically on medium and large parcels, which are currently distributed by SSUs, we observe that the share handled by the UCC ranges between 33 and 37%. In other words, under this scenario, approximately one-third of all medium and large parcels would be processed and distributed through the Urban Consolidation Centre, reducing the reliance on direct vehicle-based distribution within the city center.

Similar to the previous two scenarios, three different fleet compositions were defined based on fleet electrification levels, following the same criteria as in Scenario 1: Current Fleet Composition, 50% Electrification, and 100% Electrification. However, in this case, the Urban Consolidation Centre (UCC) is assumed to operate its own fleet, supplied by various logistics operators.

Specifically, the UCC fleet consists of 20 small vans, with their distribution across combustion and electric vehicles varying according to the three fleet alternatives: (20 combustion – 0 electric) for the Current Fleet Composition, (10 combustion – 10 electric) for 50% Electrification, and (0 combustion – 20 electric) for 100% Electrification. The baseline postal fleet remains unchanged; the UCC fleet is an exogenous capacity addition used to serve parcels within the UCC area of influence.

5 Results

5.1 Results for Scenario 1

Table 6 presents the results for Scenario 1, evaluating different fleet compositions in terms of electrification. As observed, the current fleet, which predominantly consists of conventional combustion vehicles, exhibits a high environmental impact, with associated social costs exceeding €17.75 per day. However, as the fleet's electrification level increases, emissions are significantly reduced—by approximately 85%—along with the corresponding social costs. This substantial reduction is attributed to the fact that the postal operator's fleet is oversized to accommodate peak demand, meaning that during periods of intermediate demand, 50% electrification would already achieve an 85% reduction in emissions. As expected, a fully electrified fleet (100% Electrification) results in a 100% reduction in environmental impact.

From an operational perspective, the analysis confirms that service levels remain unaffected, as the fleet

Table 6 Results for scenario 1

Indicator	Unit	Current Composition	50% Electric Vehicles	100% Electric Vehicles
CO concentration	g/day	152.58	11.01	0.00
SOx concentration	g/day	1.00	0.14	0.00
NOx concentration	g/day	589.18	108.37	0.00
NH3 concentration	g/day	2.90	0.40	0.00
PM10 concentration	g/day	27.83	4.90	0.00
CO2	g/day	197,936.30	27,555.43	0.00
CH4	g/day	14.68	1.21	0.00
NO2	g/day	11.64	1.97	0.00
Social costs of air quality and GHG emissions	€/day	17.75	2.48	0.00
Number of shipments	n./day	15,106.00	15,083.50	15,107.50
Number of routes	n./day	208.00	206.00	207.00
Total km covered (including walking)	km/day	2,722.30	2,702.36	2,722.50
Total km covered by green modes (including walking)	km/day	1,390.51	2,518.24	2,722.50
Total veh-km covered by freight vehicles	Veh-km/day	1,683.94	1,664.00	1,684.14
Total veh-km covered by green freight vehicles	Veh-km/day	352.15	1,479.88	1,684.14
Vehicle utilisation factor	%/day	0.33	0.33	0.33
Fixed costs	€/day	24,612.49	24,683.48	25,008.18
Running costs	€/day	187.99	112.47	106.57
Capital costs	€/day	1,154,800.00	1,507,800.00	1,842,800.00

electrification process does not compromise delivery capacity or efficiency. However, in terms of operational costs, the transition towards an electric fleet incurs a moderate increase in fixed costs and a more substantial increase in capital costs (up to 60%), primarily due to the higher acquisition costs of electric vehicles. Conversely, running costs are reduced by 40%, reflecting the lower energy and maintenance expenses associated with electric vehicle operations.

5.2 Results for Scenario 2

Table 7 presents the results for Scenario 2, which evaluates the impact of implementing a ZEZ in Zaragoza. In terms of environmental impact, the findings are consistent with those observed in Scenario 1, as expected, since fleet electrification remains the primary determinant of emissions reduction.

However, notable differences arise in delivery performance. The number of shipments delivered increases by approximately 1% with higher fleet electrification, compared to the baseline scenario (Scenario 1), where the current postal fleet operates without restrictions. This increase is attributed to the fact that as electrification levels rise, more vehicles are permitted to enter the Low Emission Zone, enabling a higher number of parcels to be delivered within the restricted area. Conversely, in the current fleet composition, where a significant proportion of vehicles are combustion-powered, access restrictions result in a 1% decrease in deliveries within the zone.

5.3 Results for Scenario 3

Table 8 presents the results for Scenario 3, which evaluates the impact of deploying an Urban Consolidation Centre (UCC) in Zaragoza.

One of the most significant findings in this scenario is the increase in the number of parcels delivered, which rises by 1% compared to Scenario 0. This improvement is attributed to the greater availability of vehicles within the UCC, which allows for the fulfillment of a higher number of orders. However, this increase in deliveries comes at a considerable operational cost, particularly in terms of vehicle kilometers traveled, which rise from 2722 km to 2832 km. This increase is likely due to the fact that, in Scenario 0, more remote orders were left unfulfilled, whereas in this scenario, the expanded fleet allows these orders to be delivered. However, this leads to higher travel distances, increased fuel or energy consumption, and a significant rise in emissions.

Additionally, the results indicate that the increased number of vehicles results in a lower vehicle utilization rate, reducing overall efficiency. This suggests that while a UCC-based logistics model may improve delivery coverage, it may also introduce inefficiencies if not properly optimized.

6 Conclusions

The main contribution of this work is the implementation and validation of an optimization module within the broader SENATOR system, with a strong emphasis on data preparation and system integration. Rather than introducing a new VRP metaheuristic, we build on the JSprit engine and provide the missing operational

Table 7 Results for Scenario 2

Indicator	Data/unit	Current Composition	50% Electric Vehicles	100% Electric Vehicles
CO concentration	g/day	144.58	11.01	0.00
SOx concentration	g/day	0.97	0.14	0.00
NOx concentration	g/day	566.32	108.36	0.00
NH3 concentration	g/day	2.82	0.40	0.00
PM10 concentration	g/day	26.52	4.90	0.00
CO2	g/day	190,998.58	27,553.29	0.00
CH4	g/day	13.80	1.21	0.00
NO2	g/day	11.34	1.97	0.00
Social costs of air quality and GHG emissions	€/day	17.13	2.48	0.00
Number of shipments	n./day	14,930.00	15,058.50	15,107.50
Number of routes	n./day	208.00	206.00	207.00
Total km covered (including walking)	km/day	2,695.79	2,703.98	2,722.50
Total km covered by green modes (including walking)	km/day	1,418.84	2,519.88	2,722.50
Total veh-km covered by freight vehicles	Veh-km/day	1,657.42	1,665.62	1,684.14
Total veh-km covered by green freight vehicles	Veh-km/day	380.47	1,481.52	1,684.14
Vehicle utilisation factor	%/day	0.32	0.33	0.33
Fixed costs	€/day	24,328.19	24,659.85	25,008.18
Running costs	€/day	184.26	112.59	106.57
Capital costs	€/day	1,198,800.00	1,507,800.00	1,842,800.00

Table 8 Results for scenario 3

Indicator	Data/unit	Current Composition	50% Electric Vehicles	100% Electric Vehicles
CO concentration	g/day	153.53	6.40	0.00
SOx concentration	g/day	1.07	0.09	0.00
NOx concentration	g/day	680.25	82.29	0.00
NH3 concentration	g/day	3.04	0.23	0.00
PM10 concentration	g/day	32.43	3.95	0.00
CO2	g/day	211,446.44	17,617.88	0.00
CH4	g/day	15.37	0.92	0.00
NO2	g/day	12.70	1.26	0.00
Social costs of air quality and GHG emissions	€/day	18.97	1.59	0.00
Number of shipments	n./day	15,294.00	15,270.50	15,295.00
Number of routes	n./day	213.00	212.50	212.50
Total km covered (including walking)	km/day	2,832.70	2,827.23	2,823.54
Total km covered by green modes (including walking)	km/day	1,388.22	2,699.73	2,823.54
Total veh-km covered by freight vehicles	Veh-km/day	1,794.34	1,788.87	1,785.17
Total veh-km covered by green freight vehicles	Veh-km/day	349.86	1,661.37	1,785.17
Vehicle utilisation factor	%/day	0.31	0.32	0.32
Fixed costs	€/day	25,217.52	25,415.03	25,619.48
Running costs	€/day	201.02	115.15	113.34
Capital costs	€/day	1,273,800.00	1,678,800.00	2,000,800.00

layers required in real last-mile deployments: a unified input data model, ingestion and validation routines, and interfaces that translate operational data (jobs, heterogeneous fleets, and urban access rules) into a solvable VRP instance. On top of this integration layer, we implement practical extensions and configurations, such as time-dependent travel time profiles and a configurable multi-objective cost/emissions criterion that enable scenario based analyses and policy evaluation in realistic case studies. The case studies demonstrate that these engineering and modeling contributions materially affect solution feasibility and quality in practice, while the optimization core remains compatible with established JSprit search operators.

The research is conducted within the framework of the SENATOR European project, where a pilot study in Zaragoza evaluates the impact of low-emission zones, urban consolidation centers, and fleet electrification. Through three experimental scenarios, the study demonstrates how advanced routing optimization can enhance operational efficiency, reduce emissions, and improve urban mobility. Key contributions include:

- Robust route optimization, incorporating stochastic travel and service times to improve delivery reliability.
- Multimodal logistics integration, considering alternative transport modes to improve urban delivery efficiency.
- Time-dependent travel modeling, capturing varying traffic conditions throughout the day to enhance routing accuracy.
- Evaluation of fleet electrification strategies, quantifying their impact on environmental and operational performance.

The results of the pilot study indicate that fleet electrification can reduce emissions by up to 85%, while urban consolidation strategies can improve delivery efficiency but may also increase vehicle kilometers traveled if not properly optimized. These findings provide valuable insights for urban planners and logistics operators seeking to balance efficiency, sustainability, and cost-effectiveness in last-mile delivery systems.

Our current implementation targets offline planning and does not yet include online replanning for dynamic requests or an explicit stochastic optimization model; these aspects are left for future work. For future work, we aim to incorporate a dynamic component into the solver that accounts for the continuous arrival of new deliveries and unexpected disruptions, such as vehicle breakdowns. This enhancement will enable real-time route re-optimization and better adaptability to changing conditions. Furthermore, we intend to compare single-period optimization with multi-period optimization across the day, integrating live traffic data. By embedding these dynamic capabilities, the solver will support a more responsive and resilient last-mile logistics system, capable of handling operational uncertainties while ensuring efficiency and sustainability. In addition, we are conducting pilots in real-world operational settings with diverse characteristics, which will allow us to assess the solver performance and identify its potential limitations across different problem types.

Author Contribution
All authors contributed equally to the conception and development of this article, as well as to the analysis, interpretation, writing, and revision of the manuscript. All authors read and approved the final manuscript.

Data availability
Not applicable.

Materials availability
Part of the results presented in this scientific article were obtained within the framework of the European project SENATOR.

Code availability
Not applicable.

Declarations

Competing interests
The authors declare that they have no competing interests.

Received: 15 April 2025 / Accepted: 5 March 2026
Published online: 24 March 2026

References
1. Na, H. S., Kweon, S. J., & Park, K. (2021). Characterization and design for last mile logistics: A review of the state of the art and future directions. *Applied Sciences*, *12*(1), 118.
2. Plazier, P., Rauws, W., Neef, R., & Buijs, P. (2024). Towards sustainable last-mile logistics? Investigating the role of cooperation, regulation, and innovation in scenarios for 2035. *Research in Transportation Business and Management*, *56*, 101198.
3. Mohammad, W. A., Diab, Y. N., Elomri, A., & Triki, C. (2023). Innovative solutions in last mile delivery: Concepts, practices, challenges, and future directions, in supply chain forum. *An International Journal*, *24* Taylor & Francis, 151–169.
4. Golden, B. L., Raghavan, S., & Wasil, E. A. (2008). *The vehicle routing problem: Latest advances and new challenges* (Vol. *43*). Springer Science & Business Media.
5. Toth, P., & Vigo, D. (2002). An overview of vehicle routing problems. In P. Toth & D. Vigo (Eds.), *The vehicle routing problem* (pp. 1–26). SIAM. https://doi.org/10.1137/1.9780898718515.ch1
6. Lahyani, R., Khemakhem, M., & Semet, F. (2015). Rich vehicle routing problems: From a taxonomy to a definition. *European Journal of Operational Research*, *241*(1), 1–14.
7. Lesch, V., König, M., Kounev, S., Stein, A., & Krupitzer, C. (2022). Tackling the rich vehicle routing problem with nature-inspired algorithms. *Applied Intelligence*, *52*(8), 9476–9500.
8. Mardešić, N., Erdelić, T., Carić, T., & D-Test-Urasević, M. (2023). Review of stochastic dynamic vehicle routing in the evolving urban logistics environment. *Mathematics*, *12*(1), 28.
9. Psaraftis, H. N., Wen, M., & Kontovas, C. A. (2016). Dynamic vehicle routing problems: Three decades and counting. *Networks*, *67*(1), 3–31.
10. Wang, S., Han, C., Yu, Y., Huang, M., Sun, W., & Kaku, I. (2022). Reducing carbon emissions for the vehicle routing problem by utilizing multiple depots. *Sustainability*, *14*(3), 1264.
11. Patella, S. M., Grazieschi, G., Gatta, V., Marcucci, E., & Carrese, S. (2020). The adoption of green vehicles in last mile logistics: A systematic review. *Sustainability*, *13*(1), 6.
12. Christiaens, J., & Berghe, G. V. (2020). Slack induction by string removals for vehicle routing problems. *Transportation Science*, *54*(2), 417–433.
13. Sabet, S., & Farooq, B. (2022). Green vehicle routing problem: State of the art and future directions. *IEEE Access*, *10*, 101622–101642.
14. Braekers, K., Ramaekers, K., & Van Nieuwenhuyse, I. (2016). The vehicle routing problem: State of the art classification and review. *Computers & Industrial Engineering*, *99*, 300–313.
15. Oyola, J., Arntzen, H., & Woodruff, D. L. (2017). The stochastic vehicle routing problem, a literature review, part ii: Solution methods. *EURO Journal on Transportation and Logistics*, *6*(4), 349–388.
16. Oyola, J., Arntzen, H., & Woodruff, D. L. (2018). The stochastic vehicle routing problem, a literature review, part i: Models. *EURO Journal on Transportation and Logistics*, *7*(3), 193–221.
17. Bernardo, M., & Pannek, J. (2018). Robust solution approach for the dynamic and stochastic vehicle routing problem. *Journal of Advanced Transportation*, *2018*(1), 9848104.
18. Hou, D., Fan, H., Ren, X., Tian, P., & Lv, Y. (2021). Time-dependent multi-depot heterogeneous vehicle routing problem considering temporal–spatial distance. *Sustainability*, *13*(9), 4674.
19. Figliozzi, M. A. (2012). The time dependent vehicle routing problem with time windows: Benchmark problems, an efficient solution algorithm, and solution characteristics. *Transportation Research Part E: Logistics & Transportation Review*, *48*(3), 616–636.
20. Mor, A., & Speranza, M. (2020). Vehicle routing problems over time: A survey. 4OR, *18*(2), 129–149.
21. Ghilas, V., Demir, E., & Van Woensel, T. (2016). An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows and scheduled lines. *Computers & Operations Research*, *72*, 12–30.
22. Grangier, P., Gendreau, M., Lehuédé, F., & Rousseau, L. M. (2016). An adaptive large neighborhood search for the two-echelon multiple-trip vehicle routing problem with satellite synchronization. *European Journal of Operational Research*, *254*(1), 80–91.
23. Gu, W., Cattaruzza, D., Ogier, M., & Semet, F. (2019). Adaptive large neighborhood search for the commodity constrained split delivery vrp. *Computers & Operations Research*, *112*, 104761.

Publisher's Note