Optimisation of Residential Waste Collection

Balancing Travel Time and Visual Attractiveness in Side Loader Routes

Amber van der Helm

Delft University of Technology

TU Delft

Optimisation of Residential Waste Collection

Balancing Travel Time and Visual Attractiveness in Side Loader Routes

by

Amber van der Helm

to obtain the degree of **Master of Science** in Transport, Infrastructure and Logistics

at Delft University of Technology,

to be defended publicly on Monday March 30, 2026 at 13:00.

Student number:	5164303	
Date:	March 2026	
Thesis committee:	Dr. B. Atasoy,	TU Delft, Faculty 3mE (chair)
	Dr. ir. A. Pel,	TU Delft, Faculty CEG
	MSc. K. Hauge,	AMCS Group
	MSc. K. Van Duurling,	AMCS Group
Report number:	2025.TIL.9170	
Cover:	adapted from AMCS Group (2025b)	

TU Delft amcs

Preface

This thesis marks the final step of my master's programme Transport, Infrastructure and Logistics at Delft University of Technology. From the start of my studies, I was motivated by the idea of using analytical and optimisation methods to tackle real-world problems. During the programme, I became increasingly interested in how mathematical models can support operational decision-making in logistics systems. This interest ultimately led me to the topic of residential waste collection routing, where technical optimisation meets practical constraints and human considerations.

Working on this thesis has been an intense but rewarding process. A large part of the research involved trial and error, particularly when defining visual attractiveness metrics and developing an optimisation framework around them. Translating a subjective concept such as route attractiveness into quantitative, mathematical terms proved to be challenging, and often required revisiting assumptions and redesigning models. Similarly, developing a metaheuristic that behaved as intended was very much a matter of detail, from choosing appropriate operators to tuning parameters so that the algorithm balanced exploration and performance. Looking back, I am proud of the persistence this process required and of the final optimisation framework that developed from it.

This thesis was carried out in collaboration with AMCS, whose practical perspective shaped this research. I would like to sincerely thank my supervisors at AMCS, Koen van Duurling and Kristian Hauge, for their insights, openness, and willingness to share real operational challenges. Their input ensured that this research remained grounded in practice. I am also very grateful to my supervisors at the TU Delft, Bilge Atasoy and Adam Pel, for their guidance, critical feedback, and support throughout the project. Their questions consistently helped me sharpen both the scientific and practical relevance of my work.

I would also like to thank my friends, with whom I spent many hours writing together at the university, sharing ideas, doubts, and progress over cups of coffee. Those moments made the process lighter and reminded me that I was not working through it alone. I am very thankful to my family, who always believed in me and reminded me that I could do this, even when they only partly understood what I was actually working on. Their trust and encouragement meant more than they may realise.

Finally, I would like to thank my partner, Bram Verweij. Throughout this project, he had to deal with me at my most stressed moments, helped calm me down when things felt overwhelming, and was always willing to spar and brainstorm about ideas. His patience, support, and confidence in me were invaluable.

I hope you enjoy reading this thesis.

Amber van der Helm
Delft, March 2026

Abstract

Residential kerbside waste collection using side-loading vehicles introduces routing challenges that differ fundamentally from conventional routing problems. Because side loaders can collect waste from only one side of the street per pass, mandatory double traversal of streets is required. When a single route is optimised solely for travel time, this often results in fragmented service of areas, long revisit intervals between opposite kerbsides, and structurally complex routes that are difficult to interpret and execute. While travel time efficiency remains important for operational efficiency, route quality in this context also depends on visual attractiveness. Visually attractive routes improve route clarity for drivers, operational robustness to disruptions, and service predictability for residents. Despite its operational relevance, visual attractiveness for optimisation of a single route in side loader settings has received limited attention in existing literature.

This study developed a single-route optimisation framework for residential kerbside waste collection that accommodates double traversal while balancing travel time minimisation and visual attractiveness maximisation under realistic operating conditions. The problem was formulated as an asymmetric travelling salesman problem (ATSP) on a directed network, where nodes represent directed kerbside street segments and travel times are computed as the fastest paths on the underlying road network respecting traffic directions and U-turn restrictions. Visual attractiveness was quantified using complementary metrics for regional compactness, local compactness, and route complexity. The metrics were defined on an abstract network using straight-line connections to maintain computational tractability. They were combined with travel time in a weighted sum objective and optimised by simulated annealing.

The optimisation framework was evaluated on 12 real municipal case studies. In 11 cases, the optimisation approach generated routes that substantially improved visual attractiveness compared to the travel-time-optimal routes produced by existing systems of AMCS. Neighbourhood fragmentation, revisit spans between opposite kerbsides, and intra-route crossings were consistently reduced. These findings indicate that optimising visual attractiveness metrics defined on the abstract network produces more visually attractive routes on the underlying road network.

The results also reveal trade-offs. First, trade-offs arise among the visual attractiveness metrics. Reducing neighbourhood fragmentation, revisit spans between opposite kerbsides, and intra-route crossings resulted in larger local jumps measured in directed Euclidean distance. This confirms that visual attractiveness is multi-dimensional and cannot be optimised across all dimensions simultaneously. Second, improvements in these measures were generally accompanied by increases in total travel time, reflecting the fundamental trade-off between efficiency and visual attractiveness. The sensitivity analysis confirmed this behaviour. Increasing the preference weights of the visual attractiveness metrics using a scaling factor reduced neighbourhood fragmentation, revisit spans, and intra-route crossings while increasing the total travel time and directed Euclidean distance. Very small scaling factors produced less consistent improvements in visual attractiveness, whereas larger factors yielded diminishing improvements while travel time increased substantially.

Overall, the findings demonstrated that visual attractiveness can be incorporated into an optimisation framework for a single waste collection route using a side loader. The optimisation framework suggests that visual attractiveness, captured by compactness and complexity metrics, can be balanced with travel time in a weighted sum objective function. The results indicate that the trade-offs between travel time and visual attractiveness can be managed in a controlled and interpretable manner. The selected preference weights provide a reasonable configuration for the analysed case studies. However, careful calibration remains necessary for other networks with different characteristics. Scientifically, this research extends existing metrics of visual attractiveness to the context of single side loader routes and clarifies the interaction between multiple routing objectives. Practically, it provides waste collection service providers with a configurable framework to improve route clarity and service predictability while allowing the balance between travel time efficiency and visual attractiveness to be adjusted.

Contents

List of Figures

List of Tables

1

Introduction

Global municipal solid waste generation is rising rapidly, driven by industrialisation, urbanisation, and consumerism. In 2023, global municipal solid waste production reached 2.1 billion tonnes and may nearly double to 3.8 billion tonnes by 2050 (United Nations Environment Programme, 2024). This poses challenges to infrastructure, sustainability, and public health, highlighting the need for more efficient municipal solid waste management. Waste collection is the most resource-intensive activity in municipal solid waste management, accounting for up to 80 percent of the costs (Das & Bhattacharyya, 2015). Therefore, optimising waste collection routes has become a priority for authorities and operators in order to reduce costs, traffic congestion, and environmental impact (United Nations Environment Programme, 2024).

Against this backdrop, this chapter introduces the problem, focusing on the implications of automated side-loading vehicles in residential waste collection. Then, it briefly reviews the current state of the literature and the research gap, before formulating the research objective and questions that guide this study. Finally, the chapter outlines the adopted methodology and the structure of the report.

1.1. Problem Description

Municipal solid waste collection systems are typically divided into three main categories, namely industrial, commercial, and residential waste collection (Golden et al., 2002). Industrial systems often use large roll-on/roll-off containers, where one container is moved at a time. Commercial systems rely on scheduled pickups of large bins at restaurants, offices, and retail centres. In contrast, residential collection often involves kerbside pickup, where vehicles traverse street segments to collect household waste. Among these, residential kerbside collection typically serves the largest number of stops, making it the most frequent and resource intensive. This operational complexity makes effective routing essential (Han & Ponce-Cueto, 2015; Kim et al., 2006).

In many suburban areas, municipalities are increasingly transitioning from rear-loading trucks to automated side-loading vehicles. Side loaders offer safety and labour advantages, as they require only a single operator and reduce manual handling (Archive Market Research, 2024). For instance, several Dutch municipalities have adopted side loaders in recent years (Blink Afvalbeheer, 2024; DAR NV, 2024; Gemeente Midden-Delfland, 2021). However, side loaders introduce a significant routing complication in kerbside waste collection. Since they can only collect waste from one side of the street, both kerbsides must be serviced in separate passes. This requirement is known as double traversal (Beliën et al., 2012). Double traversal effectively doubles parts of the network and forces vehicles to re-enter already serviced streets or neighbourhoods, often much later in the route. As a result, a side loader route becomes fragmented, reducing operational clarity for operators and residents.

These challenges are amplified by traffic directions and street geometry. Heavy collection vehicles often cannot make U-turns or zigzag between opposite kerbsides. Instead, they must follow detours to re-enter the same street or neighbourhood, further complicating the sequencing of collections and reducing operational clarity. In settings with multiple routes, operators prefer routes that are spatially compact, have clear service territories, and are easy to interpret. This notion is commonly referred

to as visual attractiveness in the vehicle routing literature (Kim et al., 2006; Poot et al., 2002; Sahoo et al., 2005). Visually attractive routes are not only easier for crews to follow, but also improve public acceptance and safety by providing predictable service patterns (Rossit et al., 2019). For a single side loader route, this means servicing nearby streets together, keeping time intervals between opposite kerbside visits short, and maintaining an intuitive route shape.

At the same time, efficiency remains important in residential waste collection to ensure that routes are feasible. Although many routing studies use distance minimisation as a proxy for reducing costs and emissions (Bautista et al., 2008; Rambandara et al., 2022), operators increasingly focus on travel time. The shortest routes are not necessarily the fastest, as arterial roads may enable quicker access than small streets despite longer travel distances. Companies such as Waste Management Inc. have demonstrated that optimising for travel time can substantially improve margins by enabling more collections per shift (Broughton, 2024). Thus, travel time is a key measure of efficiency and feasibility.

However, side loader routes that minimise efficiency do not necessarily display visually attractive behaviour (Kim et al., 2006; Rossit et al., 2019). Fast routes may excessively switch between neighbourhoods or postpone revisiting a street to service the other kerbside, while more attractive routes may increase the total travel time. This creates a trade-off between travel time efficiency and visual attractiveness. In particular, it highlights the complexity of designing a residential waste collection route for a side loader that is efficient and visually attractive under realistic operating conditions.

Technology providers such as AMCS, a global leader in waste management software, support route optimisation for residential waste collection (AMCS Group, 2025a). However, existing tools face difficulties when side loaders are involved, since double traversal creates challenges that are not well captured by traditional routing methods. Through collaboration with AMCS, this research addresses an industry challenge by optimising a single waste collection route for a side loader. This involves producing a feasible and logically ordered household service sequence under double traversal constraints.

In practice, AMCS routing solutions account for additional objectives and constraints, including vehicle capacities, crew schedules, heterogeneous fleets, disposal site planning, and stochastic traffic conditions, waste demand, and operational disruptions. To allow focused research, these aspects are intentionally excluded. Instead, this research restricts attention to a single-route setting with fixed demand and service requirements, enabling a detailed study of optimising routes for efficiency and visual attractiveness in side loader operations.

1.2. Scientific Relevance

The waste collection routing problem has been widely studied, in which vehicles collect waste from collection points or street segments and transport it to disposal facilities. Within this domain, a substantial body of literature has addressed visual attractiveness, a concept introduced by Poot et al. (2002) to capture the perceived clarity and coherence of routing solutions. Visually attractive routes are easier for crews to execute, improve the acceptance of routing plans by waste collection operators, and increase public acceptance by providing clear and predictable service patterns and vehicle territory separation (Poot et al., 2002; Rossit et al., 2019). Prior studies show that operators naturally prefer routes that are compact and largely free of crossings, even when these are not strictly optimal in distance or time (Van Rooij et al., 2003; Vickers et al., 2006).

In the literature, visual attractiveness is studied predominantly in vehicle routing problems, where a set of routes serves a region together. In this context, visually attractive solutions are commonly characterised by three components, namely compactness, non-overlapping, and low complexity. Compactness reflects how spatially concentrated the customers assigned to each route are and is typically measured using distance-to-centre or pairwise distance metrics. Non-overlapping focuses on the spatial separation between different routes, ensuring that vehicle service territories do not interfere with one another. This is often measured by inter-route crossings or shape indices such as the convex hull (Kim et al., 2006; Matis, 2008; Poot et al., 2002; Sahoo et al., 2005; Tang & Miller-Hooks, 2006). Low complexity captures the geometric simplicity of individual route shapes, favouring smooth and intuitive paths without unnecessary irregularities. Unlike compactness and non-overlapping, this concept is directly applicable to single route problems (Rossit et al., 2019).

For a single residential waste collection route for a side loader, unattractive behaviour shows differ-

ently. A single route lacks compactness when service is spatially fragmented, while non-overlapping is no longer defined between routes and is instead implicitly captured by compactness. This occurs when routes move to new areas before completing the current one, when nearby streets are not serviced consecutively, or when opposite kerbsides of the same street are visited far apart in time. Low complexity also remains relevant, as unnecessary crossings or irregular sequencing can disrupt intuitive route shapes. This behaviour is caused by double traversal, traffic directions, and U-turn restrictions, which force the vehicle to revisit already serviced areas and streets, often much later. However, most existing visual attractiveness metrics are defined for multiple routes, where attractiveness is assessed by comparing route territories. As a result, current metrics do not capture visual attractiveness for optimising a single side loader route, revealing a clear gap in the literature.

1.3. Research Objective and Questions

Building on the problem description and identified research gap, this study develops a single-route optimisation framework for residential kerbside waste collection that optimises visual attractiveness alongside travel time under realistic operating conditions relevant to side loaders. Scientifically, it addresses the lack of methods to quantify and model visual attractiveness for individual routes, particularly in side loader operations. Practically, integrating travel time efficiency and visual attractiveness reflects two core needs of the waste collection sector. Waste collection companies seek routes that are feasible and efficient, whereas operators and residents benefit from routes that are predictable and easy to follow. By studying these conflicting objectives together, this study aims to develop routing solutions that balance efficiency and operational clarity.

To study this research objective, the following main research question is proposed:

> *How can a residential kerbside waste collection route be designed to accommodate double traversal, while optimising visual attractiveness and travel time under relevant operating conditions?*

To answer this main research question, the following sub-questions are formulated:

SQ1. How have earlier studies identified and addressed double traversal and other relevant operational constraints in residential kerbside waste collection routing problems?
SQ2. Which performance metrics and multi-objective optimisation methods are most suitable for quantifying visual attractiveness maximisation alongside travel time minimisation?
SQ3. Which optimisation methods are appropriate for this multi-objective optimisation problem, given the need to handle large networks?
SQ4. How can the identified constraints and objectives be formulated into a mathematical optimisation model?
SQ5. How can the proposed mathematical model be solved effectively using a suitable optimisation method identified?
SQ6. How does the proposed optimisation framework perform on real municipal case studies compared to existing AMCS methods, particularly with respect to travel time and visual attractiveness?

1.4. Methodology and Report Structure

The remainder of this study follows a structured methodology to address the research questions. An overview of this methodology is provided in Figure 1.1. The study begins with a literature review in Chapter 2. This chapter establishes the scientific background and addresses modelling approaches for a single waste collection route using a side loader (SQ1), visual attractiveness objectives and multi-objective optimisation methods (SQ2), and optimisation methods for related routing problems (SQ3).

Chapter 3 introduces the empirical foundation of this study. It describes the 12 case study datasets provided by AMCS and presents descriptive statistics that characterise their size, structure, and spatial layout. These descriptives show that the case studies are sufficiently comparable to be analysed together, while still displaying enough variation to study model behaviour. The chapter also introduces a small test set, which enables exact optimisation for the selection of visual attractiveness metrics.

Building on the literature and data description, Chapter 4 focuses on the selection of visual attractiveness metrics. First, the routing problem is conceptually formulated and several candidate visual

attractiveness metrics are defined. These metrics are then translated into single-objective exact optimisation models, each isolating one metric while sharing a common routing structure (SQ4). This allows the optimisation effect of each metric on route structure to be analysed in isolation. Following this, the behaviour of the metrics is analysed in two steps. First, a correlation analysis on travel-time-optimal routes of all case studies examines how the metrics covary on realistic networks. Since solving each metric optimally on the full datasets is too computationally expensive, correlations can still provide insight into how the metrics behave on realistic networks. This is complemented by single-objective optimisation on the test set, which allows the effect of optimising for each metric to be inspected both numerically and visually. Together, these analyses support the selection of a set of complementary metrics for the optimisation framework.

Using the selected metrics, Chapter 5 proposes an optimisation framework to construct side loader routes that balance visual attractiveness and operational efficiency in terms of travel time. The chapter presents the conceptual and algorithmic framework of an optimisation method based on simulated annealing, in which the selected metrics are combined in a weighted sum objective function. Subsequently, the framework is verified and validated in Chapter 6 to ensure that it is implemented correctly and behaves as intended, after which a final set of model parameters is selected (SQ5).

Next, Chapter 7 evaluates the final optimisation framework on the case studies. Numerical and visual comparisons with the AMCS baseline are used to assess its practical performance (SQ6). A sensitivity analysis of the objective function weights is conducted to examine the trade-offs between travel time and visual attractiveness. The research concludes with a discussion and conclusion in Chapter 8 and Chapter 9, which synthesise the findings, revisit the research questions, reflect on limitations, and provide recommendations for future research and practice.

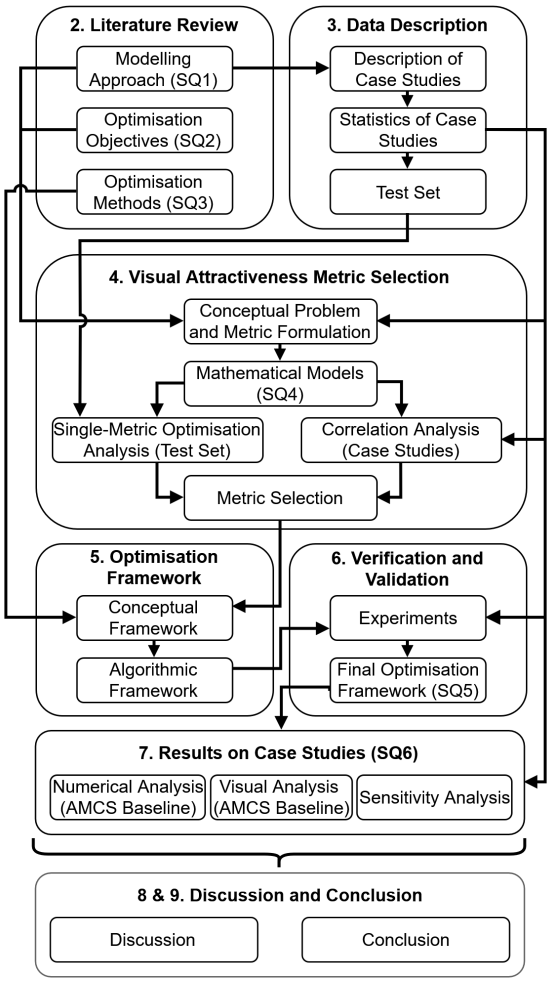

Figure 1.1: Overview of Methodology and Report Structure

2

Literature Review

Municipal solid waste collection is a routing problem in which vehicles collect waste from collection points or street segments and transport it to disposal facilities. Unlike classical distribution routing, such as the vehicle routing problem (VRP) (Dantzig & Ramser, 1959), waste collection involves reverse logistics rather than forward deliveries from a central depot to customers (Beliën et al., 2012; Han & Ponce-Cueto, 2015). This chapter reviews the modelling approaches, optimisation objectives, and optimisation methods most relevant for this study.

2.1. Modelling Approach

A key decision in waste collection routing is whether to model the problem as node- or arc-based, depending on whether waste demand is located at specific points or along street segments. This distinction reflects the differences between industrial, commercial, and residential waste collection systems (Golden et al., 2002). Industrial and commercial waste collection usually involve containers, fixed pick-up locations, and low stop density. In contrast, residential kerbside waste collection requires vehicles to drive through many short streets, stop frequently, and serve a large number of households. As a result, residential routing tends to be operationally complex (Kim et al., 2006).

Choosing a modelling approach depends on whether waste demand is assigned to specific locations (nodes) or street segments (arcs or edges). In node-based models, vehicles must visit specified locations such as containers, bins, or drop-off points. In arc-based models, they must traverse every street segment that carries demand such as in kerbside waste collection. Both approaches appear in the literature, although node-based formulations are more commonly used in practice as they align with simpler datasets and have many optimisation methods available (Beliën et al., 2012; Han & Ponce-Cueto, 2015). Additionally, both perspectives can be applied in single- and multi-route settings. Multi-route variants reflect fleet-based planning, while single-route models optimise one route at a time. In this context, node-based models correspond to the waste collection vehicle routing problem (WC-VRP), whose single-route form reduces to the travelling salesman problem (WC-TSP). Conversely, arc-based models correspond to the waste collection arc routing problem (WC-ARP), whose single-route form is the Chinese Postman Problem (WC-CPP) (Laporte, 1992).

Beyond the choice between node- and arc-based formulations in single- or multi-route settings, modelling decisions are also shaped by operational constraints from street networks and vehicle limitations. These constraints determine which movements are allowed and, for side loaders, whether streets must be traversed twice to complete service. Several studies incorporated such constraints by modifying the graph. One-way streets, turn restrictions, and undesirable left or U-turns are modelled by removing or penalising arcs (Bautista et al., 2008; Bonomo et al., 2012). Mixed-graph formulations further distinguish between different road types. For example, wide streets can be represented with parallel arcs so that each side can be serviced (Cortinhal et al., 2016). However, studies involving side loaders are rare and those that mention them usually do not explicitly model them (Angelelli & Speranza, 2002).

This study focuses on a single residential kerbside waste collection route for a side loader. Since side

loaders serve kerbsides, demand is naturally represented as directed arcs. Although an asymmetric arc routing formulation would be a natural approach, the problem is reformulated in this study as an asymmetric travelling salesman problem (ATSP). Each service arc is converted into a node, whereas the deadhead travel between two service arcs is represented by a directed edge. This choice is motivated by two reasons. First, a large body of optimisation methods exists for (A)TSP variants. Moreover, the case studies and AMCS routing tools already use node-based structures, making this formulation compatible with practical systems (see Chapter 3). A similar transformation was performed by Bautista et al. (2008) in waste collection routing.

2.2. Optimisation Objectives

In residential kerbside waste collection with side loaders, only one side of the street can be serviced per pass. Therefore, streets must often be traversed twice and vehicles may detour into distant neighbourhoods before returning to the same street. Such routes can appear incoherent and characterised by backtracking. This undermines operational clarity for both crews and residents, who may even perceive parts of the service as missed. Visual attractiveness objectives are designed to capture these issues. However, operational clarity alone is not enough, as routes must also be operationally efficient to ensure practical feasibility. Thus, travel time is incorporated as a proxy for efficiency. Finally, these objectives are not independent. Improving visual attractiveness may increase travel time, and vice versa. To address these trade-offs, multi-objective optimisation methods are reviewed to evaluate routes from both perspectives. This section discusses these three topics.

2.2.1. Visual Attractiveness

Visual attractiveness refers to the idea that route plans should not only minimise costs but also appear logical to be acceptable. The concept was introduced in VRP research by Poot et al. (2002), who found that clients often rejected cost-optimal plans when routes crossed or appeared counter-intuitive. To address this, they introduced additional evaluation criteria to capture visual attractiveness. Using a savings-based heuristic, they showed that these visual metrics could be improved alongside cost measures. Their main insight was that visually attractive plans increased acceptance among drivers and planners, even at small cost trade-offs.

The comprehensive review of Rossit et al. (2019) showed that visual attractiveness is motivated by both practical and human factors. Practical experience in parcel delivery, waste collection, and territory design indicated that compact and non-overlapping routes are easier to explain, reduce confusion, and simplify operations. In some cases, such as area-based distribution or gated communities, compactness is even required to avoid detours and preserve territorial boundaries. From a human perspective, experimental studies found that people naturally create and prefer clustered, crossing-free tours (Van Rooij et al., 2003; Vickers et al., 2006). This helps explain why operators perceive such routes as more logical, even when they are not strictly optimal in distance or time.

Visual attractiveness is studied in vehicle routing problems, where a set of routes services a collection of stops and divides the service area into territories (Rossit et al., 2019). In this context, visual attractiveness aims to ensure that routes are spatially coherent, easy to interpret, and clearly separated from each another. Therefore, the next subsections discuss the main components of visual attractiveness and the metrics used to evaluate them for multiple routes. This provides the foundation for developing adapted metrics suitable for a single side loader route, as discussed later in Chapter 4.

Components and Performance Metrics

Visual attractiveness is difficult to define because it cannot be measured directly. Instead, the literature describes it through characteristics that make routes appear logical and easy to interpret. In vehicle routing problems, three feature groups are commonly used to characterise visual attractiveness, namely compactness, non-overlapping, and low complexity. Because these components are qualitative, researchers rely on quantitative proxy metrics to assess them. Rossit et al. (2019) provided a comprehensive comparison of such metrics. Building on their study, the following subsections discuss the main components of visual attractiveness for multiple routes and the metrics most commonly used to evaluate them.

Compactness
Compactness captures the idea that each route should concentrate service activity within a limited area. In multi-route settings, stops located in the same area should be served together by a single route rather than being fragmented across multiple routes. Compact routes are easier for residents and operators to understand and reduce long detours and unnecessary travel between distant streets or neighbourhoods. According to Rossit et al. (2019), the most commonly used compactness-related metrics are geographical compactness and proximity:

- **Geographical compactness (COMP):** Measures how closely stops in a route lie to each other. There are two types of compactness metrics, which are defined for a multi-route solution consisting of a set of routes K, where each route $l \in K$ serves a set of stops T_l.
  - **Distance-to-centre:** For all routes, measures the distance of each stop to the centre of its assigned route. Smaller values indicate more compact routes around a central location. A general distance-to-centre metric is defined as the sum of distances from each stop $i \in T_l$ to the centre of its assigned route c_l:

    $$COMP = \sum_{l \in K} \sum_{i \in T_l} Dist(i, c_l)$$

    The definition of the route centre and the distance measure varies across studies. Kim et al. (2006) measured compactness as the sum of the Euclidean or Manhattan distances from each stop to their geometric centre (centroid) on all routes. When travel time or distance is used instead, the definition of the centre is the stop closest to that centroid. Poot et al. (2002) adopted a similar approach but computed the average Euclidean distance from each stop to the geometric centre instead of summing over all routes. Tang and Miller-Hooks (2006) defined the centre as the medoid, which is the stop that minimises the total distance to all others. This also allows travel time or distance to be used. Finally, Rossit et al. (2019) compared several centre definitions and concluded that the stop closest to the geometric centroid provides a good balance between accuracy and computational efficiency.
  - **Pairwise distance:** For all routes, measures the distances between all pairs of stops assigned to the same route. Unlike distance-to-centre metrics, this formulation does not require defining a route centre but is more computationally demanding. Smaller values indicate more compact routes. A general pairwise distance metric is defined as the sum of distances between all pairs of distinct stops $i, h \in T_l$ within each route:

    $$COMP = \sum_{l \in K} \sum_{i \in T_l} \sum_{\substack{h \in T_l \\ h \neq i}} \mathrm{Dist}(i, h)$$

    Poot et al. (2002) proposed this type of measure using the average pairwise Euclidean distance per route to make the metric independent of route size. Matis (2008) introduced a related measure based on the ratio of the average distance between consecutive stops and the average distance of the longest 20% consecutive distances.
- **Proximity (PROX):** For all routes, measures whether stops are assigned to the route whose centre lies closest to them. Ideally, each stop should be served by the geographically closest route. Tang and Miller-Hooks (2006) defined a proximity metric where, for each route $l \in K$, O_l' denotes the number of stops that are closer to the centre of another route than to the centre of their assigned route. They used Euclidean distances and defined route centres as medoids. Smaller values indicate more compact routes, in which stops are better allocated to routes.

  $$PROX = \sum_{l \in K} O_l^{'}$$

Non-Overlapping
In vehicle routing problems, visually attractive solutions should consist of routes that do not cross or overlap one another. Crossings or re-entries into the service territories of other routes make routes harder to follow and make it unclear which vehicle is responsible for which area. The non-overlapping component therefore aims to ensure spatial separation between routes. According to Rossit et al. (2019), the most commonly used metrics to capture this component include:

- **Convex hull overlap (CH):** Measures the degree of spatial overlap between the service areas of different routes. For each route $l \in K$, a convex hull is constructed around the locations of its stops. Stops that fall within multiple convex hulls indicate poor spatial separation. Let CH_l denote the number of such stops associated with route l. Higher values imply more spatially entangled and less distinct service territories. This metric was introduced by Poot et al. (2002) and was later applied to waste collection problems (Kim et al., 2006; Sahoo et al., 2005).

$$CH = \sum_{l \in K} CH_l$$

- **Inter-route crossings:** Counts how often the geometric paths of different routes intersect. Such intersections indicate local spatial entanglement between routes and can complicate execution and responsibilities. Unlike convex hull overlap, this metric captures local interactions rather than overall territorial overlap. Poot et al. (2002) noted that these crossings are often implemented as geometric approximations that ignore the street network. However, they are straightforward to compute with computational geometry and provide a useful proxy for visual interference between routes.

Low Complexity
Low complexity captures the idea that an individual route should have a smooth shape without irregularities or unnecessary sharp turns. Simple routes are easier for crews to execute and less prone to operational errors. Unlike compactness and non-overlapping, which are primarily defined in multi-route settings, low complexity evaluates the structure of each route independently and is therefore applicable to both multi-route and single-route problems. According to Rossit et al. (2019), the most common metrics include:

- **Bending energy (BE):** Measures the sum of turn angles along a route. For each sequence of arcs $i \to j$ and $j \to k$, let θ_{ijk} be the turning angle at node j. Higher values indicate jagged routes with many sharp turns, while lower values correspond to smoother and more natural routes.

$$BE = \sum_{(i,j,k)} \theta_{ijk}$$

- **Intra-route crossings:** Counts how often a route intersects itself. Self-intersections make a route appear illogical and harder to interpret. Poot et al. (2002) originally proposed this metric using straight-line connections between successive stops rather than actual paths according to the road network. Although this is a geometric approximation, it provides a computationally efficient proxy that penalises visually unattractive self-overlaps.

Practical Considerations

Empirical studies have indicated that many visual attractiveness metrics are strongly correlated. In particular, Rossit et al. (2019) showed that compactness and overlap measures often capture similar aspects of route structure, suggesting that a small set of metrics is sufficient to represent visual attractiveness. Therefore, they recommended combining a single compactness indicator, preferably distance-to-centre, with a single overlap-related metric. Among overlap measures, inter-route crossings were found to be computationally simpler than convex hull overlap while capturing similar forms of overlap. Moreover, convex hull overlap is closely related to compactness, whereas crossing-based measures tend to behave more independently.

An important practical distinction concerns data requirements. Most compactness and convex hull metrics are computed using stop coordinates, whereas measures such as crossings and bending energy require the actual sequence of arcs and turns along the street network. Coordinate-based approximations of crossings, as proposed by Poot et al. (2002), are computationally efficient and still provide a reasonable proxy for visual clarity. However, such approximations only partially capture operational conflicts. In contrast, bending energy cannot be approximated in this way and requires detailed path information, increasing both data requirements and computational effort.

Several waste collection studies have incorporated visual attractiveness concepts. Compactness and overlap have often been addressed through clustering-first approaches, as in Kim et al. (2006) and

Sahoo et al. (2005). They used distance-to-centre compactness and convex hull overlap as metrics. Other studies have incorporated these concepts more directly into the optimisation model. For example, Sniezek and Bodin (2006) limited geographic spread through objective function terms in arc routing, while Cortinhal et al. (2016) enforced sector connectivity to reduce overlap in sectoring models. Together, these studies demonstrate that visual attractiveness can be addressed either indirectly through clustering procedures or directly by incorporating metrics into the optimisation objective.

Many foundational studies on visual attractiveness were conducted during a period when in-vehicle navigation support was less common. Although modern navigation systems may reduce the need for routes to be explainable in a strict sense, visual attractiveness remains relevant for operational execution, planning transparency, and public acceptance. Compact and logical routes are easier to monitor, adjust, and communicate. In addition, they reduce unexpected backtracking or late revisits that may arise from operational disruptions. Although most visual attractiveness metrics were originally developed for multiple routes, their underlying principles remain relevant for a single side loader route. In this thesis, these principles are therefore used to develop adapted metrics suited to a single-route setting with double traversal, as discussed in Chapter 4.

2.2.2. Travel Time

Operational efficiency in waste collection routing is usually measured through cost, distance, or time minimisation. Among these, cost minimisation is the most commonly pursued, as it integrates various aspects such as labour, fuel, vehicle use, overtime, distance, fees, and other routing factors into monetary terms (Beliën et al., 2012; Hess et al., 2024). Travel distance and time minimisation also appear frequently, either as primary goals or as proxies for cost. Although these objectives are often correlated, they are not interchangeable. Travel distance measures physical travel, while travel time captures variable traffic conditions. Travel distance is therefore widely used to approximate costs and emissions (Bautista et al., 2008; Rambandara et al., 2022), whereas travel time provides a more direct view of routing efficiency. For instance, arterial roads may offer faster access than local streets despite longer travel distances, so the fastest route may not be the shortest in distance. It is also important to note that cost, distance, and time do not always align. For instance, Viotti et al. (2003) found that night collection reduces congestion but increases labour costs due to premium wages.

In this research, travel time is chosen as the operational efficiency measure. Operators often prioritise time because it directly determines how many collections fit within a shift. Industry cases, such as Waste Management Inc., show that travel time minimisation can substantially increase margins by enabling more collections per crew per day (Broughton, 2024). This makes travel time a more practical objective for evaluating efficiency in residential waste collection with side loaders.

2.2.3. Multi-Objective Optimisation Methods

Routes optimised for traditional objectives such as cost, distance, or travel time often differ from those optimised for visual attractiveness. Visually attractive routes tend to be longer or more expensive, while routes optimised for efficiency may appear fragmented or counter-intuitive (Rossit et al., 2019). In practice, routes that are efficient but visually confusing risk rejection, while visually attractive routes that ignore travel time may be infeasible in practice. These conflicts motivate the use of multi-objective optimisation, which explicitly accounts for competing objectives. Instead of producing a single best route, multi-objective optimisation can identify Pareto-optimal solutions, where no objective can improve without worsening another. Several methods exist to address multi-objective optimisation problems (Talbi, 2009).

The first group of multi-objective optimisation methods is scalarisation, in which multiple objectives are combined into a single objective function. The most widely used method is the weighted sum, which assigns a weight to each objective to reflect its relative importance. This approach is simple, flexible, and compatible with most optimisation methods. While it depends on the chosen weights, it allows continuous trade-offs between objectives to be explored without assuming strict priorities among objectives. An alternative scalarisation method is the ϵ-constraint approach, which optimises one objective while bounding others by predefined limits. It is well suited to exact optimisation and captures trade-offs clearly by creating a Pareto frontier. However, it requires solving multiple problems with different ϵ values.

The second group consists of the criteria-based methods, which treat the objectives separately rather than combining them. A well known example is the lexicographic or sequential method, in which objectives are ordered by priority. The most important one is optimised first and subsequent objectives are considered only once the higher-priority goals are achieved. This approach is useful when objectives differ clearly in importance, but it offers little flexibility to explore trade-offs between objectives.

Moreover, dominance-based methods compare solutions directly using Pareto dominance rather than through aggregation or prioritisation. They aim to approximate the entire Pareto frontier, identifying non-dominated solutions that represent efficient trade-offs. These are common in metaheuristics, where populations of solutions evolve toward sets of non-dominated alternatives. They provide a broad view of trade-offs in a single run but lack optimality guarantees and require more computation.

In waste collection routing, scalarisation-based approaches are commonly used in practice. For example, Kim et al. (2006) and Sahoo et al. (2005) combine travel time, compactness, and workload balance using weighted objectives within their clustering algorithm, while Poot et al. (2002) integrate distance and visual attractiveness in a savings-based heuristic. Although these studies do not construct full Pareto frontiers, they demonstrate that weighted aggregation is an effective and practical way to incorporate multiple objectives in waste collection routing problems.

In this research, travel time and visual route attractiveness are balanced without assuming strict priorities of objectives or predefined thresholds, as such preferences are not fixed in current AMCS operations. Instead, trade-offs depend on operational context and preferences of their clients. While ϵ-constraint and lexicographic methods are useful when clear service level thresholds or objective rankings exist, they are less suitable in settings where such preferences are not explicitly defined. Additionally, dominance-based methods are not used, as they are limited to population-based heuristics and require higher computational effort. Instead, a weighted sum approach is adopted as it is a simple and flexible way to explore trade-offs between objectives.

2.3. Optimisation Methods

The travelling salesman problem (TSP) is a widely studied problem in combinatorial optimisation, with many formulations and optimisation methods. Its asymmetric version (ATSP), where travel costs differ by direction, has received less attention despite strong practical relevance. Real-world routing, such as waste collection, often involves asymmetry because travel times and distances depend on direction (Cirasella et al., 2001; Odili et al., 2021). Since some classical TSP methods assume symmetry, their heuristics and metaheuristics are not always directly effective for the ATSP. For instance, the 2-opt move assumes that reversing part of a route leaves its cost unchanged, which only holds in the symmetric case. The same applies to Lin–Kernighan searches and approximation methods such as Christofides' algorithm. Optimisation methods for the ATSP are commonly discussed in three groups, namely exact methods, heuristics, and metaheuristics. The methods reviewed here primarily target distance or time minimisation, as most visual attractiveness work is studied for multiple routes rather than an individual route.

2.3.1. Exact Methods

Exact methods guarantee an optimal solution. However, they are only practical for small to medium instances, typically up to a few hundred nodes. As the ATSP is NP-hard, the computational effort of exact algorithms grows exponentially with problem size. This makes them unsuitable for larger real-world problems (Öncan et al., 2009; Roberti & Toth, 2012). Nevertheless, they are important in critical contexts that require provable optima and they often serve as benchmarks for heuristics and metaheuristics.

The ATSP can be formulated as a mixed-integer linear programming (MILP) model that minimises the total route cost, while subject to constraints preventing subtours and enforcing one incoming and one outgoing arc per node. Two well-known formulations are the DFJ model (Dantzig et al., 1954), with an exponential family of subtour elimination constraints (SECs), and the MTZ model (Miller et al., 1960), which uses order variables with a polynomial number of SECs. MTZ is easy to implement, but its relaxed solutions can deviate significantly from the true integer optimum. This forces solvers to explore more nodes through large branch-and-bound trees. DFJ produces tighter relaxations and bounds closer to

the optimum, which generally makes it preferable (Öncan et al., 2009; Roberti & Toth, 2012).

To solve these formulations, branch-and-bound and branch-and-cut are commonly applied. Branch-and-bound explores subsets of solutions using relaxations to discard infeasible or suboptimal branches of the search tree. Branch-and-cut strengthens this by dynamically adding inequality constraints such as SECs. These approaches are implemented in solvers like Gurobi, CPLEX, and Concorde, and have been applied in real-world waste collection. For example, Aransay et al. (2020) optimised municipal waste routes in Spain with Gurobi and OR-Tools, achieving time and distance reductions of approximately 30%. Benitez-Bravo et al. (2021) optimised collection routes in Mexico with a MILP in RStudio linked to GIS data, reducing distance and fuel use by 9–20%. In Argentina, Bonomo et al. (2012) modelled the routes as directed Hamiltonian paths, transformed them into symmetric TSPs, and solved them with Concorde, reducing travel distances by up to 39%.

2.3.2. Heuristics

Heuristics are problem-specific rules that quickly generate feasible tours but not necessarily optimal ones. They are valuable for larger instances where exact methods are too costly. They are often divided into construction methods, which build a route from scratch, and improvement methods, which refine an existing route through local changes. Some authors also distinguish assignment-based methods, which start from an assignment relaxation and then patch a tour. In this study, they are grouped under construction heuristics, as their main goal is to create an initial route (Cirasella et al., 2001).

Construction Heuristics

Construction heuristics build a feasible route from scratch using simple rules. They are fast and easy to implement, making them useful for large problems. However, their greedy nature usually results in suboptimal solutions. Some optimise only locally, while others rely on global relaxations. In both cases, optimality is not guaranteed. As a result, construction heuristics are rarely used alone and instead provide a starting solution for improvement heuristics or metaheuristics. The main construction heuristics for the ATSP are described in the following paragraphs (Cirasella et al., 2001).

Nearest Neighbour (NN)

Nearest neighbour (NN) is a greedy heuristic. Starting at the depot (or any node), it repeatedly visits the closest unvisited stop based on directed costs. It scales well to large instances, but because it only considers immediate distance and does not look ahead, it can easily lock into poor routes if early decisions are unlucky. This makes NN very fast but often yields suboptimal solutions. Comparative studies showed that it is outperformed by more advanced heuristics (Cirasella et al., 2001). Therefore, it is used mainly as a baseline or as an initial solution for local search. For example, Boryczka and Szwarc (2019) combined NN with greedy local search and hill climbing to establish a baseline.

Insertion heuristics

Insertion heuristics start with a small subtour, often the depot and one other node, and expand it by inserting unvisited nodes until all are included. In each step, the insertion position is chosen by the additional directed cost of replacing an arc (i, j) in the current route with two arcs (i, k) and (k, j) when inserting a new node k. The cost increase is:

$$\Delta = (c_{ik} + c_{kj}) - c_{ij}$$

The algorithm chooses the insertion position and node according to one of several rules first described by Rosenkrantz et al. (1977):

- **Cheapest insertion:** Select the node and position with the smallest added cost over all candidates. This produces relatively good routes but requires evaluating all options, making it computationally slower than simpler rules such as nearest or arbitrary insertion.
- **Farthest insertion:** Select the node farthest from the current subtour and insert it at the position that minimises the added cost. This helps avoid locally compact but globally poor routes, but the results are not always better than those of cheapest insertion. The extra step of identifying the farthest node also makes it slower than simpler rules.

- **Nearest insertion:** Select the node closest to the current subtour and insert it at the position that minimises the added cost. This is faster than cheapest insertion and is often effective, but adding distant nodes last can cause long detours and lower overall solution quality.
- **Arbitrary insertion:** Select the next unvisited node in a fixed order and insert it at the position that minimises the added cost. This is very fast and simple, but generally produces worse routes than the more selective variants.
- **Randomised insertion:** Select the next node at random, sometimes biased toward promising candidates, and insert it in the position that minimises the added cost. Although not introduced by Rosenkrantz et al. (1977), it extends arbitrary insertion (Odili et al., 2021). Randomised insertion can add diversity between runs and achieve more robust results.

Insertion heuristics have been widely applied. Benitez-Bravo et al. (2021) compared nearest, farthest, and arbitrary insertion (with edge exchange improvement) for waste collection in Mexico, reducing the route distances by 10 to 20%. Among them, nearest gave the best routes but required more computation, while farthest and arbitrary were faster but less accurate. In the ATSP literature, Brest and Žerovnik (2005) proposed randomised arbitrary insert (RAI), which repeatedly removes and reinserts nodes to diversify tours. They achieved near-optimal solutions on TSPLIB instances[1] in short runtimes. More recently, Odili et al. (2021) applied a variant that randomises node selection at each step, finding it competitive with metaheuristics such as African Buffalo Optimisation. These studies show that insertion heuristics are practical and effective starting points for further improvement.

Assignment-based heuristics

Unlike nearest-neighbour or insertion methods, assignment-based heuristics start with a global relaxation rather than building tours incrementally. They solve the linear assignment problem (AP) in polynomial time, producing a cycle cover of disjoint directed cycles. Since a Hamiltonian tour must be a single cycle, the goal is to merge them. Although more computationally demanding, they often yield stronger tours (Cirasella et al., 2001). Several variants have been developed:

- **Karp's method:** The classical approach (Karp, 1979) solves the AP once to obtain a cycle cover, then patches cycles into a Hamiltonian tour. Cycles are merged sequentially into a base cycle by deleting and reconnecting arcs to minimise added cost. It is simple and leverages the AP, but solution quality depends on the patching strategy.
- **Patched cycle cover:** A simpler variant that also starts with an AP solution but applies greedy rules, reconnecting cycles by the cheapest option until one cycle remains. It is faster and easier to implement than Karp's method, but usually produces weaker routes and is therefore used when speed matters more than accuracy (Cirasella et al., 2001).
- **Repeated assignment:** Proposed by Frieze et al. (1982), this method avoids patching by solving successive APs. After each solution, subtour constraints are added to forbid them and force cycles to break until one route remains. This method avoids heuristic patching but requires solving multiple APs, making it computationally expensive.
- **Zhang's algorithms:** Zhang (Zhang, 1993) proposed more advanced merging and replacement rules. Instead of merging two cycles at a time, they consider multiple merges and targeted arc replacements to avoid detours. In Cirasella et al. (2001), one of Zhang's variants produced the best routes among assignment-based heuristics, although at higher complexity and runtime.

All assignment-based heuristics start from the AP but differ in handling subtours, creating a trade-off between speed and quality. Patched cycle cover is fastest, while repeated assignment and Zhang's methods are slower but usually obtain better solutions. Cirasella et al. (2001) confirmed this trade-off, showing patched cycle cover as a quick baseline and the others as stronger options. They also found assignment-based heuristics generally outperform nearest neighbour and insertion, making them valuable constructive methods.

Improvement Heuristics

Unlike construction heuristics, improvement heuristics start from a feasible route and try to improve it through successive modifications. In the ATSP literature, these are local search methods, which explore the neighbourhood of the current route. A neighbourhood is the set of routes obtained by making

[1] TSPLIB is a standard library of TSP and ATSP instances derived from real-world problems. For these cases, the optimal or best-known solutions are available, making it a common benchmark for evaluating heuristics.

small changes to the current solution, such as exchanging arcs or relocating nodes. Local search iteratively applies improving moves within this neighbourhood until no further progress is possible. The resulting solution is locally optimal with respect to the defined neighbourhood structure. Improvement heuristics refine construction results or are embedded in metaheuristics to enhance performance. The main approaches for the ATSP are discussed below (Cirasella et al., 2001).

Node-Based Moves
Node-based moves directly manipulate the position of individual nodes in a tour and are among the simplest local search operators. They are usually applied through first-improvement (accepting the first improving move) or best-improvement (evaluating all moves and choosing the best) (Talbi, 2009). The most common are:

- **Relocate (reinsertion):** Removes a node from its position and reinserts it elsewhere in the route. Only the affected arcs need updating, making this operator effective for correcting misplaced nodes without recomputing the full route length. Suppose node k is moved from between i and j to between p and q, the cost change is:

$$\Delta = (c_{ij} + c_{pk} + c_{kq}) - (c_{ik} + c_{kj} + c_{pq})$$

- **Exchange (swap):** Selects and swaps two nodes, which changes their incoming and outgoing arcs. This can yield large improvements when asymmetry makes one node sequence much costlier. If k is between i and j, and m between p and q, then swapping k and m changes the cost by:

$$\Delta = (c_{im} + c_{mj} + c_{pk} + c_{kq}) - (c_{ik} + c_{kj} + c_{pm} + c_{mq})$$

- **Or-opt:** Moves a block of consecutive nodes (usually two or three) to another position in the route. This generalises the relocate move and allows larger adjustments while preserving direction. If block $(k_1, ..., k_r)$ is moved from between i and j to between p and q, the cost change is:

$$\Delta = (c_{ij} + c_{pk_1} + c_{k_r q}) - (c_{ik_1} + c_{k_r j} + c_{pq})$$

In ATSP literature, simple node-based moves are widely used as a standalone local search or within metaheuristics. The recursive arc insertion (RAI) of Buriol et al. (2004) extends relocate by repeatedly reinserting nodes in positions with the greatest cost reduction until no further improvement is possible. They showed it outperformed basic relocate in their memetic algorithm. Node swaps are also common in metaheuristics. Szwarc and Boryczka (2017) compared greedy local search and hill climbing in a memetic algorithm. Greedy local search used greedy reinsertion of one node with a first-improvement rule, whereas hill climbing used swaps with best-improvement. Both operators improved the performance of their memetic algorithm. Moreover, Boryczka and Szwarc (2019) applied greedy local search and hill climbing to routes constructed with NN, showing that reinsertion and swap moves provide improvements at low computational cost.

Node-based operators have also been applied in waste collection contexts. Karadimas et al. (2007a) used a genetic algorithm for waste collection in Athens where mutation used random swaps or subtour reversals. In both cases, the positions to modify were chosen at random from the chromosome representing the route. They concluded that swaps were more effective for route quality and population diversity. Similarly, Mekamcha et al. (2021) applied simulated annealing and tabu search to waste collection in an Algerian city, using node swaps. They found random swaps were crucial for exploring alternative routes efficiently. These studies confirm that even simple node-based operators, when embedded in metaheuristics, can deliver meaningful cost savings in real-world applications.

Edge Exchange Moves
Edge exchange operators modify a route by removing edges and reconnecting the resulting fragments differently. The idea originates from Croes (1958), who proposed 2-opt for the TSP. In this case, two edges are removed and the path is reconnected by reversing the order of the intermediate nodes. Because arc cost are symmetric ($c_{ij} = c_{ji}$), this yields feasible routes and makes 2-opt effective (Johnson & McGeoch, 1997). However, the arc costs in an ATSP are directional, so segment reversals usually produce costly routes. Instead, directed k-opt arc moves ($k \geq 3$) are often used, where three or more arcs are removed and reconnected in feasible orientations (Cirasella et al., 2001). A key example is the

4-opt bridge move (Kanellakis & Papadimitriou, 1980), which reconnects distant parts to restructure the route more globally and escape local minima. Figure 2.1 shows examples of 2-opt, 3-opt, and 4-opt bridge moves for edges instead of arcs. Smaller k-opt moves are simple to implement, but explore small neighbourhoods that risk getting stuck in local minima. Larger k-opt moves expand the search but increase the computational cost (Cirasella et al., 2001; Johnson & McGeoch, 1997).

Using k-opt moves, Kanellakis and Papadimitriou (1980) developed a directed version of the Lin–Kernighan heuristic for the ATSP. Unlike fixed k-opt methods, the K&P algorithm does not predetermine the number of arcs. Instead, it performs a variable-depth sequence of moves, extending as long as improvements are found. The process starts with an initial arc and attempts a 2-opt. If this does not reduce costs, the exchange is extended to 3-opt, 4-opt, or more. The search continues until no further improvement is found, allowing K&P to explore larger neighbourhoods than k-opt. Consequently, it consistently finds better solutions, but at higher computational costs and implementation complexity. Other experiments also confirmed that K&P outperformed simpler arc exchanges (Cirasella et al., 2001).

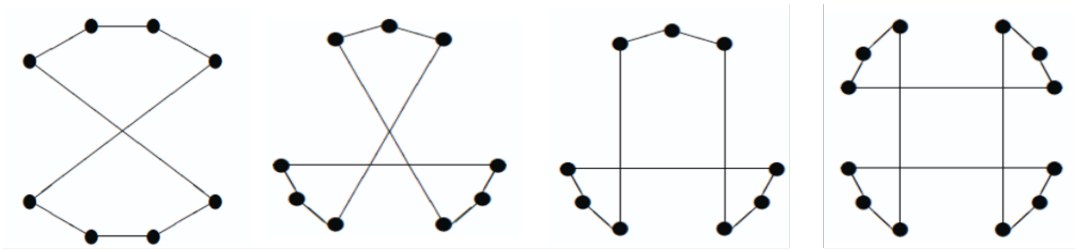

Figure 2.1: Undirected Edge Exchange Moves: 2-opt (a), 3-opt (b, c), and 4-opt bridge (d). Adapted from Johnson and McGeoch (1997).

Note: Original routes are shown as circles, where removed edges are omitted.

Arc-exchange moves are widely used in ATSP studies, often combined with efficiency techniques to keep them computationally feasible. For example, Cirasella et al. (2001) applied directed 3-opt and 4-opt moves with two speed-up techniques. The first technique, nearest-neighbour restriction, limits candidate reconnections to a small set of closest nodes, reducing neighbourhood size with little loss in solution quality. The second method, don't look-bits, temporarily skip nodes that have not produced recent improvements, preventing wasted checks in unpromising route areas. Similarly, Freisleben and Merz (1996) used fast 3-opt with don't-look-bits in their genetic algorithm, while Nagata and Soler (2012) applied a 3-opt variant with don't-look-bits on routes constructed with NN to initialise their genetic algorithm population.

2.3.3. Metaheuristics

Conventional heuristics, while computationally efficient, often rely on greedy or problem-specific rules and can become trapped in local optima. As problem size increases, they may fail to find (near-)optimal routes. Metaheuristics address these limits as general stochastic frameworks that guide heuristics to explore the solution space more effectively. Their underlying principle is balancing exploration, which diversifies the search across many possible routes, and exploitation, which refines promising routes. Metaheuristics are commonly divided into single-solution and population-based approaches (Talbi, 2009). The main metaheuristics for the ATSP are outlined in the following sections.

Single-Solution Based Metaheuristics

Single-solution based metaheuristics, also called trajectory methods, iteratively transform a single solution into a new one using a neighbourhood structure and acceptance criterion. A neighbourhood of a solution is the set of solutions reachable through simple modifications, such as swapping two nodes. At each step, the algorithm evaluates the neighbours and decides whether to accept a move. These methods intensify the search around promising regions, but risk getting stuck in local optima. Thus, diversification strategies are added to escape and explore new regions. Typical examples of this family include tabu search, simulated annealing, and variable neighbourhood search (Talbi, 2009).

Tabu Search
Tabu search (TS), proposed by Glover (1986), extends basic local search with adaptive memory by storing recently visited solutions or moves in a tabu list. This prevents the algorithm from cycling back to the same configurations and encourages exploration of new areas of the search space. It alternates between intensifying the search around promising areas and diversifying when progress stagnates. A key feature is the aspiration criterion, which allows a tabu move if it leads to a solution better than the best found so far. This ensures that promising opportunities are not blocked by short-term memory.

Simulated Annealing
Simulated annealing (SA), introduced by Kirkpatrick et al. (1983), is a fundamental single-solution metaheuristic inspired by the annealing process in metallurgy, where a material is gradually cooled to reach a stable structure. It probabilistically accepts both improving and worsening moves. Worsening moves are accepted with probability $\exp(-\Delta f/T)$, where Δf denotes the cost increase and T the temperature parameter. Since T decreases over time according to a cooling schedule, the search is more likely to accept worsening moves in early stages and gradually focuses on intensification as $T \to 0$. This mechanism enables SA to escape local minima and explore the solution space more broadly.

This balance between diversification and intensification has made SA useful in practice. Mekamcha et al. (2021) applied both SA and TS to a municipal waste collection problem modelled as an ATSP to reflect one-way streets. They optimised 21 daily truck routes with 24–122 stops. In their implementation, SA modified routes by random swaps and used Boltzmann acceptance with exponential cooling. TS explored multiple swaps, chose the best neighbour, and used a tabu list to forbid recent moves. Results showed SA achieved greater distance reductions than TS on most routes, but TS still improved routes compared to original routes.

Variable Neighbourhood Search
Variable neighbourhood search (VNS), proposed by Mladenović and Hansen (1997), is based on the idea that different neighbourhoods complement each other to escape local optima. Instead of using a single structure like SA and TS, VNS systematically changes neighbourhoods during the search. Each iteration applies a shaking step that perturbs the current solution by taking a move from one of the predefined neighbourhoods. Then a local search is applied to the perturbed solution to obtain a locally optimal neighbour. If the new solution is better, it replaces the current solution and the search restarts with the first neighbourhood. Otherwise, the algorithm moves to a larger neighbourhood. This structured approach combines diversification (shaking) with intensification (local search). Each iteration applies a shaking step that perturbs the current solution by taking a move from one of the predefined neighbourhoods.

This general strategy has also been proven to be effective for the ATSP. For instance, Burke et al. (2001) applied VNS with specialised neighbourhood operators. They implemented a hybrid local search combining a 3-opt style operator with a problem-specific heuristic. Moreover, they replaced random shaking with a guided shake that disrupted promising parts of the route to encourage exploration of new regions of the solution space. Their VNS was tested on ATSP benchmark instances from TSPLIB with up to 443 nodes. Their results showed that the guided VNS produced competitive solutions.

Population-Based Metaheuristics
Population-based metaheuristics evolve a set of candidate solutions simultaneously. Each solution in the population represents a route, and new ones are created by combining or modifying existing ones. Unlike single-solution methods, this collective approach maintains diversity while exploiting good solutions. The balance between exploration of many solutions and exploiting the best ones makes them effective for large, complex problems. Common examples in this family include genetic algorithms, ant colony optimisation, and particle swarm optimisation (Talbi, 2009).

Genetic Algorithms
Genetic algorithms (GA), introduced by Holland (1975), are metaheuristics inspired by natural evolution. A GA starts with a population of candidate solutions that evolve over generations. In each generation, selection favours low-cost solutions, using schemes such as tournament or roulette-wheel, so fitter solutions are more likely to become parents. Crossover recombines parents to form offspring, aiming to inherit good features. To avoid premature convergence, mutation introduces random changes (e.g.

node swaps) to maintain diversity and explore new regions in the search space. Finally, a replacement step creates a new generation by inserting offspring and discarding weaker solutions. Together, these steps balance exploitation of good tours with exploration of the broader solution space.

For the ATSP, several GA variants have been proposed. Freisleben and Merz (1996) combined GA with local search in genetic local search (GLS). They introduced the distance preserving crossover (DPX) that preserves common edges and refined offspring with fast 3-opt before insertion into the population. Choi et al. (2003) increased diversity by allowing crossover to produce infeasible subtours and repairing them with a patching heuristic. Their GA with roulette-wheel selection and swap mutation performed competitively on medium TSPLIB instances. Nagata and Soler (2012) developed edge assembly crossover (EAX), which preserves promising segments from both parents. Combined with tournament selection and minimal mutation, their GA outperformed earlier approaches on large TSPLIB benchmarks.

GAs have also been applied to asymmetric waste collection problems. Viotti et al. (2003) developed a GA for municipal waste collection in Rome, using rank selection, two-point crossover (i.e., take route segment of one parent and fill remaining from other), random swap mutation, and a “stranger” operator that inserts random solutions to maintain diversity. Field tests showed route improvements of up to 20%. Similarly, Karadimas et al. (2007a) used GIS-based data for Athens and evaluated selection, crossover, and two mutation operators (sub-tour reversal and node swaps). Their results indicated that swap mutation performed best, improving municipal routes by roughly 10%.

Ant Colony Optimisation

Ant colony optimisation (ACO), first proposed as the Ant System by Dorigo et al. (1996), is inspired by the foraging behaviour of ants. Artificial ants construct routes probabilistically, where the probability of choosing edge (i, j) depends on pheromone τ_{ij} and a heuristic desirability term η_{ij}, often set to the inverse distance. After routes are built, pheromone trails are updated. Shorter edges receive more pheromone, while evaporation reduces pheromone everywhere to maintain diversity and prevent premature convergence. In the next iteration, ants are more likely to follow reinforced edges, while evaporation keeps alternative paths open. Repeating this cycle gradually leads to better solutions.

The ant colony system (ACS) of Gambardella and Dorigo (1996) improves basic ACO by updating pheromone only on the best-so-far route to speed up convergence, applying local pheromone reductions to edges used to avoid stagnation, and using a pseudo-random proportional rule that balances exploiting the best edge with exploring alternatives. These changes enhanced performance on both symmetric and asymmetric TSPs. For waste collection, Karadimas et al. (2007b) applied ACS to the same Athens case as in their GA study. Compared to empirical municipal routes, ACS produced routes up to 25% shorter and reduced the average length by about 10%. These improvements exceeded those of their earlier GA study, highlighting ACS as an effective method for real-world ATSP waste collection problems.

Other Population-Based Metaheuristics

Several other population-based metaheuristics have been adapted to the ATSP. Szwarc and Boryczka (2017) systematically compared sixteen memetic algorithm (MA) variants differing in local search type (GLS, HC, TS, SA) and in the evolutionary stage where the search was applied. On random ATSP instances up to 20 cities, TS was most effective before selection. Moreover, Buriol et al. (2004) incorporated RAI in their memetic algorithm, showing it outperformed relocate. Boryczka and Szwarc (2019) adapted Harmony Search (HS) and showed that it outperformed NN and was competitive with GLS and HC on TSPLIB instances up to 170 cities. Later, Boryczka and Szwarc (2020) proposed hybrid HS variants embedding local search (TS, GLS, SA, HC), where HS with HC gave the best performance on 19 TSPLIB instances up to 443 cities. Odili et al. (2021) introduced african buffalo optimisation (ABO), inspired by herd communication. They found that it generally outperformed GA, ACO, and PSO on ATSP benchmarks, suggesting that newer bio-inspired methods can rival classical metaheuristics.

2.4. Comparative Overview

The literature on ATSP and WC-ATSP shows a clear distinction between theoretical studies and applied work on waste collection. ATSP studies typically minimise tour length, which is the sum of arc weights in a route. When weights are distances, this equals route distance, and if they represent time, it

becomes route time. Some WC-ATSP studies instead optimise costs, usually defined as a monetary or weighted function. Algorithmic papers focus on metaheuristics with enhancements such as local search or crossover, while WC-ATSP papers more often use exact models and heuristics tailored to practical constraints. Their experiments also differ. ATSP studies typically benchmark on TSPLIB instances with hundreds of nodes and extensive parameter sweeps, whereas WC-ATSP studies emphasise real-world case studies with networks from a few dozen to several hundred nodes.

Table 2.1: Overview of ATSP and WC-ATSP Literature

Reference	Objective(s)	Problem Type(s)	Method	Specifications	Experiments	Test Set
(Boryczka & Szwarc, 2019)	Tour length	ATSP	HS	Roulette node choice; greedy pitch; memory reset	Parameter sweeps; vs NN/GLS/HC	TSPLIB ATSP (n=17–443)
(Boryczka & Szwarc, 2020)	Tour length	ATSP	HS + HC	HC embedded at three insertion points	Hybridisation point study	TSPLIB ATSP (n=17–443)
(Brest & Žerovnik, 2005)	Tour length	ATSP	Construction heuristic	Randomised Arbitrary Insertion; remove/reinsert	Compared with insertion + or-opt	TSPLIB (n=17–443)
(Buriol et al., 2004)	Tour length	ATSP	MA	Recursive Arc Insertion (RAI) LS offspring	Compared with published metaheuristics	TSPLIB ATSP (n=17–443)
(Burke et al., 2001)	Route distance	ATSP	VNS & LS	Guided shaking; HyperOpt/3-opt LS	Compared VNS with LS variants	TSPLIB (n=17–443)
(Choi et al., 2003)	Tour length	ATSP	GA	Allows infeasible subtours; patch repair; swap mutation	Compared with Karp patching & B&B	TSPLIB ATSP (n=17–443)
(Dorigo et al., 1996)	Tour length	ATSP/TSP	ACO (AS)	Original ACO	Parameter tuning; vs heuristic	Small/medium ATSP/TSP
(Freisleben & Merz, 1996)	Tour length	ATSP/TSP	GA (GLS)	DPX crossover; LK/3-opt local search	Compared with GA/LS	TSPLIB (n=43–170)
(Gambardella & Dorigo, 1996)	Tour length	ATSP/TSP	ACO (ACS)	Pseudo-random rule; pheromone updates	Compared with AS/Ant-Q	TSPLIB (n=43–170)
(Kanellakis & Papadimitriou, 1980)	Tour length	ATSP	LS	Variable-depth LS; special 4-opt	—	—
(Nagata & Soler, 2012)	Tour length	ATSP	GA	EAX crossover; tournament selection; minimal mutation; 3-opt on initial population	Compared GA baselines (100 runs)	TSPLIB ATSP/TSP (n=17–443)
(Odili et al., 2021)	Tour length	ATSP	ABO/IEO/ ACO/GA	Randomised insertion (RAI)	Compared all methods	TSPLIB ATSP (n=17–443)
(Szwarc & Boryczka, 2017)	Tour length	ATSP	MA	16 MA + LS variants (GLS/HC/TS/SA)	Compared 16 variants	Random ATSP (n≤20)
(Aransay et al., 2020)	Route time	WC-ATSP	Exact	Several ATSP formulations	Compared formulations	La Rioja (n=150–300)
(Benitez-Bravo et al., 2021)	Route distance	WC-ATSP	Construction + improvement	Nearest/farthest insertion + 2-opt	Sensitivity to one-way network share	Monterrey (2 municipalities)
(Bonomo et al., 2012)	Distance & wear	WC-ATSP	Exact	Node-splitting for ATSP to TSP; min-work selection	Compared optimal tours to baseline	Buenos Aires (n=47–161)
(Das & Bhattacharyya, 2015)	Route costs	WC-ATSP	Heuristic	Three-stage routing (collection, transfer, plant; ATSP first stage)	Simulation study	Kolkata (n=50–100)
(Karadimas et al., 2007b)	Tour length	WC-ATSP	ACO (ACS)	Candidate lists; edge-exchange	27,700 ACS runs	Athens (n=72)
(Karadimas et al., 2007a)	Tour length	WC-ATSP	GA	Subtour copy crossover; swap/reverse mutation	Compared with empirical routes	Athens (n=72)
(Mekamcha et al., 2021)	Route distance	WC-ATSP	SA & TS	—	Compared SA/TS vs empirical route	Tlemcen (n=24–122)
(Viotti et al., 2003)	Route costs	WC-ATSP	GA	Minimum Length Route (MLR); rank selection; stranger operator	Compared with exact and heuristic baselines	Rome (n≥14)

2.5. Summary and Research Implications

This literature review has provided the foundation for the modelling and optimisation approach adopted in this thesis. It has identified commonly used formulations and optimisation methods for waste collection routing, as well as gaps that motivate the specific modelling choices made in this research.

Waste collection routing problems can be formulated using either node-based or arc-based network representations and in single-route or multi-route settings. Node-based formulations model demand at discrete locations, such as bins or collection points, while arc-based formulations represent demand along street segments. Multi-route problmes reflect fleet-level planning and dominate the literature, whereas single-route problems focus on sequencing the service of a single vehicle. In this thesis, a node-based single-route formulation is adopted, as the case studies consist of individual collection orders assigned

to one vehicle. Moreover, modelling kerbside collection with side loaders requires representing each street side separately, which introduces directionality and naturally leads to an asymmetric travelling salesman problem (ATSP). This reflects real operations, where travel times and feasible movements depend on direction and the road network.

Because side loaders can service only one kerbside at a time, many streets must be visited twice. This often leads to fragmented routes with long detours and revisits of streets or neighbourhoods, reducing route clarity for drivers and residents. These characteristics motivate the inclusion of visual attractiveness as an optimisation objective. While visual attractiveness has been widely studied in multi-route settings, its role in single-route kerbside collection is largely unexplored. The literature identifies three main components of visual attractiveness in multi-route problems, namely compactness, non-overlapping, and low complexity. However, these concepts and their associated metrics cannot be applied directly to a single side loader route. Consequently, Chapter 4 adapts the definitions of these components and develops metrics that capture these components for a single side loader route.

In addition to visual attractiveness, operational efficiency remains important in waste collection routing. Efficiency is commonly measured through cost, distance, or travel time. Although distance is often used as a proxy for cost and emissions, travel time provides a more direct measure of routing efficiency, as it captures traffic conditions and driving speed. For this reason, travel time is adopted as the efficiency objective in this study. However, travel time and visual attractiveness may conflict. For instance, the fastest route may appear fragmented or illogical, while a visually attractive route may require additional travel time. This trade-off motivates a multi-objective optimisation approach. Among the multi-objective methods, a weighted sum approach is adopted. This allows flexible exploration of trade-offs between objectives without assuming a strict priority order or predefined minimum or maximum service levels. Such preferences are not specified in current AMCS operations and often depend on operational context and preferences of their clients.

Finally, the literature distinguishes between exact methods, heuristics, and metaheuristics for solving routing problems. Exact methods provide optimal solutions but are computationally feasible only for small instances. In contrast, heuristics and metaheuristics are suitable for larger networks, as they deliver near-optimal solutions within reasonable computation times. In line with this distinction, exact optimisation models are applied to the small test set to analyse the behaviour of visual attractiveness metrics and support metric selection. For the full case studies, simulated annealing is selected as the solution method. As a metaheuristic, it can handle large problem instances, escape local optima through controlled acceptance of worsening moves, and operate within practical runtime limits. Moreover, simulated annealing is well suited to a weighted sum objective and can achieve good performance using simple generic neighbourhood operators, avoiding the need for highly specialised and problem-specific search rules.

3

Data Description

This chapter introduces the 12 case study datasets supplied by AMCS, covering two suburban waste collection regions. These case studies form the empirical basis for the models and analyses in this thesis. Consequently, understanding their structural and spatial characteristics is essential. The chapter proceeds in three steps. First, the main components of the case studies are described. Following this, descriptive statistics characterise their network size, composition, and spatial structure to show that the datasets are sufficiently comparable to be analysed together while retaining meaningful variation. This variation is used to study the behaviour of the visual attractiveness metrics in Chapter 4 and to evaluate the optimisation framework in Chapter 7. Finally, this chapter introduces a small test set to enable exact optimisation in Chapter 4 and verification in Chapter 6.

3.1. Description of Case Studies

In this research, two suburban residential waste collection regions, referred to as set 1 and set 2, are used. Each collection region contains several routes, that is, five in set 1 and seven in set 2, where each route forms an independent dataset. Set 1 represents a suburban area in New Zealand, where vehicles operate on the left side of the road, while set 2 corresponds to a suburban area in the United States of America, where driving occurs on the right side. In both regions, automated side loaders are used to collect household waste from the kerbside, requiring the vehicle to traverse each street segment with waste. Within each route, the orders represent individual household collection locations along the kerbside. Moreover, each route includes one or two terminals that serve as the start and end points of the route. These terminals represent depots or disposal facilities located at some distance from the residential service area. An example of the spatial layout of orders and terminals on the real road network is shown in Figure 3.1 for route 1 in New Zealand. Additional road network visualisations of other datasets are provided in Appendix B.

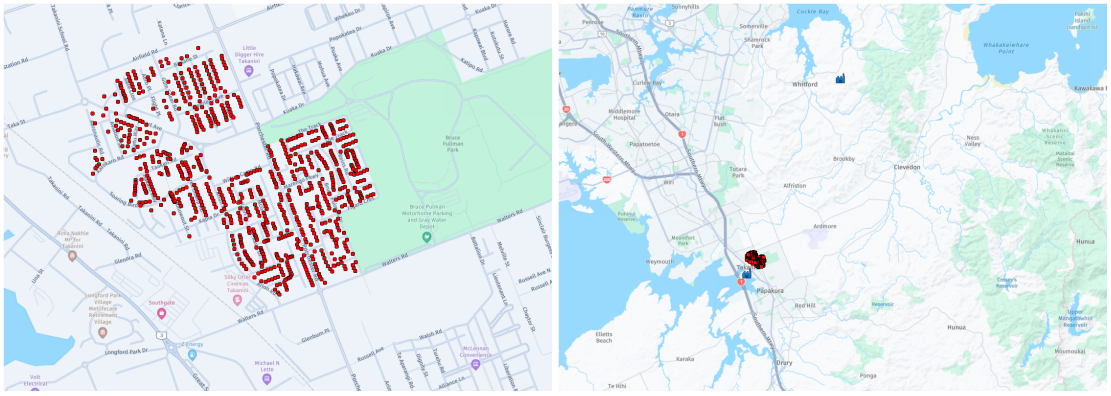

(a) Orders on the Road Network **(b)** Orders and Terminals on the Road Network

Figure 3.1: Dataset Representation on the Road Network for Route 1 of Set 1 (New Zealand)

To reduce the size of the datasets and represent the problem at a higher level of abstraction, orders are aggregated into link clusters. A link cluster represents a directed kerbside street segment between two intersections and may contain one or more orders. Each link cluster must be traversed in the permitted traffic direction, beginning at its first order node and ending at its last. Street segments requiring service on both sides are represented by two link clusters (i.e., one per kerbside direction), while street segments serviced on only one side correspond to a single link cluster.

An example of the abstract network for route 1 of set 1 is shown in Figure 3.2. Here, Figure 3.2a displays all orders to be serviced, where each point represents a collection location with coordinates in the UTM space. Moreover, Figure 3.2b illustrates the relationship between orders and link clusters. Grey dots denote individual orders, while coloured squares indicate the midpoints of the link clusters. A link cluster midpoint is computed as the average of the coordinates of the orders associated with the link cluster. The arrows indicate the directed nature of each link cluster, pointing from the first to the last order node along the street segment.

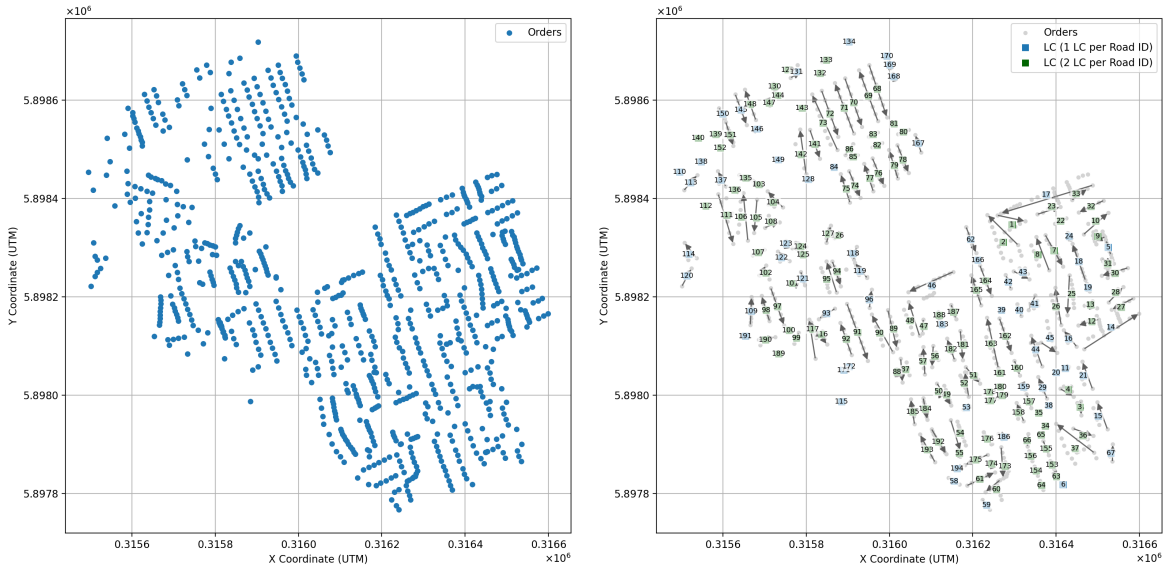

(a) Orders on the Abstract Network

(b) Link Clusters and Orders on the Abstract Network

Figure 3.2: Dataset Representation on the Abstract Network for Route 1 of Set 1 (New Zealand)

3.2. Descriptive Statistics of Case Studies

This section provides an overview of the case studies using descriptive statistics that capture both network composition and the spatial distribution of link clusters. The goal is to show that the case studies are broadly comparable in structure and spatial scale, while retaining sufficient variation to support the analyses in subsequent chapters. The structural characteristics of the network are presented first, followed by the spatial distribution patterns of the link clusters. All characteristics are summarised using the minimum, maximum, median, and interquartile range (IQR), which are preferred over the mean and standard deviation because they are less sensitive to outliers in such a small set of case studies.

3.2.1. Network Size and Composition

The scale and structural characteristics of the case studies are summarised in Table 3.1. The indicators describe the size of each network and how orders are grouped into link clusters. Reported values include the total number of orders and link clusters, the convex hull area of the link cluster midpoints, the mean and maximum number of orders per link cluster, and the ratio of single-to-dual kerbside street segments. The convex hull provides a consistent study window that excludes unused space while covering the full service region. Moreover, the single-to-dual ratio indicates the balance between street segments serviced on one or both sides, represented by one or two link clusters, respectively. Detailed values per case study dataset for all indicators are provided in Subsection B.2.1.

Across all routes, the network sizes are comparable. Each case study contains a few hundred link clusters and around a thousand collection orders. However, the convex hull areas vary substantially, ranging from compact to extensive service areas. This suggests a variation in spatial density. In addition, routes in set 2 generally cover smaller areas and contain fewer link clusters than those in set 1, indicating slightly smaller and denser service regions. In both regions, most link clusters contain only a few collection orders, with occasionally longer or more dense street segments. Each case study contains at least one link cluster with a single order. Finally, Set 1 includes slightly more street segments with two kerbsides to be serviced, whereas Set 2 contains a slightly larger share of street segments with a single kerbside. Overall, the datasets share a similar composition while still capturing meaningful variation in size, link cluster aggregation, and share of single- and dual-kerbside street segments.

Table 3.1: Network Size and Composition Statistics by Region and Overall

Set	Summary Value	Area (m^2)	Total Orders	Total LCs	Average Orders per LC	Maximum Orders per LC	Ratio Single-to-Dual
Set 1	*Minimum*	653 762	903	194	3.23	22	0.32
	Maximum	117 982 848	1815	367	5.67	41	1.99
	Median	15 815 393	1526	325	4.70	28	0.94
	IQR	110 549 190	590	43	0.15	9	0.34
Set 2	*Minimum*	872 531	933	138	4.07	17	0.78
	Maximum	3 114 554	1618	398	7.02	24	1.53
	Median	1 934 322	1070	206	4.55	20	1.09
	IQR	591 322	133	59	0.88	2.5	0.28
Overall	*Minimum*	653 762	903	138	3.23	17	0.32
	Maximum	117 982 848	1815	398	7.02	41	1.99
	Median	2 056 894	1101	253	4.68	21.5	1.03
	IQR	4 817 586	581	132	0.79	8.25	0.35

3.2.2. Spatial Distribution

In addition to the structural indicators, the spatial arrangement of each dataset was analysed using spatial distribution and point pattern statistics derived from the link cluster midpoints. These measures describe the centrality and compactness of link clusters, their global and local density, and the extent to which link clusters are regularly spaced or clustered within the service area. Summary results for both regions are presented in Table 3.2, while both detailed statistics as well as illustrative visualisations regarding centrality and local density are provided per route in Subsection B.2.2.

Centrality and Compactness

Centrality and compactness measures describe the geometric centre and spatial dispersion of each service area. Their definitions and visual illustrations are provided in Subsection B.2.2. The mean centre, geometric median, and standard deviational ellipse show that some routes form compact circular service areas, while others display elongated shapes typical of linear street networks. Compactness is quantified by the standard distance, which measures the average dispersion of link clusters around the mean centre. As shown in Table 3.2, standard distances vary considerably between routes, from a few hundred metres in the most compact networks to several kilometres in the most extensive ones. This variation aligns with the differences in area sizes. The geometric offset, defined as the relative displacement between the mean and median centres, is small for most routes. This indicates symmetric spatial distributions, with only a few cases showing elongated or uneven layouts. Overall, these measures show that most case studies share a similar spatial structure while still displaying some variation in compactness and geometric shape.

Density-Based Distribution

Density-based measures describe how intensively each service area is occupied by link clusters and how this intensity varies within a route. Global density reflects the average concentration of link clusters within the convex hull area, while local density surfaces obtained from kernel density estimation (KDE) highlight variation in link cluster concentration. Details on the KDE method and local density plots for

all routes are provided in Subsection B.2.2. As shown in Table 3.2, global densities differ by nearly two orders of magnitude. Routes in New Zealand generally show lower and more variable global densities than those in the United States, reflecting their larger and more dispersed service areas. The local density maps reveal variation between routes. Some case studies contain dense and clustered service areas, whereas others display more uniform or elongated patterns. Overall, these measures indicate comparable densities across regions while capturing differences in the spatial distribution of link clusters along routes.

Distance-Based Distribution
Distance-based metrics compare the spacing of link clusters to a benchmark of Complete Spatial Randomness (CSR). Using the nearest-neighbour statistic of Clark and Evans (1954), each route is classified along a range from clustered $(R < 1)$, through random $(R \approx 1)$, to dispersed $(R > 1)$ point patterns. The statistical significance of this deviation is evaluated with the analytical z-score. Following conventional thresholds, $|z| > 1.96$ indicates a significant difference from CSR at the 5% confidence level, with $z < -1.96$ indicating clustering and $z > 1.96$ dispersion. The methodological details are reported in Subsection B.2.2. The summaries in Table 3.2 indicate clustering of link clusters, particularly in New Zealand. In contrast, networks in the United States appear more uniformly distributed. Overall, the distance-based metrics suggest clustering across datasets, although routes vary between strong clustering, random spacing, and mild dispersion.

Table 3.2: Spatial Descriptive Statistics by Region and Overall

Set	Summary Value	Standard Distance (m)	Geometric Offset	Global Density λ (m^{-2})	NND-Ratio R	Z-Score of R
Set 1 (New Zealand)	*Minimum*	376.47	0.052	0.000003	0.59	−14.13
	Maximum	5487.70	0.243	0.000297	1.15	4.07
	Median	1555.27	0.116	0.000023	0.67	−12.10
	IQR	3973.93	0.179	0.000118	0.33	12.33
Set 2 (United States)	*Minimum*	463.84	0.011	0.000087	0.70	−11.60
	Maximum	859.12	0.304	0.000162	1.16	3.87
	Median	595.85	0.056	0.000138	1.03	0.85
	IQR	245.79	0.120	0.000035	0.13	3.66
Overall	*Minimum*	376.47	0.011	0.000003	0.59	−14.13
	Maximum	5487.70	0.304	0.000297	1.16	4.07
	Median	691.90	0.087	0.000125	0.99	−0.46
	IQR	508.76	0.149	0.000075	0.37	13.39

3.2.3. Comparison of Case Studies
Overall, the descriptive statistics show that the case studies have a consistent network composition while still exhibiting some variation in size and spatial layout to support meaningful analyses. Across both regions, the networks are constructed in a similar way. The mean number of orders per link cluster is almost identical, and the ratios of single- to dual-kerbside street segments are close to one. This indicates that the aggregation of orders into link clusters is handled consistently, which supports joint analysis of all case studies.

At the same time, some variation exists in network size and spatial distribution. Service areas differ noticeably, with New Zealand routes covering larger areas with more link clusters and orders than those in the United States. As a result, these networks are somewhat larger and sparser, although both regions still exhibit comparable global density ranges. This suggests that the main differences are driven by dataset size rather than fundamentally different spatial patterns. The spatial distribution measures confirm this, as most routes are clearly clustered, with only a few displaying slightly random spacing. In general, routes in New Zealand tend to be more clustered, whereas those in the United States are more compact and evenly distributed.

Importantly, the variation across case studies is moderate rather than extreme. Excessive variation

would undermine comparability, whereas the observed differences in service area size and spatial layout provide sufficient diversity to examine how the visual attractiveness metrics behave on travel-time-optimal routes, as discussed in Subsection 4.3.2. This balance between comparability and variation also supports the evaluation of the optimisation framework on the case studies in Chapter 6 and Chapter 7, ensuring that performance is assessed under different yet comparable conditions. Taken together, the similarity in network composition and the moderate variation in network size and spatial structure justify aggregating all 12 routes for the analyses that follow.

3.3. Test Set

To evaluate model behaviour in a controlled and interpretable setting, a small test instance is used throughout this thesis. Solving the exact optimisation models on the large case studies, is computationally expensive and often infeasible within reasonable time limits. The test set allows the exact optimisation models in Subsection 4.3.2 to be solved efficiently and enables easy inspection of how the visual attractiveness metrics influence the structure of the route. It is also used to support the verification of the optimisation framework in Section 6.1. Therefore, the purpose of this test set is not performance benchmarking, but to provide a controlled instance for understanding how the visual attractiveness metrics shape the route and for verifying optimisation behaviour.

A simple random subset of link clusters is not suitable, as random sampling disrupts the street network. This breaks spatial density patterns and results in fragmented or unrealistic networks. In addition, existing benchmark datasets, such as the ATSPLIB instances, are also not suitable because they do not capture key characteristics of residential side loader routing, including opposite kerbside representations. Instead, a spatially coherent subregion of route 1 in set 1 is selected as a test set. This case study is representative of all case studies, as shown by the descriptive statistics in Section 3.2. The selected area, shown in Figure 3.3, forms a compact neighbourhood bounded by an arterial road in the west and a park in the east. It has a consistent link cluster density and a variety of street segment types, such as short and long street segments. Importantly, entire street segments are preserved to reflect side loader operating behaviour. As discussed in Section 3.1, a street segment may contain one or two kerbsides, represented by link clusters, which must be included together.

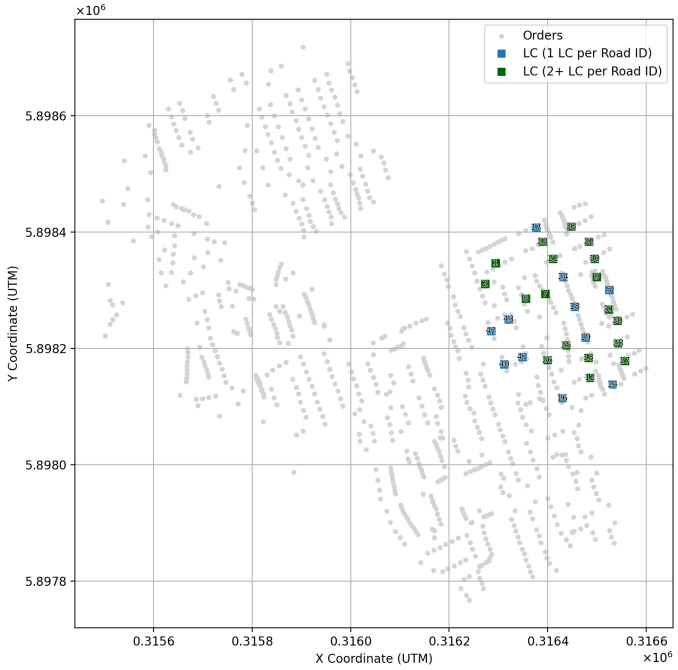

Figure 3.3: Test Set Extracted from Route 1 of Set 1 (New Zealand)

4

Visual Attractiveness Metric Selection

In this chapter, exact optimisation models are described to define, operationalise, and select visual attractiveness metrics for a single waste collection route using a side loader. The goal is to identify a set of complementary metrics that captures travel time and the most important aspects of visual attractiveness for the optimisation framework. Therefore, the chapter proceeds in two steps. First, the waste collection system and its modelling assumptions are introduced, followed by the definition of candidate visual attractiveness metrics. These concepts are then formulated in a set of single-objective Integer Linear Programming (ILP) or Mixed Integer Linear Programming (MILP) models. Each model optimises one metric while sharing the same baseline routing structure that minimises travel time.

Second, the behaviour of the metrics is analysed. A correlation analysis is performed on the travel-time-optimal routes for all 12 case studies introduced in Section 3.1. This indicates which visual attractiveness metrics behave similarly on realistic networks. Because the effects of the metrics become clear only when optimised directly, the correlation analysis is complemented by single-objective optimisation on the test set described in Section 3.3. This test set is small enough to obtain (near-)optimal solutions and to examine the effect of each metric on the route both numerically and visually. Together, these analyses support the selection of metrics for the optimisation framework in Chapter 5.

4.1. Conceptual Problem and Metric Formulation

This section describes the routing problem and introduces the candidate visual attractiveness metrics for a single side loader route. It outlines the modelling setting and the concepts behind the metrics without going into mathematical detail. This forms the basis for analysing their behaviour and selecting metrics for the optimisation framework.

4.1.1. Network Representation

The routing problem is formulated as an ATSP over the link clusters and terminals of the case studies introduced in Chapter 3. Each node in the optimisation model corresponds to a terminal or a directed link cluster, where the latter represents a directed kerbside street segment composed of one or more collection orders. The route begins at the start terminal, visits every link cluster exactly once, and ends at the end terminal. The strategic allocation of collection orders to vehicles has already been completed, accounting for capacity, demand generation, and service requirements. Consequently, the set of orders assigned to each route is fixed and capacity constraints are omitted, reducing the problem to determining an efficient and visually attractive household service sequence.

To construct the abstract network, directed arcs are defined from the start terminal to every link cluster, between every pair of link clusters, and from each link cluster to the end terminal. If a case study contains only a single terminal, the start and end terminals are modelled as two separate nodes located at the same position. This results in a network that is fully connected among link clusters, ensuring that each link cluster is reachable from any other via at least one feasible path. Travel times on the arcs are derived from the underlying street network using AMCS's fastest-path algorithm, which respects the allowed travel directions and prohibits U-turns unless unavoidable, in which case a penalty of 300

seconds is applied (e.g., at dead-end streets). For each directed pair of link clusters (c, d), the arc travel time corresponds to the fastest path from the last collection order of link cluster c to the first order of link cluster d. Consequently, the arcs in the optimisation network do not represent individual street segments. Instead, each arc aggregates potentially complex paths on the road network into a single travel time value. This abstraction preserves realistic routing constraints while keeping the optimisation tractable. However, it also prevents the use of street geometry and driving path data when computing the metrics. A simplified example network is illustrated in Figure 4.1, showing both the complete set of directed arcs and an example of a feasible route.

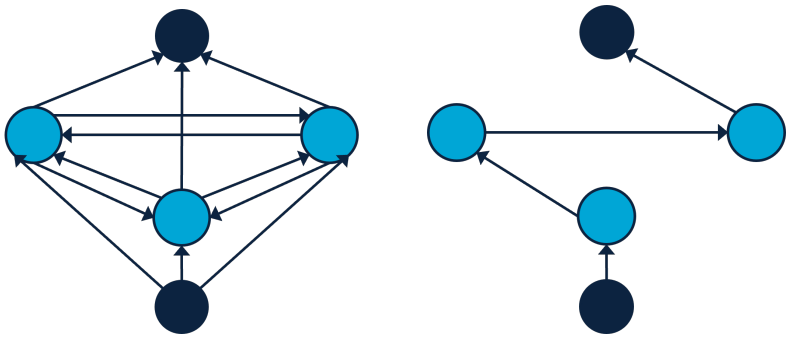

(a) All Directed Arcs in the Simplified Example Network

(b) One Feasible Route in the Simplified Example Network

Figure 4.1: Simplified Example of the Abstract Network Representation

4.1.2. Candidate Visual Attractiveness Metrics

As discussed in Subsection 2.2.1, visual attractiveness originates from VRPs and reflects how compact, non-overlapping, and easy-to-follow multiple routes appear to operators. For multiple routes, compactness keeps each route geographically cohesive, non-overlapping separates vehicle territories, and low complexity ensures smooth and intuitive shapes.

For a single side loader route, these principles translate differently. With only one vehicle, territorial separation is irrelevant and the only meaningful form of overlap is intra-route, such as revisiting areas that have already been serviced. This is captured by compactness, which expresses that nearby link clusters should be served together and that the vehicle should avoid unnecessary switching between distant or previously serviced areas. Thus, non-overlapping can be treated as part of compactness. Conversely, low complexity remains a distinct component, as a route may be spatially compact but still confusing. Therefore, visual attractiveness for a single side loader route consists of two components:

1. **Compactness:** Routes should concentrate service within a limited area before moving on, and nearby link clusters should be served together. This reduces detours and returns to previously serviced streets or areas. Thus, compactness is considered at two scales: *regional compactness* to encourage one area to be fully serviced before moving on to another, and *local compactness* to limit long moves between successive link clusters.
2. **Low complexity:** Routes should have a simple and intuitive shape without excessive zigzag movements and crossing over themselves.

These components provide different practical benefits in residential waste collection operations. Regional compactness improves operational robustness because service activity remains concentrated within a limited area. If a route cannot be completed due to unforeseen circumstances, such as vehicle breakdowns or delays, the remaining work is confined to a small region rather than scattered across multiple neighbourhoods. Both regional and local compactness also improve service predictability for residents by ensuring that nearby streets and opposite kerbsides are serviced within a short time window. When collection vehicles move through an area in a more logical and predictable pattern, residents can better anticipate when the vehicle will pass. This reduces situations in which large collection trucks unexpectedly re-enter previously serviced streets or areas. This contributes to safer neighbourhoods, particularly in suburban residential areas where large side loaders operate on streets shared with pedestrians and children. Finally, routes with low complexity are easier for drivers to interpret and execute, improving operational clarity and reducing the likelihood of navigation errors.

Although visual attractiveness is qualitative, its components can be approximated using quantitative indicators. All metrics and their required inputs are summarised in Table 4.1 and tailored to a single side loader route. Metrics are defined as sums or counts of relevant distances or events along the route, allowing them to be used directly as optimisation objectives. All metrics are evaluated on link clusters only, as terminals do not affect route coherence.

With compactness as the primary component, several metrics are defined to capture its complementary aspects of regional and local compactness. Regional compactness is captured by grouping-based metrics, including medoid spread, pairwise spread, and neighbourhood fragmentation. Medoid and pairwise spread are computed over equal-sized route segments along the visitation sequence. This allows compactness to be assessed locally along the route without defining spatial areas in advance. Therefore, the segments act as a sliding window over the route. Within each segment, dispersion is measured either relative to the medoid or through all pairwise distances. Neighbourhood fragmentation penalises re-entries into the same neighbourhood, with penalties increasing as re-entries occur later in the route. This encourages each neighbourhood to be serviced in a single contiguous block.

The top longest jumps and local jumps metrics address local compactness by directly penalising transitions between consecutive link clusters. Top longest jumps limits extreme transitions by penalising a fixed fraction of the largest distances along the route, while local jumps penalises the total distance between consecutive link clusters, allowing occasional longer moves if compensated by many short ones. This distinction is relevant when routes must traverse between distant neighbourhoods while remaining locally coherent within each area. Unlike grouping-based metrics, these objectives directly influence the visitation sequence of link clusters. Finally, revisit span captures the most local form of compactness specific to side loaders by penalising time intervals between servicing both kerbsides of the same street segment only when they exceed a specified revisit tolerance.

For medoid spread, pairwise spread, and neighbourhood fragmentation, undirected Euclidean distances between link cluster midpoints are used, as these metrics compare groups of link clusters spatially. In contrast, top longest jumps and local jumps evaluate consecutive moves along the route. Therefore, they use directed Euclidean distances between link cluster endpoints. This choice is consistent with the definition of travel time and better reflects spatial proximity along the road network. Travel time or distance is not used as a distance proxy for these metrics, as restricted U-turns can cause nearby link clusters to appear far apart in travel time or distance. For example, two link clusters on opposite kerbsides may be close in Euclidean distance but far in travel time or distance.

For the low complexity component, metrics requiring trajectory information, such as bending energy, cannot be applied because data on driving paths and road geometry are unavailable. Consequently, only intra-route crossings are included. This metric approximates route complexity by detecting geometric intersections between straight lines connecting consecutive link clusters by their endpoints. As noted by Poot et al. (2002), this simplification provides a computationally efficient approximation of visual route complexity without requiring detailed road geometry data. Although only one low complexity metric is defined, some aspects of complexity are indirectly captured by the travel time matrix, where prohibited U-turns already limit zigzagging.

Table 4.1: Visual Attractiveness Metrics for a Single Route using a Side Loader

Metric	Description and Interpretation	Additional Inputs
Compactness		
Medoid spread	Measures compactness of link clusters within each route segment by summing their distances to the segment medoid (adapted from Tang and Miller-Hooks (2006)). Segments are defined as equal-sized consecutive blocks along the visitation order. The total metric value is obtained by summing the segment spreads over all segments. Lower values indicate stronger compactness within segments.	Set of route segments

Continued on next page

Table 4.1: Visual Attractiveness Metrics for the Single-Route Setting *(continued)*

Metric	Description and Interpretation	Additional Inputs
Pairwise spread	Measures compactness of link clusters within each route segment by summing all pairwise distances between link clusters in the segment (adapted from Poot et al. (2002)). Segments are defined as equal-sized consecutive blocks along the visitation order. The total metric value is obtained by summing the segment spreads over all segments. Lower values indicate stronger compactness within segments.	Set of route segments
Neighbourhood fragmentation	Measures fragmentation of neighbourhood service by penalising re-entries into the same neighbourhood. Neighbourhoods serviced in a single contiguous block incur no penalty, while re-entries are penalised more heavily the further apart they occur along the route. The total metric value is obtained by summing these penalties over all neighbourhoods. Lower values indicate stronger compactness of neighbourhoods.	Set of neighbourhoods; assignment of link clusters to neighbourhoods
Top longest jumps	Measures the sum of the longest fraction of directed Euclidean distances between consecutive link clusters along the route (adapted from Matis (2008)). Lower values indicate stronger local compactness through limited long-distance moves.	Fraction of longest consecutive distances
Local jumps	Measures the total directed Euclidean distance between consecutive link clusters along the route. Lower values indicate stronger local compactness through shorter average moves.	–
Revisit span	Measures the excess time between servicing both kerbsides of the same street segment beyond a specified revisit tolerance, summed over all revisit pairs. Lower values indicate that opposite kerbsides are serviced within a shorter time window.	Set of revisit pairs; revisit span tolerance
Low complexity		
Intra-route crossings	Counts the total number of geometric intersections between route segments, detected using straight-line connections between link cluster endpoints. Lower values indicate a simpler and more visually coherent route shape.	Set of directed link cluster arc pairs with a geometric intersection

Note. Medoid spread, pairwise spread, and neighbourhood fragmentation use undirected Euclidean distances between link cluster midpoints. Top longest jumps and local jumps use directed Euclidean distances between the end order node of one link cluster to the start order node of another link cluster.

4.2. Mathematical Optimisation Models

This section presents the mathematical formulation of the routing problem for a side loader route. It first introduces the notation for all models, followed by the objective functions and constraints. The baseline model minimises total travel time and provides the routing structure for all visual attractiveness model variants. Each model builds on this baseline by modifying the objective function and, where necessary, introducing additional sets, parameters, variables, or constraints. All models are constructed from the metrics introduced in Section 4.1 and do not use elements from other visual attractiveness models.

4.2.1. Notations and Definitions

The sets, parameters, and decision variables for all models are given in Table 4.2 and described below.

Baseline Travel Time (TT) Model

The residential waste collection area is modelled as a directed graph $G = (V, A)$, where V denotes the set of nodes (terminals and link clusters), and A the set of feasible directed arcs. Each arc $(i, j) \in A$ has a precomputed travel time tt_{ij} representing the travel time from the end of link cluster i to the start of link cluster j. Every link cluster $c \in C$ has a fixed service time st_c to collect all orders, while terminals have no service time $st_{z_s} = st_{z_e} = 0$. Since all nodes are visited exactly once, the total service time is constant and does not affect optimisation. The routing structure is defined by binary variables x_{ij}, which equals 1 if node j is visited immediately after node i. Subtour elimination is enforced using integer order variables u_i following the MTZ formulation.

Medoid Spread (MS) Model

The medoid spread model penalises spatial dispersion within route segments to promote regional compactness. The route is divided into segments indexed by $l \in L$, defined as equal-sized blocks along the visitation order with bounds (α_l, β_l). Segment sizes differ by at most one cluster and together cover all link clusters. A binary variable y_{cl} indicates whether link cluster c belongs to route segment l. Within

each segment l, one link cluster p is selected as the medoid, defined as the link cluster with the smallest total distance to all others in that segment. This is modelled through a binary variable m_{pl}. Moreover, every link cluster c assigned to segment l is then linked to its segment medoid through assignment variables a_{cpl}. To ensure consistency between segment membership and route order, the model links y_{cl} to the MTZ variables u_c using big-M constraints (see Subsection 4.2.3). By grouping nearby link clusters around a representative medoid in route segments, the model encourages compact route segments.

Pairwise Spread (PS) Model
As in the medoid spread model, the pairwise spread model penalises spatial dispersion within route segments to promote regional compactness. Similarly, the route is divided into segments indexed by $l \in L$ defined by visitation order bounds (α_l, β_l), with segment membership indicated by variables y_{cl}. What differs is how dispersion is measured. Instead of using a representative medoid, dispersion is measured by summing the distances between all unordered pairs of link clusters within each segment. For every segment l and each unordered pair (c, d) with $c < d$, a binary variable p_{cdl} equals 1 if both link clusters belong to segment l. Using unordered pairs avoids double-counting distances, since the Euclidean distance between two link cluster midpoints is symmetric ($d_{cd} = d_{dc}$). Segment membership is again linked to the route order using big-M constraints (see Subsection 4.2.3). By minimising pairwise distances in route segments, the model encourages compact route segments.

Neighbourhood Fragmentation (NF) Model
The neighbourhood fragmentation model penalises re-entries into the same neighbourhood to promote regional compactness. As neighbourhood data are unavailable, neighbourhoods are either specified manually or constructed using a k-means clustering procedure (see Appendix C). Although this provides only approximate neighbourhoods, it offers a simple and reproducible division of the network. Each link cluster $c \in C$ is assigned a neighbourhood label $\nu(c) \in H$, where H denotes the set of neighbourhoods. Based on this assignment, a binary parameter δ_{cd} is defined, which equals 1 if link clusters c and d belong to different neighbourhoods and 0 otherwise. Rather than penalising all arcs that switch neighbourhoods, fragmentation is captured at the neighbourhood level by propagating a neighbourhood transition index q_c along the route. This index increases when the route switches neighbourhoods. For arcs not selected in the route, the propagation constraints are relaxed using a big-M constraints (see Subsection 4.2.3). For each neighbourhood $h \in H$, variables F_h and L_h represent the minimum and maximum values of the transition index among all link clusters assigned to a neighbourhood h. If a neighbourhood is served in a single contiguous block, all its link clusters share the same transition index, yielding $L_h - F_h = 0$. If the route leaves and later re-enters the same neighbourhood, this span increases. The objective minimises the total span across all neighbourhoods, discouraging later re-entries while allowing each neighbourhood to be entered once without penalty.

Top Longest Jumps (TJ) Model
The top longest jumps model penalises the longest moves between successive link clusters in the route. Let dd_{cd} denote the directed Euclidean distance from the end order node of link cluster c to the start order node of link cluster d. Given a fraction $f \in (0, 1)$, the number of penalised jumps is defined as $k^{\text{top}} = \lceil f(|C|-1) \rceil$. To identify the k^{top} longest successive distances in a linear formulation, a threshold variable τ and non-negative slack variables s_{cd} are introduced. At optimality, the objective penalises the k^{top} largest values. This encourages routes that limit unnecessary large spatial jumps and improve local compactness.

Local Jumps (LJ) Model
The local jumps model penalises long moves between successive link clusters in the route. The metric penalises the total directed Euclidean distance between successive link clusters along the route, encouraging shorter average movements between successive link clusters. Unlike the top longest jumps model, this formulation does not require additional parameters and penalises total route length rather than focusing on extreme values to allow longer jumps when beneficial.

Revisit Span (RS) Model
The revisit span model penalises long time intervals between servicing both kerbsides of the same street segment. Let $R = \{(c, d) : c < d\}$ denote the set of unordered pairs of link clusters belonging to the same street segment. Since each link cluster represents one kerbside segment between two intersections, the metric is evaluated for street segments rather than across entire streets. For each

pair $(c, d) \in R$, a continuous variable r_{cd} captures the excess revisit span beyond an allowed tolerance Δ_{RS}, where the revisit span is defined as the absolute difference in arrival times between the two link clusters. This formulation penalises only long undesirable delays between servicing opposite kerbsides, while allowing short acceptable time intervals.

To compute revisit spans, arrival time variables t_i are introduced for all nodes $i \in V$ and propagate along the route based on travel and service times. Arrival times are constrained to lie within precomputed bounds $[E_i, L_i]$, derived from the baseline travel time model. For each link cluster c, the earliest arrival time is $E_c = tt_{z_s c}$, while the latest feasible arrival time is $L_c = T^{\text{path}}_{\max} - (st_c + tt_{cz_e})$, where $T^{\text{path}}_{\max}$ is the maximum feasible total travel time obtained by maximising the baseline model. For the start terminal, arrival times are fixed to zero ($Ez_s = L_{z_s} = 0$), and for the end terminal they are bounded by the minimum and maximum feasible travel times (($E_{z_e} = T^{\text{path}}_{\min}, L_{z_e} = T^{\text{path}}_{\max}$). Here, $T^{\text{path}}_{\min}$ is the minimum feasible total travel time obtained by minimising the baseline model. These bounds enable the use of arc-specific big-M coefficients M^{LB}_{ij} and M^{UB}_{ij} in the arrival time constraints (see Subsection 4.2.3), ensuring tight linear relaxations.

Intra-Route Crossings (IC) Model

The intra-route crossings model measures how often the route geometrically crosses itself. Crossings are detected by comparing all route segments between link clusters, represented as straight lines from the end order node of one link cluster to the start order node of the next. Two lines intersect if they cross within their interiors, rather than touching at endpoints. To test for intersections, the point orientation test from computational geometry is applied (Berg et al., 2008; Goodrich & Tamassia, 2014). For three points a, b, and c, the orientation test determines whether the point c lies left, right, or exactly on the directed line from a to b by computing the signed area of the triangle formed by the points:

$$\text{orient}(a, b, c) = \text{sign}\big((b_x - a_x)(c_y - a_y) - (b_y - a_y)(c_x - a_x)\big)$$

A positive value indicates that c lies left of the line from a to b, a negative value indicates that it lies right, and zero means the points are collinear. Thus, two link cluster lines $\overline{p_c p_d}$ and $\overline{p_p p_q}$ intersect properly if the endpoint of each line lies on opposite sides of the line supporting the other:

$$\text{orient}(p_c, p_d, p_p)\,\text{orient}(p_c, p_d, p_q) < 0 \quad \text{and} \quad \text{orient}(p_p, p_q, p_c)\,\text{orient}(p_p, p_q, p_d) < 0$$

The first condition ensures points p_p and p_q lie on opposite sides of the line through p_c and p_d, while the second ensures that p_c and p_d are on opposite sides of the line through p_p and p_q. Applying this test to all pairs of directed link cluster arcs $((c, d), (p, q))$ with $(c, d) \neq (p, q)$, excluding arcs with terminals, yields the set of potential crossings $\mathcal{X}$. Although more efficient sweep-line algorithms exist (Berg et al., 2008; Goodrich & Tamassia, 2014), a brute-force comparison is used here for simplicity. In the optimisation model, crossings are penalised only when both corresponding arcs are selected in the route, ensuring that the objective accounts exclusively for geometric self-intersections of the route. Routes with fewer intra-route crossings are visually smoother and easier to interpret.

Table 4.2: Mathematical Notation

Symbol	Description	Domain	Unit	Models[a]
Sets and indices				
C	Set of link cluster nodes ($\|C\| = m$)	$c \in C$	–	All
Z	Set of start and end terminal nodes ($Z = \{z_s, z_e\}$)	$z \in Z$	–	All
V	Set of all nodes ($V = C \cup Z;\ \|V\| = n$)	$i \in V$	–	All
A	Set of directed arcs ($A = \{(z_s, d) : d \in C\} \cup \{(c, d) : c, d \in C, c \neq d\} \cup \{(c, z_e) : c \in C\}$)	$(i, j) \in A$	–	All
L	Index set of route segments	$l \in L$	–	MS, PS
H	Set of neighbourhoods	$h \in H$	–	NF
R	Set of unordered link cluster pairs on the same street segment	$(c, d) \in R$	–	RS

Continued on next page

Table 4.2: Mathematical Notation *(Continued)*

Symbol	Description	Domain	Unit	Models[a]
$\mathcal{X}$	Set of pairs of directed link cluster arcs whose endpoint-based straight lines intersect	$\big((c,d),(p,q)\big) \in \mathcal{X}$	–	IC
Parameters				
tt_{ij}	Travel time from node i to node j	$tt_{ij} \in \mathbb{R}_{\geq 0}$	$[s]$	TT, RS
st_i	Service time for node i	$st_i \in \mathbb{R}_{\geq 0}$	$[s]$	RS
d_{cd}	Undirected Euclidean distance between link cluster midpoints c and d	$d_{cd} \in \mathbb{R}_{\geq 0}$	$[m]$	MS, PS, NF
dd_{cd}	Directed Euclidean distance from the end order node of link cluster c to the start order node of link cluster d	$dd_{cd} \in \mathbb{R}_{\geq 0}$	$[m]$	LJ, TJ
α_l	Start position of route segment l in visit order	$\alpha_l \in \{1,\ldots,\lvert C\rvert\}$	–	MS, PS
β_l	End position of route segment l in visit order	$\beta_l \in \{1,\ldots,\lvert C\rvert\}$	–	MS, PS
$\nu(c)$	Neighbourhood label of link cluster c	$\nu(c) \in H$	–	NF
δ_{cd}	Neighbourhood transition indicator, equals 1 if $\nu(c) \neq \nu(d)$; 0 otherwise	$\delta_{cd} \in \{0,1\}$	–	NF
k^{top}	Number of link cluster arcs that are considered longest	$k^{top} \in \mathbb{Z}_{\geq 0}$	–	TJ
E_i	Earliest feasible arrival time at node i	$E_i \in \mathbb{R}_{\geq 0}$	$[s]$	RS
L_i	Latest feasible arrival time at node i	$L_i \in \mathbb{R}_{\geq 0}$	$[s]$	RS
M_{ij}^{LB}	Arc-specific big-M for lower bound time propagation	$M_{ij}^{\mathsf{LB}} \in \mathbb{R}_{\geq 0}$	$[s]$	RS
M_{ij}^{UB}	Arc-specific big-M for upper bound time propagation	$M_{ij}^{\mathsf{UB}} \in \mathbb{R}_{\geq 0}$	$[s]$	RS
Δ_{RS}	Revisit span tolerance	$\Delta_{RS} \in \mathbb{R}_{\geq 0}$	$[s]$	RS
Decision variables				
x_{ij}	Equals 1 if node j is visited immediately after node i; 0 otherwise	$x_{ij} \in \{0,1\}$	–	All
u_i	Order position of node i (MTZ formulation)	$u_i \in \{0,\ldots,n-1\}$	–	All
y_{cl}	Equals 1 if link cluster c is in route segment l; 0 otherwise	$y_{cl} \in \{0,1\}$	–	MS, PS
m_{pl}	Equals 1 if link cluster p is the medoid of route segment l; 0 otherwise	$m_{pl} \in \{0,1\}$	–	MS
a_{cpl}	Equals 1 if link cluster c is assigned to medoid p in route segment l; 0 otherwise	$a_{cpl} \in \{0,1\}$	–	MS
p_{cdl}	Equals 1 if link clusters c and d are both in route segment l ($c < d$); 0 otherwise	$p_{cdl} \in \{0,1\}$	–	PS
q_c	Neighbourhood transition index for link cluster c	$q_c \in \{0,\ldots,\lvert C\rvert\}$	–	NF
F_h	First neighbourhood transition index for neighbourhood h	$F_h \in \{0,\ldots,\lvert C\rvert\}$	–	NF
L_h	Last neighbourhood transition index for neighbourhood h	$L_h \in \{0,\ldots,\lvert C\rvert\}$	–	NF
τ	Distance threshold for longest jumps	$\tau \in \mathbb{R}_{\geq 0}$	$[m]$	TJ
s_{cd}	Excess distance of link cluster arc (c,d) in route above threshold τ	$s_{cd} \in \mathbb{R}_{\geq 0}$	$[m]$	TJ
t_i	Arrival time at node i	$t_i \in [E_i, L_i]$	$[s]$	RS
r_{cd}	Excess revisit span between two link clusters c and d on the same street segment above revisit tolerance	$r_{cd} \in \mathbb{R}_{\geq 0}$	$[s]$	RS
$\chi_{cd,pq}$	Equals 1 if link cluster arcs (c,d) and (p,q) are used and their endpoint-based straight lines intersect; 0 otherwise	$\chi_{cd,pq} \in \{0,1\}$	–	IC

Note. $\mathbb{R}_{\geq 0}$ denotes the set of non-negative real numbers, $\mathbb{Z}_{\geq 0}$ denotes the set of non-negative integers, $[a,b]$ denotes the closed interval of real numbers from a to b, and $\{a,\ldots,b\}$ denotes the discrete set of consecutive integers from a to b.
[a] TT: travel time; MS: medoid spread; PS: pairwise spread; NF: neighbourhood fragmentation; LJ: local jumps; TJ: top longest jumps; RS: revisit span; IC: intra-route crossings.

4.2.2. Optimisation Objectives

The objective function minimises one of the optimisation objectives below, depending on the chosen model. Travel time captures operational efficiency, while the remaining objectives each represent a

distinct aspect of visual attractiveness. The objectives are as follows:

$$\text{Travel Time (TT):} \quad \min \sum_{d \in C} tt_{z_s d}\, x_{z_s d} + \sum_{c \in C} \sum_{\substack{d \in C \\ d \neq c}} tt_{cd}\, x_{cd} + \sum_{c \in C} tt_{cz_e}\, x_{cz_e} \tag{4.1}$$

$$\text{Medoid Spread (MS):} \quad \min \sum_{l \in L} \sum_{c \in C} \sum_{\substack{p \in C \\ p \neq c}} d_{cp}\, a_{cpl} \tag{4.2}$$

$$\text{Pairwise Spread (PS):} \quad \min \sum_{l \in L} \sum_{c \in C} \sum_{\substack{d \in C \\ c < d}} d_{cd}\, p_{cdl} \tag{4.3}$$

$$\text{Neighbourhood Fragmentation (NF):} \quad \min \sum_{h \in H} (L_h - F_h) \tag{4.4}$$

$$\text{Top Longest Jumps (TJ):} \quad \min k^{top}\, \tau + \sum_{c \in C} \sum_{\substack{d \in C \\ d \neq c}} s_{cd} \tag{4.5}$$

$$\text{Local Jumps (LJ):} \quad \min \sum_{c \in C} \sum_{\substack{d \in C \\ d \neq c}} dd_{cd}\, x_{cd} \tag{4.6}$$

$$\text{Revisit Span (RS):} \quad \min \sum_{(c,d) \in R} r_{cd} \tag{4.7}$$

$$\text{Intra-Route Crossings (IC):} \quad \min \sum_{\big((c,d),(p,q)\big) \in \mathcal{X}} \chi_{cd,pq} \tag{4.8}$$

The first objective (4.1) minimises the total travel time of the route. The first term covers travel from the start terminal to the link clusters, the second term represents travel between link clusters, and the last term covers travel from the link clusters to the end terminal. Following this, objectives (4.2) and (4.3) measure compactness within route segments. Across all route segments, the medoid spread objective minimises the distance from each link cluster to their segment medoid, while the pairwise spread objective minimises all pairwise distances between link clusters within the segment. Additionally, objective (4.4) penalises neighbourhood fragmentation by minimising the span between the first and last occurrences of each neighbourhood along the route. Local compactness is addressed by objectives (4.5) and (4.6). The top longest jumps objective limits long moves by minimising the k^{top} largest directed distances along the route. This is achieved through a threshold variable τ, where the term $k^{\text{top}}\tau$ penalises the threshold itself and $\sum s_{cd}$ penalises distances exceeding it, causing τ to adjust to the k^{top} longest arcs. In contrast, the local jumps objective minimises the total directed Euclidean distance between consecutive link clusters, penalising long transitions on average. Additionally, objective (4.7) minimises the total excess revisit time interval of all revisit pairs beyond the predefined revisit tolerance. Finally, objective (4.8) minimises the total number of geometric self-intersections of the route.

4.2.3. Constraints

This subsection defines the constraints that determine the feasible region of each optimisation model. The constraints are grouped by their modelling purpose. The routing and subtour elimination constraints apply to all models, as they define the basic path structure of the route. The remaining constraints are specific to the corresponding model and are included only when that model is solved. Each group of constraints is introduced conceptually before its mathematical formulation is given.

Baseline Constraints

The baseline constraints, used in all model variants, consist of routing and subtour elimination constraints. The routing constraints ensure that the route forms a single continuous path from the start terminal z_s to the end terminal z_e, visiting each link cluster exactly once. Constraints (4.9) and (4.10) impose the degree conditions for link clusters, requiring exactly one incoming and one outgoing arc. Constraints (4.11) and (4.12) enforce the degrees for the start and end terminals. Together, these

constraints guarantee route continuity and full coverage of all link clusters.

$$\sum_{\substack{j \in V \\ j \neq c, j \neq z_s}} x_{cj} = 1 \qquad \forall c \in C \tag{4.9}$$

$$\sum_{\substack{i \in V \\ i \neq c, i \neq z_e}} x_{id} = 1 \qquad \forall d \in C \tag{4.10}$$

$$\sum_{d \in C} x_{z_s d} = 1 \tag{4.11}$$

$$\sum_{c \in C} x_{c z_e} = 1 \tag{4.12}$$

The subtour elimination constraints use the MTZ formulation to prevent disconnected cycles by assigning each node $i \in V$ an order position u_i. Constraint (4.13) and (4.14) fix the order position of the start and end terminal, reducing symmetry and possibly improving solver performance. Moreover, constraint (4.15) enforces consistent ordering between visited nodes so that the solution forms a single path rather than multiple disconnected subtours.

$$u_{z_s} = 0 \tag{4.13}$$

$$u_{z_e} = n - 1 \tag{4.14}$$

$$u_i - u_j + n\, x_{ij} \leq n - 1 \qquad \forall i, j \in V \setminus \{z_s\},\ i \neq j \tag{4.15}$$

Medoid Spread Constraints

The medoid spread constraints ensure that every route segment $l \in L$ contains exactly one medoid and that all link clusters assigned to that segment are linked to it. Constraint (4.16) assigns each link cluster to exactly one route segment. Additionally, constraint (4.17) selects exactly one medoid per route segment, while constraint (4.18) restricts medoid selection to link clusters belonging to that segment. Moreover, constraints (4.19)–(4.21) define the assignment variables a_{cpl}, ensuring that a link cluster c can be assigned to medoid p only if both belong to segment l, and that the assignment is active if both y_{cl} and m_{pl} equal one. To ensure that segment membership reflects the visitation order, constraints (4.22)–(4.23) link the MTZ order variables u_c to segment membership variables y_{cl}. If a link cluster belongs to segment l, its position in the route must lie within the predefined bounds (α_l, β_l). These constraints are relaxed through a big-M constant $m = |C|$, which deactivates them when a link cluster is not assigned to the route segment ($y_{cl} = 0$).

$$\sum_{l \in L} y_{cl} = 1 \qquad \forall c \in C \tag{4.16}$$

$$\sum_{p \in C} m_{pl} = 1 \qquad \forall l \in L \tag{4.17}$$

$$m_{pl} \leq y_{pl} \qquad \forall p \in C,\ l \in L \tag{4.18}$$

$$a_{cpl} \leq y_{cl} \qquad \forall c, p \in C,\ c \neq p,\ l \in L \tag{4.19}$$

$$a_{cpl} \leq m_{pl} \qquad \forall c, p \in C,\ c \neq p,\ l \in L \tag{4.20}$$

$$a_{cpl} \geq y_{cl} + m_{pl} - 1 \qquad \forall c, p \in C,\ c \neq p,\ l \in L \tag{4.21}$$

$$u_c \leq \beta_l + m(1 - y_{cl}) \qquad \forall c \in C,\ l \in L \tag{4.22}$$

$$u_c \geq \alpha_l - m(1 - y_{cl}) \qquad \forall c \in C,\ l \in L \tag{4.23}$$

Pairwise Spread Constraints

The pairwise spread constraints identify which link clusters are jointly included within each route segment $l \in L$. As in the medoid spread model, each link cluster is assigned to exactly one segment. Constraint (4.24) enforces this partition of the route through the membership variables y_{cl}. For each segment l and unordered pair (c, d) with $c < d$, constraints (4.25)–(4.27) ensure that a pair (c, d) is active in segment l if both clusters belong to that segment, as indicated by $p_{cdl} = 1$. This avoids double counting while correctly capturing pairwise dispersion within segments. To ensure that segment membership reflects the visitation order, constraints (4.28)–(4.29) link the MTZ order variables u_c to segment

membership variables y_{cl}. If a link cluster belongs to segment l, its position in the route must lie within the predefined bounds (α_l, β_l). These constraints are relaxed through a big-M constant $m = |C|$, which deactivates them when a link cluster is not assigned to the route segment ($y_{cl} = 0$).

$$\sum_{l \in L} y_{cl} = 1 \qquad \forall c \in C \tag{4.24}$$

$$p_{cdl} \leq y_{cl} \qquad \forall l \in L,\ c, d \in C,\ c < d \tag{4.25}$$

$$p_{cdl} \leq y_{dl} \qquad \forall l \in L,\ c, d \in C,\ c < d \tag{4.26}$$

$$p_{cdl} \geq y_{cl} + y_{dl} - 1 \qquad \forall l \in L,\ c, d \in C,\ c < d \tag{4.27}$$

$$u_c \leq \beta_l + m(1 - y_{cl}) \qquad \forall c \in C,\ l \in L \tag{4.28}$$

$$u_c \geq \alpha_l - m(1 - y_{cl}) \qquad \forall c \in C,\ l \in L \tag{4.29}$$

Neighbourhood Fragmentation Constraints

The neighbourhood fragmentation constraints capture whether the route leaves and later re-enters the same neighbourhood. Rather than penalising every neighbourhood switch directly, the formulation tracks neighbourhood changes along the route and measures how fragmented each neighbourhood is. An integer transition index q_c is introduced for each link cluster $c \in C$. Constraints (4.30)–(4.31) propagate this index along selected route arcs. When an arc (c, d) is used and switches neighbourhoods ($\delta_{cd} = 1$), the transition index increases by one. For arcs not selected in the route, the constraints are relaxed through a big-M constant $m = |C|$.

For each neighbourhood $h \in H$, variables F_h and L_h denote the minimum and maximum transition index values among its link clusters, enforced by constraints (4.32)–(4.33). If a neighbourhood is served in a single contiguous block, all link clusters share the same transition index and the difference $L_h - F_h$ equals zero. If the route leaves and re-enters the neighbourhood, this span increases. Because the transition index increases cumulatively along the route, later re-entries incur higher penalties than short or immediate re-entries.

$$q_d \geq q_c + \delta_{cd} - m(1 - x_{cd}) \qquad \forall c, d \in C,\ c \neq d \tag{4.30}$$

$$q_d \leq q_c + \delta_{cd} + m(1 - x_{cd}) \qquad \forall c, d \in C,\ c \neq d \tag{4.31}$$

$$F_h \leq q_c \qquad \forall h \in H,\ c \in C : \nu(c) = h \tag{4.32}$$

$$L_h \geq q_c \qquad \forall h \in H,\ c \in C : \nu(c) = h \tag{4.33}$$

Top Longest Jumps Constraints

These constraints introduce auxiliary variables s_{cd} and a threshold variable τ to identify and penalise the k^{top} longest inter-cluster connections. The variable τ represents the distance limit separating normal arcs from very long ones, while s_{cd} measures how much a selected arc (c, d) exceeds that limit. Constraint (4.34) ensures that s_{cd} measures the positive deviation of each selected arc length d_{cd} from the threshold τ, such that only the longest arcs contribute to the penalty in the objective function. If d_{cd} is smaller than τ, then s_{cd} becomes zero and the arc is not penalised. If d_{cd} is larger, s_{cd} captures the difference $(d_{cd} - \tau)$, and this excess distance is added to the objective. As a result, only the k^{top} longest arcs contribute to the penalty term, preventing the route from making unnecessary long jumps.

$$s_{cd} \geq dd_{cd}\, x_{cd} - \tau, \qquad \forall\, c, d \in C,\ c \neq d \tag{4.34}$$

Revisit Span Constraints

These constraints introduce the continuous arrival time variables t_i, which represent the service start time at each node. The time-propagation constraints (4.35)–(4.36) ensure that if an arc (i, j) is selected ($x_{ij} = 1$), service at node j only start after service at node i has finished and the vehicle has travelled to j. For arcs not selected in the route ($x_{ij} = 0$), these constraints are relaxed using a big-M formulation.

To tighten the formulation, arc-specific big-M coefficients are used. For each arc (i, j), the lower and upper bounds are defined as $M_{ij}^{\mathsf{LB}} = \max(0, L_i + st_i + tt_{ij} - E_j)$ and $M_{ij}^{\mathsf{UB}} = \max(0, L_j - E_i - st_i - tt_{ij})$. These values represent the minimum slack required to deactivate the respective constraint when the arc is not selected and are derived from the earliest and latest feasible arrival times E_i and L_i. For the lower bound constraint (4.35), the worst-case scenario occurs when t_i is at its maximum ($t_i = L_i$) and

t_j is at its minimum ($t_j = E_j$), yielding $M_{ij}^{\mathsf{LB}} = L_i + st_i + tt_{ij} - E_j$. For the upper bound constraint (4.36), the worst case occurs when t_i is at its minimum ($t_i = E_i$) and t_j is at its maximum ($t_j = L_j$), yielding $M_{ij}^{\mathsf{UB}} = L_j - E_i - st_i - tt_{ij}$. Negative values are capped at zero, since in such cases the constraint is already satisfied by variable bounds. Using these arc-specific coefficients yields a tighter linear relaxation than using a single global big-M.

Finally, constraints (4.37)–(4.38) define the revisit span variables r_{cd} as the excess absolute difference in arrival times between link clusters c and d belonging to the same street segment, beyond the revisit tolerance Δ_{RS}. Consequently, revisit spans smaller than the revisit tolerance incur no penalty.

$$t_j \;\geq\; t_i + st_i + tt_{ij} - M_{ij}^{LB}\,(1 - x_{ij}) \qquad \forall i \in V, j \in V \setminus \{z_s\}, i \neq j \tag{4.35}$$

$$t_j \;\leq\; t_i + st_i + tt_{ij} + M_{ij}^{UB}\,(1 - x_{ij}) \qquad \forall i \in V, j \in V \setminus \{z_s\}, i \neq j \tag{4.36}$$

$$r_{cd} \;\geq\; t_c - t_d - \Delta_{RS} \qquad \forall (c,d) \in R \tag{4.37}$$

$$r_{cd} \;\geq\; t_d - t_c - \Delta_{RS} \qquad \forall (c,d) \in R \tag{4.38}$$

Intra-Route Crossing Constraints

The intra-route crossing constraint determines whether two directed route segments geometrically intersect. The set $\mathcal{X}$ contains all ordered pairs of directed link cluster arcs $\big((c,d),(p,q)\big)$ whose straight lines between link cluster endpoints intersect, excluding arcs from and to terminals. For each such arc pair, a binary variable $\chi_{cd,pq}$ is introduced. This variable can take the value 1 if both corresponding directed arcs (c,d) and (p,q) are selected in the route and intersect. Constraint (4.39) links the intra-route crossing variables to the routing variables. It ensures that $\chi_{cd,pq}$ is activated only if both arcs (c,d) and (p,q) are selected in the route.

$$\chi_{cd,pq} \geq x_{cd} + x_{pq} - 1 \qquad \forall\big((c,d),(p,q)\big) \in \mathcal{X} \tag{4.39}$$

Domain Definitions of Decision Variables

Finally, the domains of all decision variables are defined in constraints (4.40)–(4.52), consistent with the notation in Table 4.2. Binary and integer variables represent discrete routing, assignment, and ordering decisions, while continuous variables capture time and distance.

$$x_{ij} \in \{0,1\} \qquad \forall (i,j) \in A \tag{4.40}$$

$$u_i \in \{0,\ldots,n-1\} \qquad \forall i \in V \tag{4.41}$$

$$y_{cl} \in \{0,1\} \qquad \forall c \in C,\ l \in L \tag{4.42}$$

$$m_{pl} \in \{0,1\} \qquad \forall p \in C,\ l \in L \tag{4.43}$$

$$a_{cpl} \in \{0,1\} \qquad \forall c,p \in C,\ c \neq p,\ l \in L \tag{4.44}$$

$$p_{cdl} \in \{0,1\} \qquad \forall c,d \in C,\ c < d,\ l \in L \tag{4.45}$$

$$q_c \in \{0,\ldots,|C|\} \qquad \forall c \in C \tag{4.46}$$

$$F_h \in \{0,\ldots,|C|\}, \quad L_h \in \{0,\ldots,|C|\} \qquad \forall h \in H \tag{4.47}$$

$$s_{cd} \geq 0 \qquad \forall c,d \in C,\ c \neq d \tag{4.48}$$

$$\tau \geq 0 \tag{4.49}$$

$$E_i \leq t_i \leq L_i \qquad \forall i \in V \tag{4.50}$$

$$r_{cd} \geq 0 \qquad \forall (c,d) \in R \tag{4.51}$$

$$\chi_{cd,pq} \in \{0,1\} \qquad \forall\big((c,d),(p,q)\big) \in \mathcal{X} \tag{4.52}$$

4.3. Metric Behaviour Analysis

This section analyses how the proposed visual attractiveness metrics behave in order to support metric selection for the optimisation framework. The goal is to determine which metrics capture distinct aspects of route structure and which behave as near-substitutes. By examining their numerical relationships and their impact on route structure, a small and non-redundant set of metrics can be identified.

The analyses are performed by solving the mathematical optimisation models introduced in Section 4.2. All models were implemented in Python and solved using the Gurobi Optimiser (version 12). All computations were performed on a computer equipped with a 3.6 GHz 8-core (physical) processor and 32

GB of RAM. The analysis consists of two parts. First, an exploratory correlation analysis is performed on travel-time-optimal routes of all 12 case studies introduced in Section 3.1 to examine numerical relationships between the metrics on realistic networks. Because correlations alone do not reveal how optimising a metric reshapes a route, this is followed by single-objective optimisation on the test set introduced in Section 3.3. The test set is small enough to allow (near-)optimal solving and enables detailed inspection of the effect of each metric on the route, both numerically and visually. Together, these analyses support the selection of complementary metrics for the optimisation framework.

4.3.1. Correlation Analysis on Case Studies

This section analyses how the proposed visual attractiveness metrics behave across all 12 case studies introduced in Section 3.1. The goal is to examine whether the metrics capture distinct or overlapping aspects of visual attractiveness on realistic road networks.

Experimental Setup

The case studies are sufficiently comparable to be analysed jointly, as shown in Section 3.2. Therefore, the correlation analysis is performed on the aggregated set of case studies. These routes were obtained by solving the travel time minimisation model for each case study with an hour time limit. Four instances solved within seconds and the remaining runs reached very small optimality gaps (0.008-7.305%), which is sufficient for correlation analysis. All parameters described in Subsection 4.2.1 remained constant across datasets for comparability. Specifically, spread metrics used 10 route segments, neighbourhoods were derived using a target size of 50, the top longest jumps metric considered the longest 20% of link cluster arcs, and the revisit tolerance was set at 600 seconds. Because these routes are optimised for travel time, they provide a baseline for evaluating whether the attractiveness metrics capture distinct or overlapping properties.

For the correlation analysis, each visual attractiveness metric is computed post-hoc on the route optimised for travel time and normalised to obtain comparable average values (see (4.53)–(4.59)). Notations not explicitly mentioned here follow Section 4.2. Let A^{LC} denotes the set of link cluster arcs in the route and S_l the set of route segment sizes with medoids m_l. For medoid spread, distances are averaged over all non-medoid link clusters in each segment. Pairwise spread is averaged over all unordered link cluster pairs within segments. For the jump-based metrics, directed Euclidean distances dd_{cd} are used, which are sorted for the top longest jumps metric in descending order $dd_{(1)} \geq \cdots \geq dd_{|A^{LC}|}$. Let n_{cross} denote the number of arc pairs in the route that cross, detected using the straight line connections between link cluster endpoints. Because each crossing involves two route arcs, this number is multiplied by two to obtain the average number of crossings per route arc.

$$\text{Average Medoid Spread (MS):} \quad \mathrm{MS}_{\text{avg}} = \frac{\sum_{l\in L}\sum_{c\in S_l} d_{cm_l}}{\sum_{l\in L}(|S_l|-1)}, \tag{4.53}$$

$$\text{Average Pairwise Spread (PS):} \quad \mathrm{PS}_{\text{avg}} = \frac{\sum_{l\in L}\sum_{c\in S_l}\sum_{\substack{d\in S_l \\ c<d}} d_{cd}}{\sum_{l\in L}\frac{|S_l|(|S_l|-1)}{2}}, \tag{4.54}$$

$$\text{Average Neighbourhood Fragmentation (NF):} \quad \mathrm{NF}_{\text{avg}} = \frac{\sum_{h\in H} L_h - F_h}{|H|}, \tag{4.55}$$

$$\text{Average Top Longest Jumps (TJ):} \quad \mathrm{TJ}_{\text{avg}} = \frac{\sum_{k=1}^{k^{top}} dd_{(k)}}{k^{top}}, \tag{4.56}$$

$$\text{Average Local Jumps (LJ):} \quad \mathrm{LJ}_{\text{avg}} = \frac{\sum_{(c,d)\in A^{LC}} dd_{cd}}{|A^{LC}|}, \tag{4.57}$$

$$\text{Average Revisit Span (RS):} \quad \mathrm{RS}_{\text{avg}} = \frac{\sum_{(c,d)\in R}\max\{0, |t_c - t_d| - \Delta_{RS}\}}{|R|}, \tag{4.58}$$

$$\text{Average Intra-Route Crossings (IC):} \quad \mathrm{IC}_{\text{avg}} = \frac{2\,n_{\text{cross}}}{|A^{LC}|} \tag{4.59}$$

Three correlation coefficients are used to examine relationships between the visual attractiveness metrics. Pearson's coefficient r measures linear association, but is sensitive to scale and outliers (Pearson,

1896). Therefore, the rank-based coefficients of Spearman (1904) and Kendall (1938) are also reported. Spearman's coefficient ρ captures any relationship using the ranks of the data, while Kendall's coefficient τ quantifies the probability that two metrics order the routes in the same way. All three coefficients are reported to assess whether observed relationships are consistent across different correlation measures. All coefficients range from -1 to 1, where 1 indicates perfect positive association, 0 no association, and -1 perfect negative association.

Correlation Matrix Analysis

The correlation matrix in Table 4.3 shows how the visual attractiveness metrics relate across all case study routes. Several metrics capture similar aspects of spatial compactness, while others reflect different structural properties of the route. These correlations describe associations rather than causal effects and reflect how the metrics covary on travel-time-optimal routes, not how they behave when used as optimisation objectives. Several correlation patterns become clear.

First, most compactness metrics (MS, PS, TJ, and LJ) show moderate to strong positive correlations. Medoid spread (MS) and pairwise spread (PS), as well as top longest jumps (TJ) and local jumps (LJ), are almost perfectly correlated. This indicates that they behave as near-substitutes and is consistent with their definitions. MS and PS both assess compactness through regional compactness by grouping link clusters, while TJ and LJ measure compactness through local compactness by penalising long moves. Moreover, TJ and LJ are moderately positive correlated with MS and PS, suggesting that routes with fragmented service within areas also tend to contain longer moves.

Second, neighbourhood fragmentation (NF) and revisit span (RS) are only weakly related to the other compactness metrics. NF shows very weak correlations with MS, PS, TJ, and LJ, indicating that it behaves largely independently from other compactness measures. NF is moderately positive correlated with RS, which is intuitive, as routes that leave and later re-enter neighbourhoods also tend to revisit street segments after longer time intervals. Apart from this relationship, RS shows only weak associations with the remaining compactness metrics.

Finally, intra-route crossings (IC) behave differently from the other metrics. IC shows weak to moderate negative correlations with MS, PS, TJ, LJ, and RS. This likely reflects density effects rather than a conflict with compactness. Spatially compact routes offer more opportunities for geometric crossings simply because link clusters visited after each other in the route lie closer together. Conversely, more dispersed routes naturally reduce such crossings. Moreover, IC shows a weak positive correlation with neighbourhood fragmentation (NF), indicating that routes with more neighbourhood re-entries tend to exhibit slightly more crossings.

Table 4.3: Correlation Matrix of Normalised Visual Attractiveness Metrics Across Case Study Routes

	MS	PS	NF	TJ	LJ	RS	IC
MS	1.00; 1.00; 1.00						
PS	0.96; 0.88; 1.00	1.00; 1.00; 1.00					
NF	0.02; -0.03; -0.01	0.11; 0.03; -0.01	1.00; 1.00; 1.00				
TJ	0.48; 0.33; 0.99	0.52; 0.39; 0.99	-0.11; -0.09; -0.07	1.00; 1.00; 1.00			
LJ	0.59; 0.49; 1.00	0.66; 0.55; 1.00	0.03; 0.00; -0.01	0.93; 0.85; 0.99	1.00; 1.00; 1.00		
RS	0.13; 0.15; 0.18	0.13; 0.09; 0.18	0.27; 0.21; 0.34	0.09; 0.09; 0.13	0.16; 0.12; 0.18	1.00; 1.00; 1.00	
IC	-0.22; -0.12; -0.13	-0.19; -0.12; -0.13	0.14; 0.06; 0.11	-0.44; -0.30; -0.13	-0.43; -0.27; -0.12	-0.18; -0.18; -0.19	1.00; 1.00; 1.00

Note. Each cell reports ρ, τ, r for Spearman, Kendall, and Pearson correlations. Values $\geq |0.70|$ typically indicate strong correlations.

4.3.2. Single-Objective Optimisation Analysis on Test Set

This section presents a numerical and visual analysis of the exact single-objective models evaluated on the test set described in Section 3.3. Each model optimises one metric, while all other metrics are computed post-hoc using the definitions in Subsection 4.2.2. This shows the effect of optimising for each individual metric on the route structure.

Experimental Setup

All models were solved using a common set of parameters as described in Subsection 4.2.1. Neighbourhoods for the test set were defined manually rather than using k-means. Because the test set is

small and extracted from a larger case study, automatically generated neighbourhoods may group link clusters unrealistically. Instead, link clusters were grouped into four neighbourhoods based on spatial proximity, as identified using AMCS software. Revisit pairs were assigned to the same neighbourhood. To ensure comparable aggregation for the spread-based objectives, four route segments were used. The top longest jumps metric considered the longest 20% of link cluster arcs and the revisit span tolerance was set to 120 seconds. It is important to note that the results are specific to the test set. Metric values and route structures may depend on the network layout and parameter choices. Thus, the findings should be interpreted as descriptive of model behaviour rather than as general conclusions.

Numerical Analysis

The results of the single-objective models on the small test set are summarised in Table 4.4. For each model, the optimised metric is reported together with all other metrics computed post-hoc, allowing direct comparison of their effects on the route. The reported travel time refers to deadhead travel time, as defined in Section 4.2. Service times of link clusters represent travel within clusters and are therefore fixed and excluded from the optimisation. Runtime and optimality gaps are also reported. The medoid spread (MS) and revisit span (RS) models reached the time limit. For MS, the small optimality gap and clear improvement in the targeted metric suggest that the solution is close to optimal. The RS model stopped with a large optimality gap due to a weak LP relaxation, but the solution still shows substantial improvements in the revisit span metric. Therefore, the results are sufficient to analyse model behaviour. All other models were solved to optimality.

The travel time (TT) model serves as the efficiency baseline. As expected, it achieves the lowest travel time but performs poorly on most visual attractiveness metrics. This confirms that routes optimised purely for efficiency lack spatial and geometric coherence. TT performs relatively well on the jump-based metrics, which is plausible since these metrics also affect the visitation sequence and use directed distances that are partially aligned with travel time. Optimising for any visual attractiveness metric substantially increases travel time, with all routes exceeding the baseline by more than 50%. The visual attractiveness objectives can be grouped into three categories, namely grouping-based, sequence-based, and single-purpose objectives. These groups differ in how they influence the link cluster visitation order and in whether improvements extend beyond the optimised metric.

The grouping-based objectives, consisting of medoid spread (MS), pairwise spread (PS), and neighbourhood fragmentation (NF), improve regional spatial coherence. MS and PS behave consistent with their strong positive correlation, as optimising one also improves the other. However, they have little effect beyond these metrics. This was expected, as they enforce regional compactness by grouping link clusters, but do not control the visitation sequence of link clusters within and between areas. Similarly, NF focuses on neighbourhood areas and slightly improves MS, PS, and revisit span (RS). However, it performs poorly on other metrics. Overall, grouping-based objectives primarily enforce regional compactness and have limited influence on local compactness.

In contrast, the sequence-based objectives, namely top longest jumps (TJ) and local jumps (LJ), directly penalise undesirable moves between link clusters and therefore restructure the entire route. The two models display identical behaviour on the test set, reflecting their near-perfect correlation. In addition to their own objectives, TJ and LJ consistently improve the other visual attractiveness metrics.

The remaining objectives, RS and intra-route crossings (IC), act as single-purpose metrics. The RS model strongly reduces revisit span but performs poorly on all other metrics, even worse than the travel time baseline. This was expected, as optimising for side-loader-specific local compactness only will not result in improvements in other metrics. Similarly, the IC model eliminates crossings but shows little improvement elsewhere. Unlike the sequence-based objectives, neither RS nor IC meaningfully restructures the entire visitation order of link clusters, confirming their specialised role.

Table 4.4: Results of Single-Objective Exact Models for the Small Test Set

Model	TT (s)	MS (m)	PS (m)	NF	TJ (m)	LJ (m)	RS (s)	IC	Runtime (s)	Gap (%)
TT	2555	2758	12 464	11	921	1887	2922	17	0	0.0
MS	4237	1536	7453	18	1250	2748	3153	27	10 800	1.2

Continued on next page

Table 4.4: Results of Single-Objective MIP and MILP Models on the Small Test Set *(Continued)*

Model	TT (s)	MS (m)	PS (m)	NF	TJ (m)	LJ (m)	RS (s)	IC	Runtime (s)	Gap (%)
PS	4764	1549	7434	19	1257	2698	3698	28	120	0.0
NF	5136	2068	9487	0	1148	2879	2770	33	22	0.0
TJ	3850	2066	9577	12	309	847	1599	1	2	0.0
LJ	3850	2066	9577	12	309	847	1599	1	1	0.0
RS	5172	3012	13 191	54	1561	3680	155	54	10 801	100
IC	4683	2943	13 351	27	1145	2393	7431	0	2	0.0

Note. Objective function values are computed as described in Subsection 4.2.2. All other metric values, that is, those not optimised by the corresponding model, are computed post-hoc using the same definitions.

Visual Analysis

While the previous subsection analysed the metrics numerically, visual attractiveness is inherently spatial. To understand how the single-objective models shape a route, this subsection inspects the routes visually. For clarity, only the travel-time-optimal route is shown here, while the remaining routes are provided in Appendix D. All plots use a consistent visual encoding. Link clusters are shown at their midpoints with visitation sequence labels and their start and end order nodes are displayed explicitly, with directed arcs connecting successive link clusters. The route segments are colour-coded, revisit pairs are connected with green lines, and the longest 20% of arcs are highlighted in purple.

The travel-time-optimal route shown in Figure 4.2 lacks both spatial compactness and geometric coherence. The route is fragmented, with frequent detours and many intra-route crossings. Both spatially close link clusters and revisit pairs are often visited far apart in the sequence. This clearly shows the absence of visual attractiveness when optimising purely for travel time.

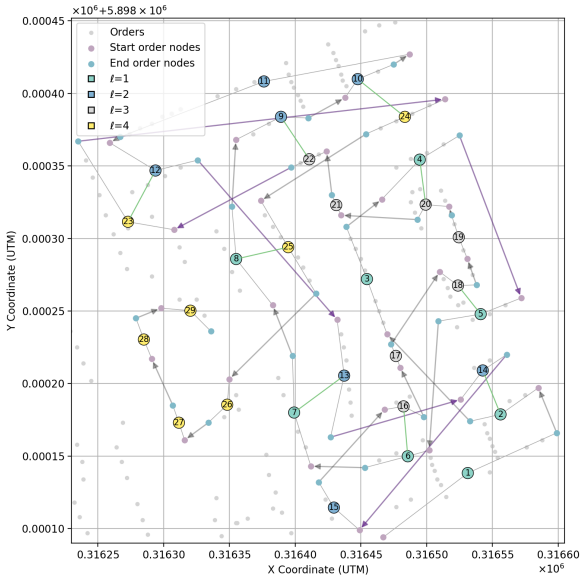

Figure 4.2: Route of Model Optimised for Travel Time on the Abstract Network for the Test Set

Across the remaining models, three distinct behaviours become apparent, corresponding to the grouping-based, sequence-based, and single-purpose objectives identified in the numerical analysis. The grouping-based objectives, medoid spread (MS), pairwise spread (PS), and neighbourhood fragmentation (NF), mainly affect which link clusters are served together. Routes optimised for MS and PS form spatially compact segments of link clusters served together, but there is no ordering of link clusters within and between segments. This leads to long jumps, large revisit spans, and several geometric crossings. Similarly, NF keeps the route within the same neighbourhood for longer periods. Nevertheless, moves between neighbourhoods seem illogical and the visitation order within neighbourhoods is unstructured. As a result, grouping-based objectives enforce regional compactness by grouping nearby link clusters together and completing one area before moving on. However, they do not improve local compactness

as they leave the sequence of link clusters within and between areas unstructured. Additionally, the fixed-size route segments used by MS and PS do not necessarily align with geographically meaningful areas and may group distant link clusters or split dense areas. In contrast, the manually defined neighbourhoods reflect logical groupings based on spatial proximity and road network connectivity.

The sequence-based objectives, top longest jumps (TJ) and local jumps (LJ), directly influence the visitation sequence of link clusters. By penalising undesirable moves, they enforce local compactness along the entire route. Visually, these models produce routes with stronger neighbourhood coherence, fewer detours, and fewer crossings than the travel time baseline. The models result in identical routes, reflecting their closely related formulations and near-perfect correlation. Although the two objectives differ conceptually, the results indicate no clear trade-off in this test set between limiting long jumps and minimising average jump length. The route can therefore be optimised to reduce both the longest moves and the overall route length simultaneously.

Finally, the revisit span (RS) and intra-route crossings (IC) objectives behave as single-purpose models. The RS model reduces the time between revisit pairs, but this metric is specific to side loaders and its effect is highly localised. The route remains spatially fragmented, with many long jumps and geometric crossings. Similarly, the IC model eliminates crossings entirely, but the route remains spatially dispersed and poorly grouped. Unlike the sequence-based objectives, neither RS nor IC achieve meaningful reordering of the route or grouping of link clusters, confirming their specialised role.

4.4. Summary and Metric Selection

This chapter used MILP models to define, analyse, and select visual attractiveness metrics for the optimisation framework. The goal was not to compare these models as final solutions, but to identify metrics that capture distinct and meaningful aspects of visual attractiveness. Metric behaviour was analysed using two complementary approaches. First, a correlation analysis on travel-time-optimal routes of the 12 case studies examined how the metrics covary on realistic networks. Because correlations do not show how metrics shape a route when optimised, this was complemented by single-objective optimisation on a small test set. The test set allowed (near-)optimal solving and enabled numerical and visual inspection of how each metric influences route structure.

The correlation results show that several metrics behave as near-substitutes. In particular, medoid spread (MS) and pairwise spread (PS), as well as top longest jumps (TJ) and local jumps (LJ), capture closely related aspects of compactness. In contrast, neighbourhood fragmentation (NF) and revisit span (RS) show weaker relationships with other metrics, while intra-route crossings (IC) behaves largely independently. Single-objective optimisation reveals three behavioural categories. Grouping-based objectives (MS, PS, NF) improve regional compactness but have limited influence on the global visitation order. Sequence-based objectives (TJ and LJ) restructure the full route and provide the broadest improvements across visual attractiveness metrics. Single-purpose objectives (RS and IC) improve their target characteristic, but offer no additional benefit compared to the travel time (TT) baseline.

Based on these insights, five metrics are selected for the optimisation framework, namely TT, NF, LJ, RS, and IC. TT is retained as the efficiency metric. LJ is selected as the sequence-based metric to help structure the full route and enforce local compactness. Although LJ and TJ produced identical routes on the test set, they differ conceptually. LJ minimises the average distance between consecutive link clusters and allows occasional longer moves, whereas TJ penalises the longest moves. This distinction may become more relevant for larger routes with distant neighbourhoods, while LJ also avoids additional parameter choices. Moreover, NF is included to enforce regional compactness, ensuring that service within one area is completed before moving to the next. This behaviour is not guaranteed by sequence-based metrics alone. NF is preferred over MS and PS because neighbourhoods derived from k-means adapt to spatial density, unlike fixed-size route segments. Finally, RS and IC are kept as specialised metrics that capture revisit behaviour and route complexity. Together, these complementary metrics capture efficiency, regional and local compactness, and route complexity.

The findings in this chapter are descriptive and depend on the experimental setup, particularly the use of the small test set with manually defined neighbourhoods and specific parameter choices. Nevertheless, the analyses provide clear guidance for metric selection. Chapter 5 builds on these results by introducing the optimisation framework and integrating the selected metrics.

5

Optimisation Framework

This chapter presents the optimisation framework proposed to construct residential waste collection routes that balance travel time with visual attractiveness. Since routing problems with multiple interdependent objectives are NP-hard, exact optimisation becomes computationally infeasible for realistic problem instances. Therefore, the framework adopts the simulated annealing (SA) metaheuristic, which enables exploration of the solution space and effective trade-offs between solution quality and computational effort. This choice follows directly from the literature review summarised in Section 2.5, which highlighted the need for a method that can handle large asymmetric networks, competing objectives, and escaping local optima. Building on the five optimisation objectives selected in Section 4.4, the framework is described in two complementary ways. First, a conceptual framework outlines the overall flow of the optimisation procedure. Second, a detailed algorithmic framework specifies the exact steps, parameters, and decision rules used in the implementation.

5.1. Conceptual Optimisation Framework

The optimisation framework is based on a SA procedure that starts from an initial route and iteratively explores alternative solutions. SA is well suited for optimisation problems in which multiple metrics are combined in a weighted sum objective function, where improving one metric may deteriorate others. In such settings, greedy or targeted local search methods often converge prematurely to local optima. SA addresses this limitation by allowing the probabilistic acceptance of worsening moves. This mechanism enables the search to escape local optima and explore the solution space more broadly. As the algorithm progresses, the acceptance probability gradually decreases through a cooling schedule, causing the search behaviour to shift from exploration towards exploitation.

Candidate routes are generated using random neighbourhood operators rather than highly specialised problem-specific rules. In SA, it is common to generate neighbouring solutions by applying simple random moves, after which the acceptance mechanism determines whether the new solution is retained (Talbi, 2009). Using generic operators instead of targeted local search operators allows a broad exploration of the solution space without imposing assumptions about which nodes or arcs should be modified. In contrast, highly targeted operators may restrict modifications to specific parts of the route structure, which can bias the search and reduce the ability of the algorithm to escape local optima. The objective function guides the search towards improved routes, while the combination of random exploration and probabilistic acceptance enables SA to balance diversification and intensification in a controlled manner.

The optimisation process is initialised with a route that is (near-)optimal with respect to travel time, providing an operationally feasible baseline from which the SA search is initiated. This initial route is obtained by solving the exact travel time minimisation model presented in Section 4.2. For this route, the values of the five selected metrics, travel time (TT), neighbourhood fragmentation (NF), revisit span (RS), intra-route crossings (IC), and local jumps (LJ), are computed, as described in Chapter 4.

Following this, all parameters are initialised. These include the SA parameters, metric-specific param-

eters, operator-related parameters, and preference weights that indicate the relative importance of the different metrics in the objective function. Route quality is evaluated using a weighted sum objective function, where the preference weights are normalised using the metric values of the initial solution to ensure comparability across metrics. A lower objective value indicates a better route. The use of a weighted sum is motivated by three considerations, consistent with the discussion in Section 2.5. First, the problem involves soft trade-offs between conflicting route characteristics rather than strict service constraints. Second, a single objective is compatible with SA's acceptance mechanism. Third, travel time acts as a regularising term that connects the optimisation to the underlying road network. Without travel time in the objective, the remaining metrics alone may permit unrealistic routing behaviour.

The preference weights in the weighted sum objective are fixed. Preference weights for TT, NF, RS, IC, and LJ are set to 2.0, 1.5, 1.0, 0.5, and 0.5, respectively. These values are chosen as informed estimates that reflect both the relative importance of the objective components and their observed behaviour in the exact optimisation results presented in Subsection 4.3.2. Fixing the weights is a deliberate modelling choice. In practice, the relative importance of route characteristics is an operational decision that depends on the preferences of the customer. Therefore, the selected weights represent a reasonable baseline.

Travel time is assigned the highest weight, as it is the only metric that explicitly accounts for the road network and therefore ensures route feasibility. Neighbourhood fragmentation is weighted strongly because it is the only metric that groups link clusters and enforces that they are visited together, which is the main goal of compactness. The revisit span and local jumps metrics aim to improve compactness within neighbourhoods, with the revisit span weighted higher because long time intervals between revisiting the same street segment are an issue specific to side loaders. Local jumps and intra-route crossings are based on straight-line approximations and do not fully reflect road network constraints. Therefore, they are assigned low weights and mainly act as tie-breaking criteria.

At the start of the search, the current solution and the best-so-far solution are both set equal to the initial route. Additionally, the current temperature T is set to the initial temperature T_0. The optimisation loop iterates through temperature levels. During the optimisation loop, the framework iteratively applies random local search operators to generate alternative routes. In each iteration, an operator is selected according to a predefined probability distribution and used to generate a set of n candidate routes. All candidates are evaluated using the weighted objective function, and the best candidate is selected for consideration. Sampling multiple candidates reduces sensitivity to individual random moves while preserving the exploratory nature of SA. Let F_{cur} denote the objective value of the current route and F_{cand} that of the selected candidate. The change in objective value is computed as

$$\Delta F = F_{\text{cand}} - F_{\text{cur}}.$$

The selected candidate is accepted deterministically if it improves or maintains the objective value ($\Delta F \leq 0$). Otherwise, it may still be accepted with a probability that depends on both the deterioration ΔF and the current temperature T. In this case, the acceptance probability is computed as

$$p = \exp\left(-\frac{\Delta F}{T}\right).$$

The candidate solution is accepted if a random draw $r \in [0, 1]$ satisfies $r < p$. This mechanism enables occasional acceptance of worsening moves, which helps the search procedure to escape local optima, particularly at higher temperatures. Accepted candidates replace the current solution, and whenever an accepted solution improves upon the best-so-far objective value, the best-so-far solution is updated accordingly. This separation between the current solution and the best recorded solution allows the algorithm to explore intermediate solutions without losing track of the best route found.

The algorithm performs N_T iterations at each temperature level T before reducing the temperature according to a geometric cooling schedule $T = \alpha T$ and resetting the iteration counter per temperature N to zero. Performing multiple local search iterations at a fixed temperature ensures that the neighbourhood is explored sufficiently at each temperature level before the search becomes more exploitative. The optimisation process stops when the temperature falls below a predefined minimum threshold T_{final}. The conceptual flow of the framework is illustrated in Figure 5.1, while the detailed algorithmic formulation is presented in the following section.

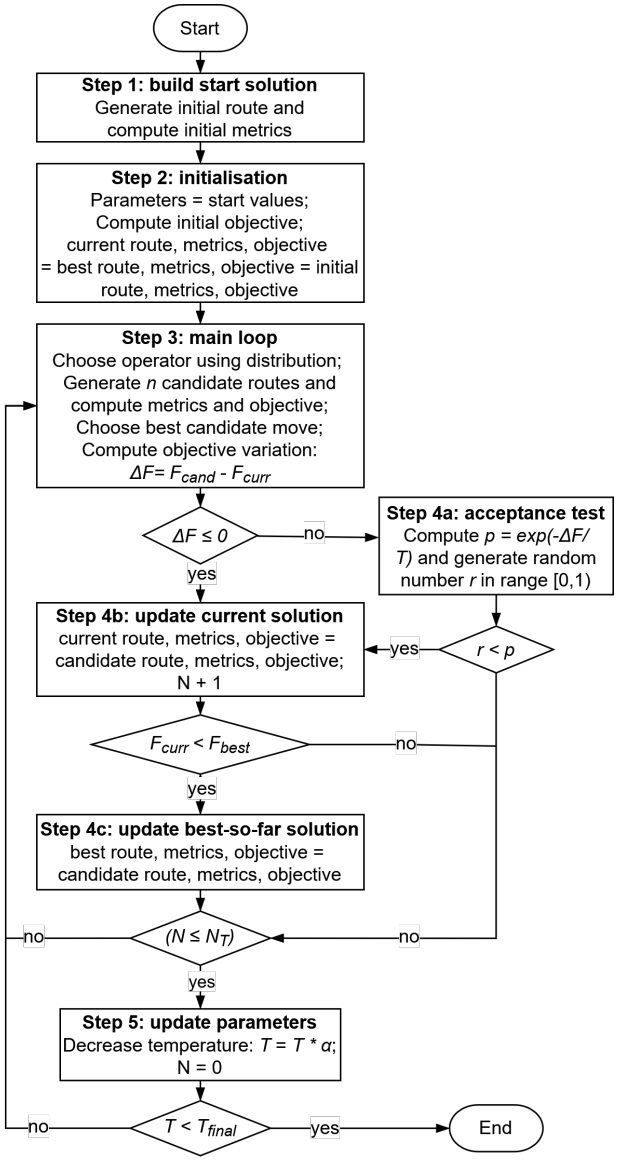

Figure 5.1: Conceptualisation of the Optimisation Framework

5.2. Algorithmic Optimisation Framework

This section translates the conceptual framework into an explicit algorithmic formulation. The optimisation procedure is specified through the pseudocode presented in Algorithm 1, which defines the required inputs, parameter initialisation, and update rules used in the implementation.

The procedure takes as input the SA parameters, including the initial temperature T_0, the cooling rate α, the number of iterations per temperature level N_T, and the final temperature T_{final}. In addition, the preference weights for the metrics $\hat{w}_{\text{metric}}$, the probability distribution of the local search operators $P(\text{op})$, the Or-Opt block length, the number of candidate moves per iteration n, the revisit span tolerance Δ_{RS}, and the number of neighbourhoods k are specified. In this thesis, the neighbourhood operators, the Or-Opt block length, the operator probability distribution, the number of candidate moves, and the final temperature are treated as fixed design choices. Their values and motivation are discussed in the following paragraphs. Moreover, the revisit span tolerance and the number of neighbourhoods are set per experiment based on the problem instance size. To ensure reproducibility, the random seed is fixed at 42. The remaining SA parameters are selected by validation experiments, as described in Chapter 6.

Before the heuristic search begins, all required inputs are prepared and organised. The problem data consist of the link cluster and terminal sets, the directed arc set with travel and service times, the

neighbourhood labels used for computing neighbourhood fragmentation (created by k-means procedure explained in Section C.1), the directed and undirected Euclidean distances between link clusters, the revisit pairs, and geometric information required for detecting intra-route crossings. Together, these data allow candidate routes to be constructed and evaluated across all metrics.

Following this, the exact travel time minimisation model is solved to obtain an initial route that is optimal with respect to travel time. The associated metric values are computed and stored. During the subsequent initialisation phase, the normalised weights are obtained by dividing the preference weights by their metric value of the initial solution. Using these normalised weights, the objective function F is defined as a weighted sum of the metric values and evaluated for the initial solution. Finally, the route, objective value, and metric values of both the current solution and the best-so-far solution are initialised to those of the initial route. In addition, the current temperature T is set to the initial temperature T_0.

The optimisation procedure consists of a nested loop structure, with an outer loop iterating over decreasing temperature levels and an inner loop performing N_T iterations at a fixed temperature. Within each inner iteration, a local search operator is selected probabilistically from the set containing the relocate, swap, and Or-Opt operators. The relocate operator removes a single link cluster and reinserts it at a different position, the swap operator exchanges the positions of two link clusters, and the Or-Opt operator relocates a contiguous block of three link clusters. As discussed in Subsection 2.3.2, these operators are widely used in routing heuristics and provide simple but effective modifications of the route. These operators modify the route without reversing subsequences, which makes them suitable for asymmetric routing problems, where arc costs depend on direction. In contrast, operators such as 2-opt reverse route segments and often perform poorly in asymmetric settings. Larger arc exchange operators, such as extended k-opt, were not considered because they increase computational effort.

The operator probability distribution is fixed at 0.4 for relocate and 0.3 for both swap and Or-Opt. Relocate is slightly favoured as it corrects small inefficiencies, while swap and Or-Opt enable slightly larger adjustments. For Or-Opt, a block length of three is used as a balance between improvement and disruption, as shorter blocks offer limited impact while longer blocks increase the likelihood of worsening candidates. Together, these operators enable exploration of neighbourhoods with varying impact, ranging from local to slightly larger route modifications. The detailed pseudocode of the operators is provided in Section C.2.

The selected operator generates a set of n candidate routes based on the current route. Here n is fixed to 10, as it provides a compromise between solution quality and computational effort. Each candidate is evaluated by computing its metric values and the corresponding weighted sum objective value. Among the candidate routes generated in a given iteration, the candidate with the lowest objective value is selected as the best candidate move. For the selected candidate, the change in objective value relative to the current solution is computed. If the candidate improves or maintains the objective value, it is accepted deterministically. Otherwise, it may still be accepted according to the acceptance probability criterion $r < p$, which depends on both the objective deterioration Δ_F and the current temperature T.

When a candidate move is accepted, the current route, its metric values, and the objective value are updated accordingly. In addition, the updated current solution is compared to the best-so-far solution maintained throughout the search. If the current objective value improves on the best-so-far objective value, the best-so-far route and its associated metrics are updated. This separation between the current solution and the best-so-far solution ensures that the algorithm can explore intermediate solutions without losing track of the best route encountered.

After completing N_T iterations at a temperature level, the temperature is reduced according to a geometric cooling schedule $T = \alpha T$. The iteration counter N is then reset, and the search continues at the new temperature level. The optimisation stops when either the temperature falls below the final threshold T_{final}. The final temperature is fixed at $T_{\text{final}} = 10^{-4}$. This value aligns with the magnitude of the weighted sum objective and serves only as a stopping criterion rather than affecting the search behaviour. Upon termination, the best-so-far route is returned as the final solution.

Algorithm 1 SA with Multiple Random Operators and Candidate Sampling

procedure SA($T_0, \alpha, T_{final}, N_T, \hat{w}_{\mathsf{TT}}, \hat{w}_{\mathsf{NF}}, \hat{w}_{\mathsf{IC}}, \hat{w}_{\mathsf{LJ}}, \hat{w}_{\mathsf{RS}}, P(\mathsf{op}), n, \Delta_{RS}, k$)

▷ **Step 0:** Preprocessing of data

$C \leftarrow \{\mathsf{LC}\},\ Z \leftarrow \{\text{terminals}\},\ V \leftarrow C \cup Z,\ A \leftarrow \{\text{arcs}\},\ R \leftarrow \{\text{revisit pairs}\}$

$tt_{ij} \leftarrow$ travel time on each arc $(i,j) \in A$, $\quad s_i \leftarrow$ service time at each node $i \in V$

$dd_{cd} \leftarrow$ directed Euclidean distance for all $c, d \in C, c \neq d$

$d_{cd} \leftarrow$ undirected Euclidean distance for all $c, d \in C, c \neq d$

$H \leftarrow$ neighbourhoods, $\nu(c) \leftarrow$ neighbourhood labels $\forall c \in C$ (obtained from KMeans($P^0_{\mathsf{LC}}, k, 42$))

$(x_c^{\mathsf{start}}, y_c^{\mathsf{start}}), (x_c^{\mathsf{end}}, y_c^{\mathsf{end}}) \leftarrow$ endpoint coordinates for each $c \in C$

▷ **Step 1:** Obtain initial solution from travel time minimisation model

$P^0_{\mathsf{nodes}} \leftarrow$ ordered node sequence of travel time optimal route

$P^0_{\mathsf{LC}} \leftarrow$ ordered link cluster sequence obtained by removing terminals Z from P^0_{nodes}

$TT_0, NF_0, IC_0, LJ_0, RS_0 \leftarrow$ initial metrics using $P^0_{\mathsf{LC}}, Z, tt_{ij}, s_i, \nu(c), (x_c^{\mathsf{start}}, y_c^{\mathsf{start}}), (x_c^{\mathsf{end}}, y_c^{\mathsf{end}}), dd_{cd}, R, \Delta_{RS}$

▷ **Step 2:** Initialisation

$w_{\mathsf{TT}} \leftarrow \hat{w}_{\mathsf{TT}}/TT_0,\ w_{\mathsf{NF}} \leftarrow \hat{w}_{\mathsf{NF}}/NF_0,\ w_{\mathsf{IC}} \leftarrow \hat{w}_{\mathsf{IC}}/IC_0,\ w_{\mathsf{LJ}} \leftarrow \hat{w}_{\mathsf{LJ}}/LJ_0,\ w_{\mathsf{RS}} \leftarrow \hat{w}_{\mathsf{RS}}/RS_0$ ▷ Normalisation

$F(TT, NF, IC, LJ, RS) \leftarrow w_{\mathsf{TT}}TT + w_{\mathsf{NF}}NF + w_{\mathsf{IC}}IC + w_{\mathsf{LJ}}LJ + w_{\mathsf{RS}}RS$ ▷ Define objective function

$P_{\mathsf{LC}} \leftarrow P^0_{\mathsf{LC}}, \quad F \leftarrow F(TT_0, NF_0, IC_0, LJ_0, RS_0)$ ▷ Current route and objective

$P^\star_{\mathsf{LC}} \leftarrow P_{\mathsf{LC}}, \quad F^\star \leftarrow F$ ▷ Best route and objective

$TT^*, NF^*, IC^*, LJ^*, RS^* \leftarrow TT, NF, IC, LJ, RS \leftarrow TT_0, NF_0, IC_0, LJ_0, RS_0$ ▷ Best & current metrics

$T \leftarrow T_0$ ▷ Current temperature

▷ **Step 3:** SA loop with N_T iterations per temperature T

while $T > T_{\mathsf{final}}$ **do**

  $N \leftarrow 0$

  while $N < N_T$ **do**

    $N \leftarrow N + 1$

    choose operator op $\sim P(\mathsf{op})$ where op $\in \{\mathsf{Relocate}, \mathsf{Swap}, \mathsf{OrOpt}\}$

    $P_{\mathsf{LC}}^{\mathsf{cand}}, TT^{\mathsf{cand}}, NF^{\mathsf{cand}}, IC^{\mathsf{cand}}, LJ^{\mathsf{cand}}, RS^{\mathsf{cand}}, F_{\mathsf{cand}} \leftarrow$ **None**

    for $i \in n$ **do** ▷ Sample n candidates and choose best candidate

      if $op =$ Relocate **then**

        $P_{\mathsf{LC}}^{'} \leftarrow \mathsf{Relocate}(P_{\mathsf{LC}}, 42)$

      else if $op =$ Swap **then**

        $P_{\mathsf{LC}}^{'} \leftarrow \mathsf{Swap}(P_{\mathsf{LC}}, 42)$

      else

        $P_{\mathsf{LC}}^{'} \leftarrow \mathsf{OrOpt}(P_{\mathsf{LC}}, 3, 42)$

      end if

      $TT^{'}, NF^{'}, IC^{'}, LJ^{'}, RS^{'} \leftarrow$ metrics using $(Z, tt_{ij}, s_i, \nu(c), (x_c^{\mathsf{start}}, y_c^{\mathsf{start}}), (x_c^{\mathsf{end}}, y_c^{\mathsf{end}}), dd_{cd}, R, \Delta_{RS})$

      $F^{'} \leftarrow F(TT^{'}, NF^{'}, IC^{'}, LJ^{'}, RS^{'})$

      if $F' < F_{\mathsf{cand}}$ **then**

        $P_{\mathsf{LC}}^{\mathsf{cand}}, TT^{\mathsf{cand}}, NF^{\mathsf{cand}}, IC^{\mathsf{cand}}, LJ^{\mathsf{cand}}, RS^{\mathsf{cand}}, F_{\mathsf{cand}} \leftarrow P_{\mathsf{LC}}^{'}, TT^{'}, NF^{'}, IC^{'}, LJ^{'}, RS^{'}, F^{'}$

      end if

    end for

    $\Delta_F \leftarrow F_{\mathsf{cand}} - F$ ▷ Compute objective change w.r.t. current objective

    if $\Delta_F \leq 0$ **then** ▷ Acceptance test

      $accept \leftarrow$ true

    else

      $p \leftarrow \exp(-\Delta_F/T)$, draw $r \sim U[0,1]$

      $accept \leftarrow (r < p)$

    end if

    if $accept$ **then** ▷ Update current solution

      $P_{\mathsf{LC}} \leftarrow P_{\mathsf{LC}}^{\mathsf{cand}}, \quad F_{\mathsf{cur}} \leftarrow F_{\mathsf{cand}}, \quad TT, NF, IC, LJ, RS \leftarrow TT^{\mathsf{cand}}, NF^{\mathsf{cand}}, IC^{\mathsf{cand}}, LJ^{\mathsf{cand}}, RS^{\mathsf{cand}}$

      if $F < F^\star$ **then** ▷ Update best-so-far solution

        $P^\star_{\mathsf{LC}} \leftarrow P_{\mathsf{LC}}, \quad F^\star \leftarrow F, \quad TT^*, NF^*, IC^*, LJ^*, RS^* \leftarrow TT, NF, IC, LJ, RS$

      end if

    end if

  end while

  $T \leftarrow \alpha T$ ▷ Cooling update after N_T iterations

end while

return $P^\star_{\mathsf{LC}}, F^*, TT^*, NF^*, IC^*, LJ^*, RS^*$ ▷ Return best-so-far solution

end procedure

6

Verification and Validation

This chapter assesses the correctness and usefulness of the optimisation framework. Since the framework combines multiple visual attractiveness metrics with travel time, it is essential to demonstrate that the implementation behaves as specified and that the resulting routes are meaningful for side loader operations. Therefore, the evaluation is structured into two complementary parts.

First, verification addresses the question “are we building the model right?” and focuses on model correctness, including the construction of the weighted sum objective, the behaviour of the neighbourhood operators, and the implementation of the SA search process. Second, validation addresses the question “are we building the right model?” and focuses on the behaviour of the SA search process. In particular, the validation examines whether the chosen SA parameter settings lead to consistent and interpretable search behaviour and produce improved routes across case studies. Subsequently, the validated parameter configuration is fixed and used for the evaluation of case studies presented in Chapter 7.

6.1. Verification

Verification concerns internal correctness by ensuring that the optimisation framework behaves as intended. The purpose of this section is to confirm that the weight normalisation and objective values are computed correctly, that neighbourhood operators modify the routes as expected, and that the SA search process responds correctly to changes in control parameters. This step is essential, as small implementation errors may not be visible in aggregate results but can lead to incorrect conclusions about model behaviour.

6.1.1. Experimental Setup

All verification experiments are conducted on the small test set introduced in Section 3.3, which allows for detailed numerical inspection and clear visual interpretation of individual algorithmic components. As in the exact optimisation experiments described in Subsection 4.3.2, neighbourhoods are defined manually and the revisit span tolerance is fixed at 120 seconds to ensure consistency across methods.

All experiments in this section were implemented in Python. The initial route was generated by solving the exact travel time minimisation model described in Section 4.2 using the Gurobi Optimiser (version 12), which yields an optimal solution in less than a second for this test instance. This route provides an operationally feasible baseline from which the SA search is initiated. All computations were performed on a computer equipped with a 3.6 GHz 8-core (physical) processor and 32 GB of RAM. All remaining parameters and weights are fixed to the values specified in Section 5.2, except for the parameters that explicitly control SA search behaviour, which vary when required for verification purposes. These include the initial temperature T_0, the cooling factor α, and the number of iterations per temperature level N_T.

6.1.2. Verification of the Objective Function and Operators

The first experiments verify the correct construction of the weighted sum objective function through weight normalisation and the correct functioning of the neighbourhood operators. Since the optimisation framework combines multiple metrics, errors in weight normalisation or objective computation could affect the search without being immediately apparent. Moreover, the operators must be visually verified before analysing larger experiments, as the solution space is explored through local search moves.

Starting from the initial route, a single candidate route is generated independently for each random neighbourhood operator (i.e., relocate, swap, and Or-Opt). For the initial route and each modified route, the raw metric values are computed, the preference weights are normalised based on the initial metric values, and the weighted sum objective value is calculated. The raw metrics are computed as described in Subsection 4.2.2, whereas the normalised weights and objective values follow the pseudocode in Section 5.2. The resulting metric values, normalised weights, and objective values are reported in Table 6.1. Recomputing the normalised weights from the raw weights and initial metric values reproduces the weights used by the model. Similarly, recomputing the objective using these weights and the corresponding metric values yields the reported objective values. This confirms that both the weight normalisation and the objective function are correctly constructed.

Table 6.1: Metric, Normalised Weight, and Objective Values for Individual Operator Moves

Route	TT (s)	NF	RS (s)	IC	LJ (m)	F
Initial (travel time optimal)	2555	11	2922	17	1887	4.90
After one relocate	2768	16	3135	25	2315	5.96
After one swap	2928	18	3590	17	1946	6.37
After one Or-Opt	2895	18	4616	25	2312	6.83
Normalised weights	0.00078	0.13636	0.00034	0.01176	0.00011	

In addition to numerical correctness, neighbourhood operators must perform the intended modifications to the route. To verify this, the routes created by each operator are visualised in Figure 6.1. Directed arcs between link cluster endpoints represent the route, with removed arcs highlighted in red and added arcs shown in green. Nodes are labelled by visit order and circled nodes indicate those directly affected by the operator. The visualisations confirm the expected behaviour. Relocate reinserts a single node elsewhere in the tour, swap exchanges two nodes, and Or-Opt relocates a contiguous block while preserving its internal order. The corresponding deterioration in metric values reported in Table 6.1 are consistent with these modifications.

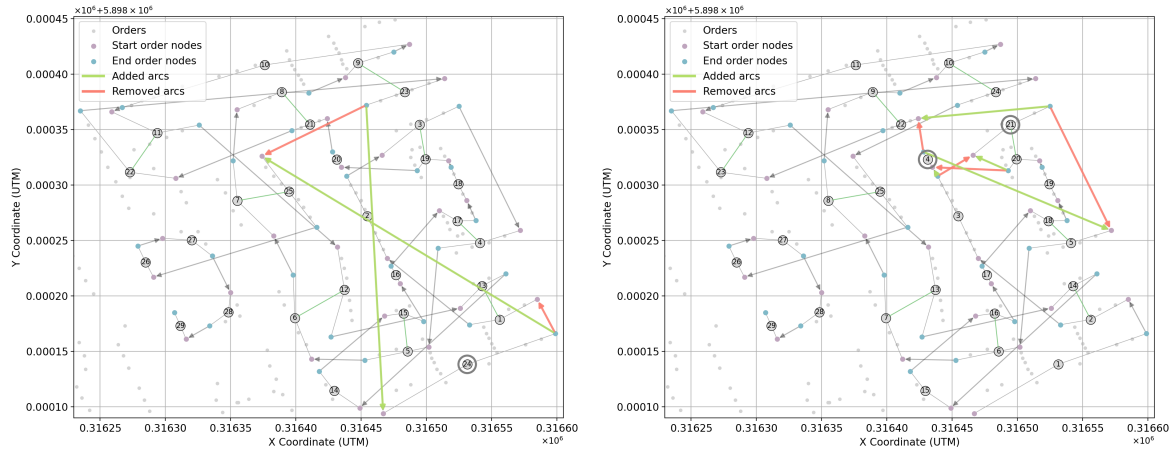

(a) Initial Route after One Random Relocate Move

(b) Initial Route after One Random Swap Move

Figure 6.1: Initial Routes after One Random Operator Move

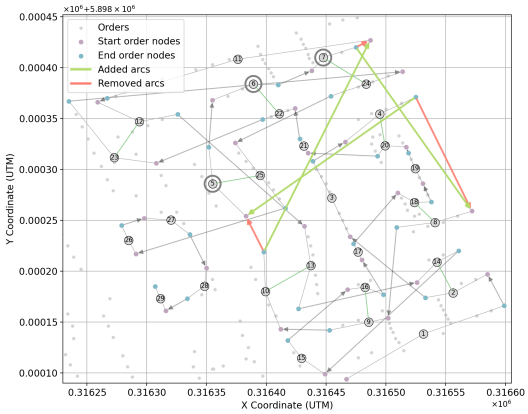

(c) Initial Route after One Random Or-Opt Move

Figure 6.1: Initial Routes after One Random Operator Move *(Continued)*

6.1.3. Verification of Search Behaviour

After verifying the correctness of the objective function and operators, this section verifies whether the SA search process behaves as intended. Even with correct metrics and operators, SA may be ineffective if the temperature schedule is incorrectly implemented. To evaluate the search behaviour, parameters that control the balance between exploration and exploitation are varied. These include the initial temperature T_0, the cooling factor α, and the number of iterations per temperature N_T.

To set up the experiment, a baseline configuration of $T_0 = 1.0$, $\alpha = 0.95$, and $N_T = 10$ is adopted. These values are selected for verification purposes on the small test set only and do not correspond to the final parameter settings used in the case studies. The cooling factor $\alpha = 0.95$ is a commonly used setting. The initial temperature T_0 is chosen to be of the same order of magnitude as the objective value to allow initial acceptance of worsening moves. The number of iterations per temperature level N_T is selected as a balance between neighbourhood exploration and runtime for the test set. For each parameter, one lower and one higher value is tested while all other parameters remain unchanged. For each configuration, the search behaviour is evaluated using plots of the current and best-so-far objective value over the iterations.

Effect of the Initial Temperature T_0

The temperature determines the acceptance probability of worsening moves. Therefore, the initial temperature controls exploration at the start of the search. A high initial temperature results in behaviour close to a random local search, whereas a low initial temperature leads to a more greedy, first-improving search. To verify this effect, T_0 is varied between 0.1, 1.0, and 5.0.

As shown in Figure 6.2, the expected effect of varying T_0 is reflected in the current and best-so-far objective values over time. A low initial temperature results in greedy behaviour with fast stabilisation of the objective value. Conversely, a high initial temperature produces objective fluctuations due to frequent acceptance of worsening moves, resulting in slower convergence but greater exploration of the solution space. The baseline configuration shows intermediate behaviour. These results are consistent with the hypothesis and confirm that the implementation responds correctly to changes in T_0.

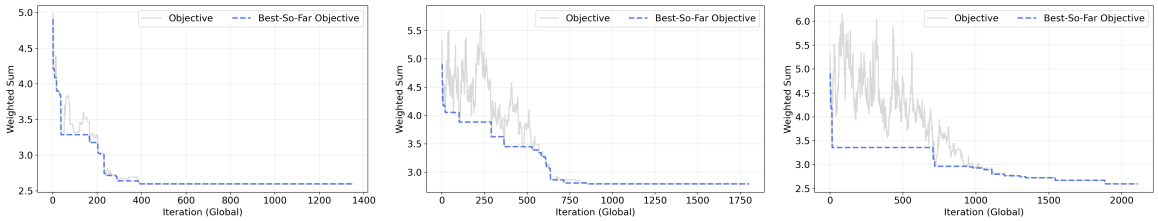

(a) Low Initial Temperature ($T_0 = 0.1$) **(b)** Baseline Initial Temperature ($T_0 = 1.0$) **(c)** High Initial Temperature ($T_0 = 5$)

Figure 6.2: Effect of the Initial Temperature T_0 on Search Behaviour

Effect of the Cooling Factor α
The cooling factor controls the rate at which the temperature is reduced. Thus, it governs the convergence speed of SA and the duration of the exploratory phase. A smaller value leads to faster cooling, causing the search to become greedy earlier. Conversely, a larger value results in slower cooling and longer exploration. To verify this effect, α is varied between 0.90, 0.95, and 0.99.

As shown in Figure 6.3, the expected effect of varying α is reflected in the current and best-so-far objective values over time. A low cooling factor leads to rapid temperature decrease, causing earlier convergence and limited exploration. In contrast, a high cooling factor results in slower cooling, which prolongs exploration and delays convergence. The baseline configuration again shows intermediate behaviour. These observations are consistent with the expected role of α in controlling the cooling schedule and confirm that the implementation responds correctly to changes in α.

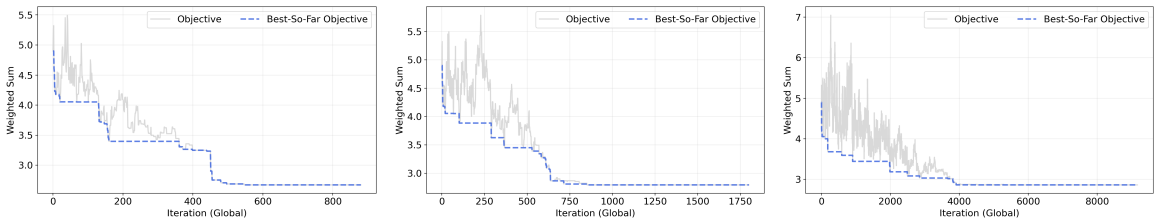

(a) Low cooling factor ($\alpha = 0.90$) **(b)** Baseline cooling factor ($\alpha = 0.95$) **(c)** High cooling factor ($\alpha = 0.99$)

Figure 6.3: Effect of the Cooling Factor α on Search Behaviour

Effect of the Number of Iterations per Temperature N_T
The number of iterations per temperature controls how extensively the neighbourhood is explored at each temperature level. A larger value allows the search to reach an equilibrium at each temperature, while a smaller value may result in insufficient sampling of the neighbourhood. Unlike T_0 and α, this parameter does not change the acceptance probabilities. To verify this effect, N_T is varied between 5, 10, and 25.

As shown in Figure 6.4, the expected effect of varying N_T is reflected in the current and best-so-far objective values over time. A small value leads to fast progression through temperature levels and early stabilisation, leading to fast convergence but limiting exploitation at each temperature. Conversely, a larger value enables more thorough exploration of the neighbourhood at each temperature level, which is reflected as prolonged objective fluctuations and slower convergence of the best-so-far objective. The baseline setting again shows intermediate behaviour. These results are consistent with the hypothesis and confirm that the implementation correctly controls the search depth per temperature.

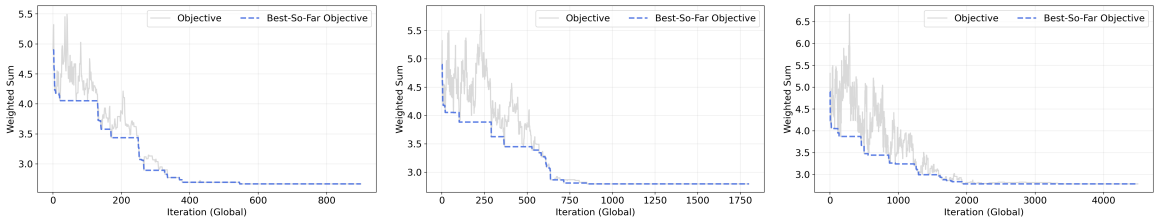

(a) Low Number of Temperature Iterations ($N_T = 5$) **(b)** Baseline Number of Temperature Iteration ($N_T = 10$) **(c)** High Number of Temperature Iteration ($N_T = 25$)

Figure 6.4: Effect of the Number of Iterations per Temperature N_T on Search Behaviour

6.2. Validation

Validation concerns practical usefulness by checking whether the right model is built. Whereas verification establishes that the framework behaves according to its specification, validation examines whether it produces routes that are meaningful for the intended purpose. In the context of side loader routing, this involves evaluating whether the framework achieves improvements in the objective function and if these improvements lead to interpretable changes in route structure.

Therefore, this section focuses on validating the parameters that control the SA search process, namely the initial temperature T_0, the cooling factor α, and the number of iterations per temperature level N_T. The goal is to identify parameter settings that lead to robust and consistent optimisation behaviour while reliably producing high-quality solutions across realistic problem instances. Subsequently, the validated parameter configuration forms the basis for the final case study evaluations used to assess the practical usefulness of the optimisation framework in Chapter 7.

6.2.1. Experimental Setup

All validation experiments are implemented in Python and performed on the 12 residential waste collection case studies introduced in Chapter 3. Using case studies rather than the test set allows the practical usefulness of the optimisation framework to be evaluated on realistic problem instances. Moreover, this ensures that the selected parameters are not tuned to a single instance but generalise across different road networks. The neighbourhood target size for the k-means clustering is set to 50 and the revisit span tolerance is fixed at 600 seconds. Initial routes were generated by solving the exact ILP model that minimises travel time with a time limit of 300 seconds using the Gurobi Optimiser (version 12). Optimal solutions were found for four instances, while the remaining case studies terminated with small optimality gaps, ranging from 0.008 to 8.6%. All experiments for each case study were executed on a separate virtual machine (VM) equipped with a 2.0 GHz 2-core processor and 8 GB of RAM.

All remaining parameters and weights are fixed to the values specified in Section 5.2, except for parameters that explicitly control SA search behaviour. A parameter sweep is conducted for different values of the parameters that control SA search behaviour, namely T_0, α, and N_T. The following values are tested, which results in 18 different configurations per case study:

- $T_0 \in \{0.5,\ 1.0,\ 2.0\}$
- $\alpha \in \{0.95, 0.99\}$
- $N_T \in \{30,\ 50,\ 70\}$

The range of cooling factor α values covers commonly used cooling factors, from relatively fast cooling to gradual cooling. The initial temperature T_0 values are selected relative to the magnitude of the weighted sum objective. For the initial solution, the objective value equals 5.5 under the fixed preference weights. Varying T_0 therefore allows assessment of more greedy versus more exploratory initial behaviour. Finally, N_T is varied to study the trade-off between exploring more candidate moves per temperature level and progressing more quickly through the temperatures.

The parameter configurations are evaluated using criteria that capture solution quality, computational effort and search behaviour. Solution quality is assessed using the best-so-far objective value, defined as the minimum weighted sum objective encountered during a run. First, heatmaps of the best-so-far objective value are constructed across the different parameter combinations. For each configuration, results are aggregated over the 12 case studies using the median and interquartile range (IQR). The median provides a robust measure of central tendency, while the IQR captures variability across instances. These measures are preferred over the mean and standard deviation, which can be disproportionately influenced by individual case studies in such a limited set of case studies.

Based on these heatmaps, a subset of promising parameter configurations is selected for more detailed inspection before selecting the best-performing parameter combination(s). For these configurations in this detailed inspection, best-so-far objective values are examined per case study using line plots, allowing differences in performance across individual datasets to be assessed. In addition, corresponding runtime plots are analysed to avoid selecting configurations that yield only small quality improvements at disproportionate computational cost.

Finally, the search behaviour of the most promising parameter configuration is examined in detail for a single case study. Two diagnostic plots are used for this purpose. The first plot shows the evolution of the current objective value, the best-so-far objective value, and the individual weighted metric contributions over the iterations, providing insight into how improvements are accumulated during the search. The second plot shows the acceptance behaviour across the temperature schedule, distinguishing between the overall acceptance rate and the acceptance rate of worsening moves. At each temperature level, the acceptance rate is computed as the fraction of accepted candidate moves over all iterations at that specific temperature level. The same procedure is applied for worsening moves, considering only

candidate moves that increase the objective value. These plots are used to validate that the selected configuration shows the intended transition from exploration to exploitation.

6.2.2. Validation of Search Behaviour

The aggregated results across all case studies are summarised in Figure 6.5, which reports the median and IQR of the best-so-far objective value for each of the 18 SA parameter combinations. The objective value of the initial travel-time-optimal solution equals 5.5. Parameter configurations for which the median best-so-far objective remains equal or close to this value, indicate that the search does not produce meaningful improvements beyond the initial solution. Therefore, such configurations are excluded from further consideration. Applying this criterion reduces the set of candidate configurations from 18 to 9. In particular, all configurations with a cooling factor of $\alpha = 0.95$ fail to consistently improve the objective value and are discarded. The remaining configurations using $\alpha = 0.99$ are analysed in more detail in the subsequent analysis of individual case studies.

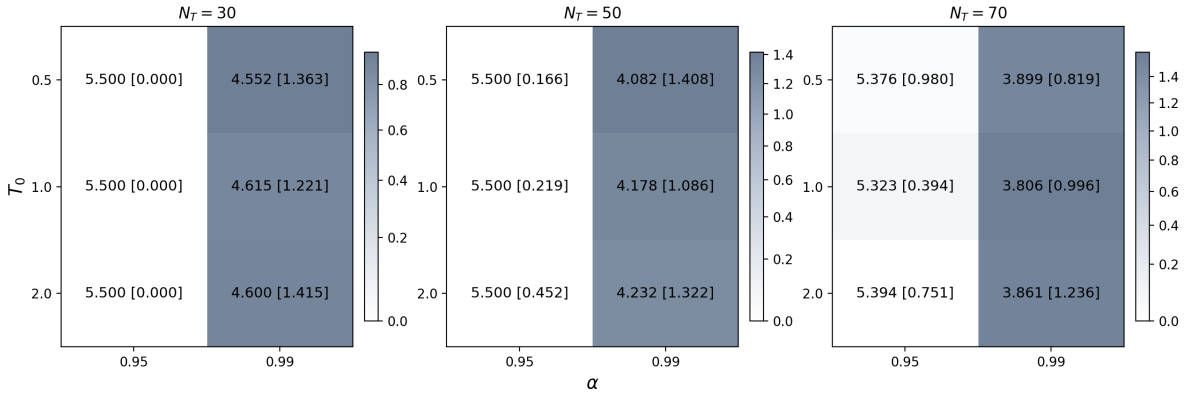

Figure 6.5: Heatmaps of Best Objective Values Across Case Studies (Median and IQR)

The remaining configurations with $\alpha = 0.99$ are analysed in more detail by inspecting the best-so-far objective value per case study across different initial temperatures T_0 and numbers of iterations per temperature N_T, as shown in Figure 6.6. For most case studies, increasing N_T leads to improved best-so-far objective values, regardless of the initial temperature. This trend is particularly pronounced for the larger case studies, which typically require more local search iterations to realise meaningful structural changes in the route. In contrast, smaller case studies already show improvements for $N_T = 30$, with diminishing marginal improvements as N_T increases further.

Across all values of T_0, increasing N_T from 30 to 50 yields a substantial improvement in solution quality, while the additional gains from increasing N_T from 50 to 70 are generally smaller. All initial temperatures are able to find improved solutions, indicating that T_0 has a limited influence on the best-so-far objective value within the tested range. One large route 2 in set 2 shows unusually strong improvements even for moderate values of N_T, suggesting instance-specific behaviour rather than a systematic effect.

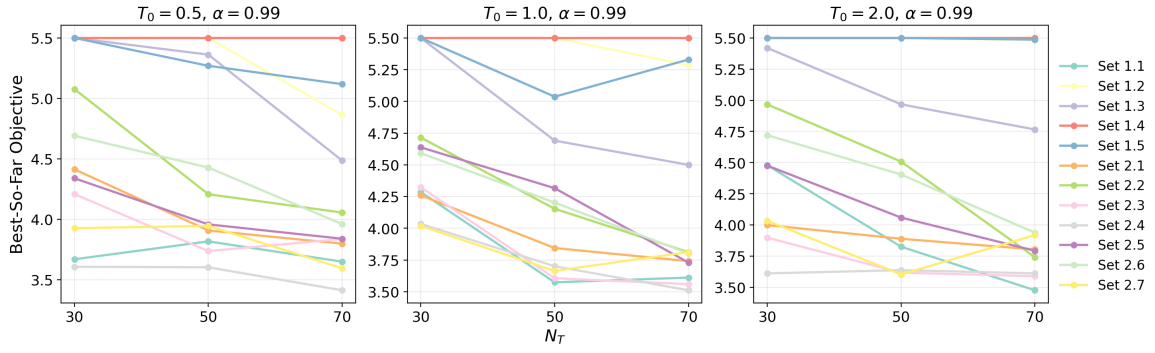

Figure 6.6: Best Objective Value per Case Study for $\alpha = 0.99$ with Varying T_0 and N_T

Based on these observations, multiple parameter configurations are selected for the final case study evaluation. The cooling factor is fixed at $\alpha = 0.99$, as it consistently yields good solution quality. The initial temperature is fixed at $T_0 = 0.5$, since it achieves comparable or better best-so-far objective values than higher temperatures while resulting in lower runtimes, as displayed in Figure 6.7.

The number of iterations per temperature N_T is selected based on case study size to balance solution quality and computational effort. An overview of the case study sizes is given in Subsection B.2.1. For case studies with about 100–200 link clusters (route 1 in set 1 and routes 3 and 4 in set 2), $N_T = 50$. For 200–300 link clusters (routes 1, 5, 6, and 7 in set 2), N_T is increased to 70. For more than 300 link clusters (routes 2, 3, 4, and 5 in Set 1 and route 2 in set 2), N_T is set to 90. Although this setting is not explicitly validated, the observed trends suggest that more iterations can improve large instances.

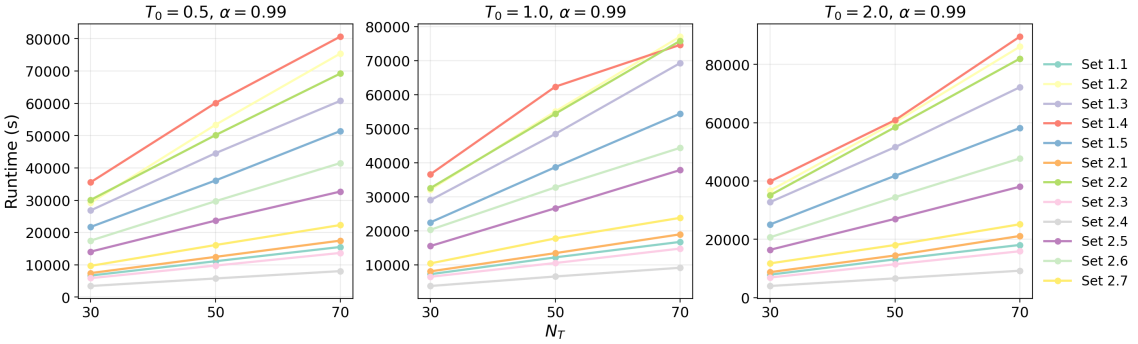

Figure 6.7: Runtime per Case Study for $\alpha = 0.99$ with Varying T_0 and N_T

Among the selected parameter configurations ($T_0 = 0.5,\ \alpha = 0.99,\ N_T \in \{50, 70, 90\}$), the search behaviour of a single configuration is examined in more detail. Results are shown for one representative case study (route 1 in set 1, United States) using the parameter combination ($T_0 = 0.5,\ \alpha = 0.99,\ N_T = 50$), as the goal is to validate that it achieves improvements in different metric and displays stable search behaviour that transitions from exploration to exploitation.

Two diagnostic plots in Figure 6.8 are used for this purpose. Figure 6.8a shows the evolution of the current objective value, the best-so-far objective value, and the weighted metric contributions over the iterations, while Figure 6.8b shows the overall acceptance rate and the acceptance rate of worsening moves across the temperature levels. At higher temperatures, worsening moves are frequently accepted, leading to increases in the weighted metrics and current objective values and reflecting exploratory behaviour. As the temperature decreases, acceptance of worsening moves becomes less likely and both the current and best-so-far objective values stabilise, indicating convergence. The weighted metric contributions illustrate the trade-off between travel time and visual attractiveness, with travel time increasing and most visual attractiveness metrics eventually decreasing relative to the initial solution. Consistently, Figure 6.8b shows high acceptance rates at early temperature levels that gradually decrease as the temperature is lowered. Together, these plots confirm that the selected configuration shows the intended transition from exploration to exploitation.

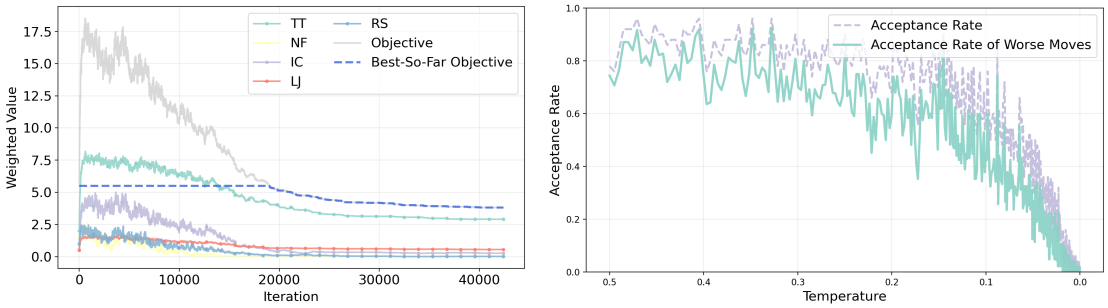

(a) Weighted Contributions Against Iterations

(b) Acceptance Rates Against Temperature Levels

Figure 6.8: Search Behaviour of Parameter Configuration ($T_0 = 0.5,\ \alpha = 0.99,\ N_T = 50$) for Route 1 in Set 1 (New Zealand)

6.3. Summary and Research Implications

This chapter verified the correctness of the optimisation framework and validated key configuration choices for its practical application. Verification was conducted on the small test set, while validation was performed on the residential waste collection case studies introduced in Chapter 3.

Verification focused on confirming the correct implementation of the optimisation framework. The weighted sum objective and weight normalisation were checked by computing these values for the initial route and for routes obtained after a single operator move. The neighbourhood operators were verified through controlled experiments and visual inspection, confirming that they produce the intended modifications. Additional experiments varied the parameters controlling the search behaviour, namely T_0, α, and N_T. The observed behaviour matched expectations: larger objective fluctuations at higher temperatures indicate exploration, while stabilisation at lower temperatures reflects exploitation. Together, these results confirm that the optimisation framework is correctly implemented and behaves as intended.

Validation focused on practical usefulness and on identifying robust parameter settings. Different configurations of T_0, α, and N_T were evaluated using the aggregated median best-so-far objective across the case studies. The results show that configurations with a cooling factor of $\alpha = 0.99$ consistently achieve meaningful improvements in the best-so-far objective value, whereas configurations with $\alpha = 0.95$ do not. Based on both solution quality and runtime per case study, three parameter configurations were selected. The initial temperature was fixed at $T_0 = 0.5$. The number of iterations per temperature level N_T was chosen as a function of case study size to balance solution quality and computational effort. Specifically, N_T is set to 50, 70, and 90 for case studies with 100–200, 200–300, and 300–400 link clusters, respectively. Consequently, $N_T = 50$ is used for route 1 in set 1 and routes 3 and 4 in set 2, $N_T = 70$ for routes 1, 5, 6, and 7 in set 2, and $N_T = 90$ for routes 2, 3, 4, and 5 in set 1 and route 2 in set 2. Detailed inspection of search and acceptance behaviour for one case study confirmed a stable transition from exploration to exploitation.

The results of this chapter justify the use of the optimisation framework in the final evaluation presented in Chapter 7. The verified implementation ensures that observed improvements can be attributed to the optimisation process itself, while the validated SA parameters provide a robust search strategy across different networks. In the next chapter, the framework is applied to all case studies and compared with the current operational routes of AMCS.

7

Results of Case Studies

This chapter presents the final evaluation of the proposed SA optimisation framework on real residential waste collection case studies. Using the verified and validated framework described in Chapter 5, the framework is applied to all 12 case studies introduced in Chapter 3. The objective is to assess how the optimisation framework performs relative to existing AMCS routing solutions, with a particular focus on improvements in visual attractiveness and the associated trade-offs with travel time.

The evaluation is structured in three parts. First, a numerical comparison is conducted across all case studies to quantify changes in travel time and visual attractiveness metrics. Second, the resulting routes are analysed visually using the road network and abstract network, with a detailed discussion provided for a single representative case study to support interpretation. Finally, a sensitivity analysis is performed to examine how variations in preference weights affect the balance between travel time and visual attractiveness for the same representative case study.

7.1. Experimental Setup

This section describes the remaining experimental setup in three parts. First, the parameter configuration used throughout this chapter is specified. To ensure a consistent evaluation, all parameters and modelling choices are fixed across all experiments. Second, the experimental setup and evaluation metrics for the numerical and visual analysis are defined. Finally, the experimental setup for the sensitivity analysis of the preference weights in the objective function is described.

All optimisation experiments use the verified and validated optimisation framework described in Chapter 6. The experiments were implemented in Python and performed on a computer equipped with a 3.6 GHz 8-core (physical) processor and 32 GB of RAM. Initial routes were obtained by solving the exact ILP model that minimises travel time using the Gurobi Optimiser (version 12).

7.1.1. Parameter Configuration

All experiments in this chapter follow the general parameter configuration described in this subsection. The neighbourhood target size used for the k-means clustering is set to 50 and the revisit span tolerance is fixed at 600 seconds. Initial routes are obtained by solving the exact ILP model that minimises travel time, with a time limit of 300 seconds. Optimal solutions were found for four instances, while the remaining case studies terminated with small optimality gaps, ranging from 0.008 to 8.6%. The parameters that control the SA search behaviour are fixed to the configuration identified as most robust in the validation study in Section 6.2. Specifically, the initial temperature is set to $T_0 = 0.5$ and the cooling factor to $\alpha = 0.99$. The number of iterations per temperature, N_T, is selected as a function of case study size in order to balance solution quality and computational effort. Concretely, N_T is set to 50, 70, and 90 for case studies with 100–200, 200–300, and 300–400 link clusters, respectively. Thus, $N_T = 50$ is used for route 1 of set 1 and routes 3 and 4 of set 2, $N_T = 70$ for routes 1, 5, 6, and 7 of set 2, and $N_T = 90$ for routes 2, 3, 4, and 5 of set 1 and route 2 of set 2. All remaining algorithmic parameters and objective function preference weights are fixed to the values specified in Chapter 5.

7.1.2. Experimental Setup for the Numerical and Visual Analysis

For the numerical analysis across case studies, two routes are evaluated for each case study, namely the travel-time-optimal route generated by the existing AMCS method and the route produced by the optimisation framework based on SA. This enables a comparison between the current operational solution and the more visually attractive alternative generated by the framework. For each route, the raw metrics are computed as described in Subsection 4.2.2, whereas the normalised weights and objective values follow Section 5.2. For the optimisation framework, these values correspond to the best route found during the run. To compare with the AMCS solution, relative changes in all metrics are computed with respect to the AMCS baseline. For a given metric M, the relative change is defined as:

$$\Delta M_{\%} = \frac{M_{\text{SA}} - M_{\text{AMCS}}}{M_{\text{AMCS}}} \times 100$$

Negative values indicate an improvement relative to the AMCS route, while positive values indicate deterioration. In addition to the individual route quality metrics, the relative change in objective value is reported to quantify overall optimisation performance. The objective value change is computed relative to the initial solution used to initialise the search procedure, with objective value F_0, as:

$$\Delta F_{\%} = \frac{F_{\text{SA}} - F_0}{F_0} \times 100$$

Moreover, the runtime of both the AMCS approach and the optimisation framework is reported to assess computational effort. To capture overall performance across the full set of case studies, summary statistics including the median, interquartile range (IQR), minimum, and maximum values are reported for all evaluation metrics. These measures are preferred over the mean and standard deviation, which can be disproportionately influenced by individual case studies in such a limited set of case studies.

To support the interpretation of the quantitative results, the routes generated by the optimisation framework are visually compared to the corresponding AMCS routes using the road network and abstract network. A detailed visual analysis is conducted for a single representative case study, while visualisations for all remaining case studies are provided for completeness.

7.1.3. Experimental Setup for the Sensitivity Analysis of Preference Weights

In addition to the final evaluation across all case studies, a sensitivity analysis is performed on the selected preference weights in the objective function, computed as explained in Section 5.2. This analysis is conducted for the same case study used in the visual analysis for consistency. Its purpose is not to exhaustively explore all preference weight combinations or fully identify the Pareto-optimal solution set, but to assess the robustness of the selected weights and illustrate how varying the relative importance of visual attractiveness with respect to travel time influences routing outcomes.

To facilitate this trade-off analysis, the visual attractiveness component of the objective is scaled uniformly relative to travel time, rather than varying all preference weights independently. Specifically, a single scaling parameter λ is introduced, while the travel time preference weight is kept fixed. The analysis considers the set $\lambda \in \{0.25, 0.5, 1, 2.0, 4.0\}$, generating a small number of scenarios that increasingly prioritise visual attractiveness. For each value of λ, the optimisation framework is executed once and the raw route quality metrics are computed as described in Subsection 4.2.2. The corresponding normalised weights and objective values follow the procedure outlined in Section 5.2 and correspond to the best-found route returned during the run. To support interpretation, the raw values of travel time and each visual attractiveness metric are plotted for the different scaling parameters.

7.2. Numerical Analysis Against AMCS Baseline

This section presents a quantitative comparison between the travel-time-optimal routes produced by the existing AMCS methods and the routes generated by the optimisation framework. It follows the setup described in Section 7.1 and the results are summarised in Table 7.1. Relative changes in route quality metrics are computed with respect to the AMCS baseline route for each case study, while relative changes in the objective value are computed with respect to the initial solution used to initialise the search procedure. All values are derived from the raw metric and objective values reported in Section E.1. The reported travel time refers to deadhead travel time, as defined in Section 4.2. Service times of link clusters represent travel within clusters and are therefore fixed and excluded from the optimisation. Runtime values for both approaches are included to highlight the increased computational cost associated with the iterative search procedure of the optimisation framework.

Runtime differences across case studies mainly arise from differences in the number of simulated annealing iterations. While all case studies use the same temperature schedule, the number of iterations per temperature level varies between 50, 70, and 90, which increases the total number of search iterations. In each iteration, ten candidate routes are generated and evaluated, requiring computation of all route quality metrics. In addition, larger case studies contain more link clusters, which likely increases the computational effort required for metric evaluation. Together, these factors lead to longer runtimes for larger instances.

For 11 of the 12 case studies, the optimisation framework identifies a route with an improved objective value compared to the initial solution. Across these case studies, the objective function decreases substantially, with a median reduction of -30.4% and a relatively small interquartile range of 11.8%. In one case study (route 4, set 1), the algorithm does not improve the initial route. Thus, the reported values correspond to this solution. This case shows that the initial solution can differ from the AMCS solution, which is also observed for other case studies as shown in Section E.1.

For the 11 case studies with an improved objective value, the optimisation framework consistently improves several aspects of visual attractiveness compared to the AMCS baseline. Neighbourhood fragmentation and revisit span are reduced for all these cases, with a median reduction of 90.2% and 92.0%, respectively. This indicates that the optimisation framework successfully reduces repeated returns to the same neighbourhood and shortens the time interval between servicing opposite kerbsides. Intra-route crossings are also reduced, but the magnitude of these improvements is lower and varies more between case studies. This is indicated by a lower median reduction of -47.7% and a larger IQR of 37.9%. This variability reflects the lower preference weight assigned to intra-route crossings compared to neighbourhood fragmentation and revisit span, resulting in less consistent improvements in route complexity than in regional compactness and local compactness specific to side loader operations.

Relative changes in local jump length show a different pattern. For all case studies, local jump length increases relative to the AMCS baseline, with a median increase of 19.2%. This reflects that local jump length and travel time are both distance-based measures derived from the spatial separation between link cluster endpoints. In contrast, the other metrics capture more structural or very specific aspects of the route. As a result, improvements in regional compactness, side-loader-specific local compactness, or route simplicity do not necessarily translate into shorter distances between successive link clusters. This highlights a trade-off between the other visual attractiveness and local jumps metrics and is consistent with the lower preference weight assigned to local jumps.

As expected, these improvements in visual attractiveness are accompanied by increased travel times. For all case studies with an improved objective value, travel time increases relative to the AMCS baseline, with a median of 40.5%. This reflects the trade-off made between efficiency and visual attractiveness. The relative reductions in objective value confirm that, within the chosen objective formulation, the improvements in visual attractiveness outweigh the associated increases in travel time.

Overall, the numerical results show that the optimisation framework generally finds routes that are more visually attractive than the AMCS baseline, while accepting increases in travel time. The results also suggest trade-offs between different aspects of visual attractiveness. The framework primarily improves regional compactness, side-loader-specific local compactnes, and route simplicity through reducing neighbourhood fragmentation, revisit spans of opposite kerbsides, and intra-route crossings. In contrast, local distance-based compactness measured by local jumps deteriorates similar to travel time. This reflects that local jump length and travel time are distance-based measures derived from spatial separation between link cluster endpoints, whereas the other three metrics capture structural and specific route characteristics that are more directly targeted by the objective function under the chosen preference weights. Improvements in route complexity are generally smaller than those in regional and side-loader-specific local compactness, consistent with the lower preference weight assigned to intra-route crossings. Within the chosen objective formulation, these findings confirm that the optimisation framework trades off efficiency for improved visual attractiveness.

Table 7.1: Relative Metric and Objective Changes for the Optimisation Framework Routes

Case Study	ΔTT (%)	ΔNF (%)	ΔRS (%)	ΔIC (%)	ΔLJ (%)	$\Delta\mathbf{F}$ (%)	Runtime AMCS (s)	Runtime SA (s)
Set 1 – Route 1	45.2	−95.7	−98.2	−55.8	6.9	−30.5	93	7681
Set 1 – Route 2	71.7	−83.3	−83.8	−19.8	34.5	−10.2	193	49 278
Set 1 – Route 3	65.8	−91.5	−92.2	−34.6	46.5	−22.7	133	39 301
Set 1 – Route 4	−1.3	21.7	22.6	−15.4	−1.0	0.0	361	49 812
Set 1 – Route 5	57.8	−88.0	−92.9	−23.9	44.9	−19.5	1156	37 591
Set 2 – Route 1	40.6	−93.3	−91.9	−72.1	21.7	−30.9	5	11 757
Set 2 – Route 2	40.5	−85.5	−94.0	−39.5	15.7	−33.8	32	59 602
Set 2 – Route 3	37.3	−94.1	−82.7	−62.5	24.3	−32.0	4	6624
Set 2 – Route 4	36.0	−92.3	−90.0	−73.1	10.2	−34.5	4	4001
Set 2 – Route 5	39.7	−88.9	−94.4	−69.7	−0.3	−30.2	13	15 777
Set 2 – Route 6	41.4	−87.5	−88.6	−38.2	19.5	−28.0	34	19 582
Set 2 – Route 7	36.4	−98.2	−94.0	−70.0	18.9	−34.7	10	11 627
Median	40.5	−90.2	−92.0	−47.7	19.2	−30.4	–	–
IQR	11.2	6.5	6.5	37.9	17.5	11.8	–	–
Min	−1.3	−98.2	−98.2	−73.1	−1.0	−34.7	–	–
Max	71.7	21.7	22.6	−15.4	46.5	0.0	–	–

Note. All metric changes ΔTT, ΔNF, ΔRS, ΔIC, and ΔLJ are reported relative to the AMCS baseline route for each case study. The objective change $\Delta\mathbf{F}$ is computed relative to the initial solution with objective value $F_0 = 5.5$.

7.3. Visual Analysis Against AMCS Baseline

While the previous subsection analysed the case studies numerically, visual attractiveness is spatial. To complement the numerical results and understand how the optimisation framework shapes the route structure, this section provides a visual comparison of the routes generated by the AMCS baseline and the optimisation framework following the setup described in Section 7.1.

Although visualisations on the abstract network are available for all 12 case studies (see Section E.2), a detailed visual analysis is presented for a single representative case study only. Providing a detailed discussion for all case studies would result in repetition, as routing patterns are largely consistent across instances. Therefore, focusing on one representative case study enables a clearer and more interpretable comparison between the baseline and optimised routes. The selected case study is route 1 of set 1 (New Zealand). This instance is representative of the entire set of case studies, as the relative changes in the metrics closely match the median values reported in Section 7.2. In addition, its descriptive statistics are consistent with the overall statistics presented in Section 3.2.

Figure 7.1 compares the visualisations of the AMCS baseline route and the route generated by the optimisation framework for route 1 of set 1. Link clusters are shown at their midpoints and labelled by visitation order, with grey arrows indicating travel direction. Revisit pairs are marked with green lines, crossing arcs in blue, and terminal nodes and arcs are omitted. Figure 7.1a and Figure 7.1b colour link clusters by neighbourhood assignment, whereas Figure 7.1c and Figure 7.1d use a colour gradient to represent the visitation sequence along the route.

Several structural differences become visible between the two routing solutions. In Figure 7.1a, the AMCS baseline route shows frequent transitions between different neighbourhoods. This results in repeated returns to previously visited neighbourhoods and several long cross-neighbourhood movements. Moreover, street segments with two kerbsides are often serviced far apart in the visitation sequence of the AMCS baseline route, which can be observed in Figure 7.1b. In addition, the AMCS baseline route displays multiple long arcs that cross several other parts of the route, indicating a relatively high number of intra-route crossings and increased route complexity.

In contrast, the route generated by the optimisation framework shows a more attractive structure. As displayed in Figure 7.1b, switching between neighbourhoods occurs less frequently than in the AMCS baseline, indicating that link clusters belonging to the same neighbourhood are visited more consecutively. Additionally, Figure 7.1d shows that large revisit spans between opposite kerbsides of the same street segment are substantially reduced. Furthermore, the number of long arcs crossing multiple other

arcs is lower, indicating a reduction in intra-route crossings. These visual observations are consistent with the numerical results reported in Section 7.2, which show strong reductions in neighbourhood fragmentation, revisit span, and intra-route crossings relative to the AMCS baseline.

Despite the observed improvements in regional compactness, side-loader-specific local compactness, and route simplicity, the optimisation framework produces routes that are less locally compact with respect to the local jumps metric. Visually, the resulting route is slightly longer in terms of directed Euclidean distance than the AMCS baseline, reflecting longer moves between successive link clusters. This behaviour corresponds to the small increase observed in the numerical analysis in Section 7.2. This trade-off arises because, under the chosen preference weights, the optimisation framework prioritises reductions in neighbourhood fragmentation, revisit spans of opposite kerbsides, and intra-route crossings over distance-based local compactness. As a result, slightly longer moves are accepted to achieve improvements in the other visual attractiveness metrics. Note that changes in travel time are not directly visible in the abstract network, as the arcs represent directed Euclidean distances rather than the actual paths taken by the vehicle on the road network.

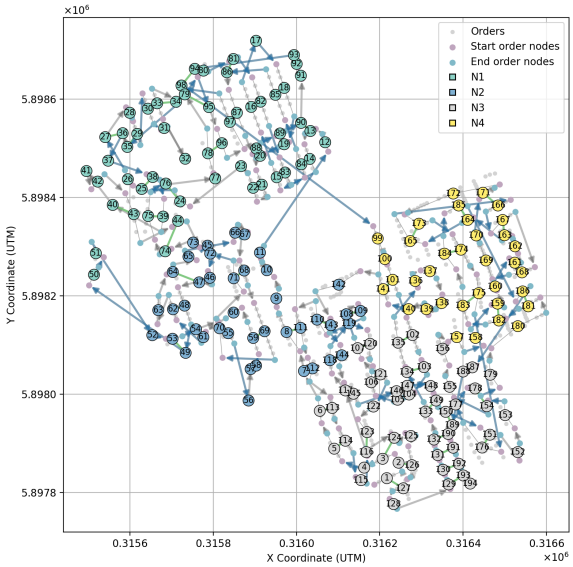

(a) AMCS Baseline Route with Neighbourhood Colouring

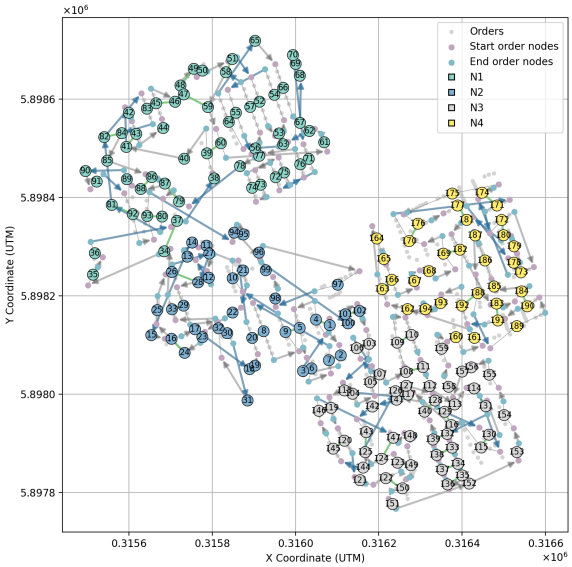

(b) Optimisation Framework Route with Neighbourhood Colouring

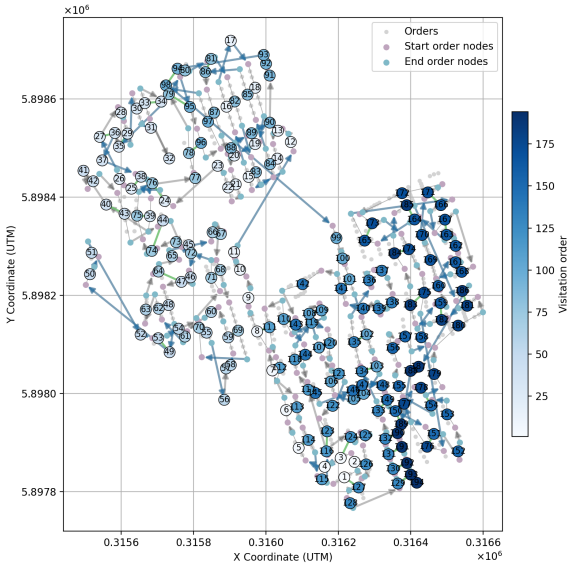

(c) AMCS Baseline Route with Visitation Order Colouring

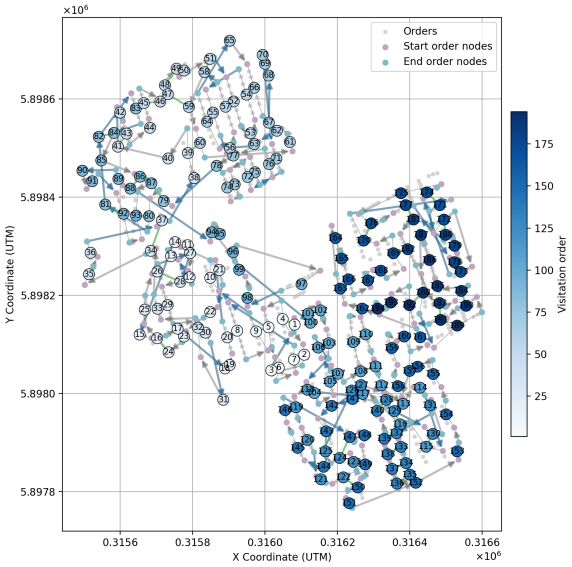

(d) Optimisation Framework Route with Visitation Order Colouring

Figure 7.1: Comparison of Routing Solutions on the Abstract Network for Route 1 of Set 1 (New Zealand)

To examine the routing behaviour in more detail, the routes are also visualised on the road network in Figure 7.2. In contrast to the abstract network views, these figures show the order sequence rather than the link cluster sequence. Additionally, order nodes are coloured by neighbourhood assignment, allowing moves between neighbourhoods to be inspected in more detail. Moreover, the first 55 order nodes are highlighted in Figure 7.2c and Figure 7.2d, illustrating differences in the first part of the route. Within this initial segment, the AMCS baseline already traverses a large part of the service area and spans multiple neighbourhoods. In contrast, the optimisation framework services orders more compact, staying within a single neighbourhood before moving to others and servicing kerbsides of the same street segment closer in time. This provides an example of the improved regional compactness and side-loader-specific local compactness achieved by the optimisation framework. Intra-route crossings are not easily identifiable in these road network visualisations, as this metric is defined based on straight-line connections between orders rather than the actual driving paths shown on the road network.

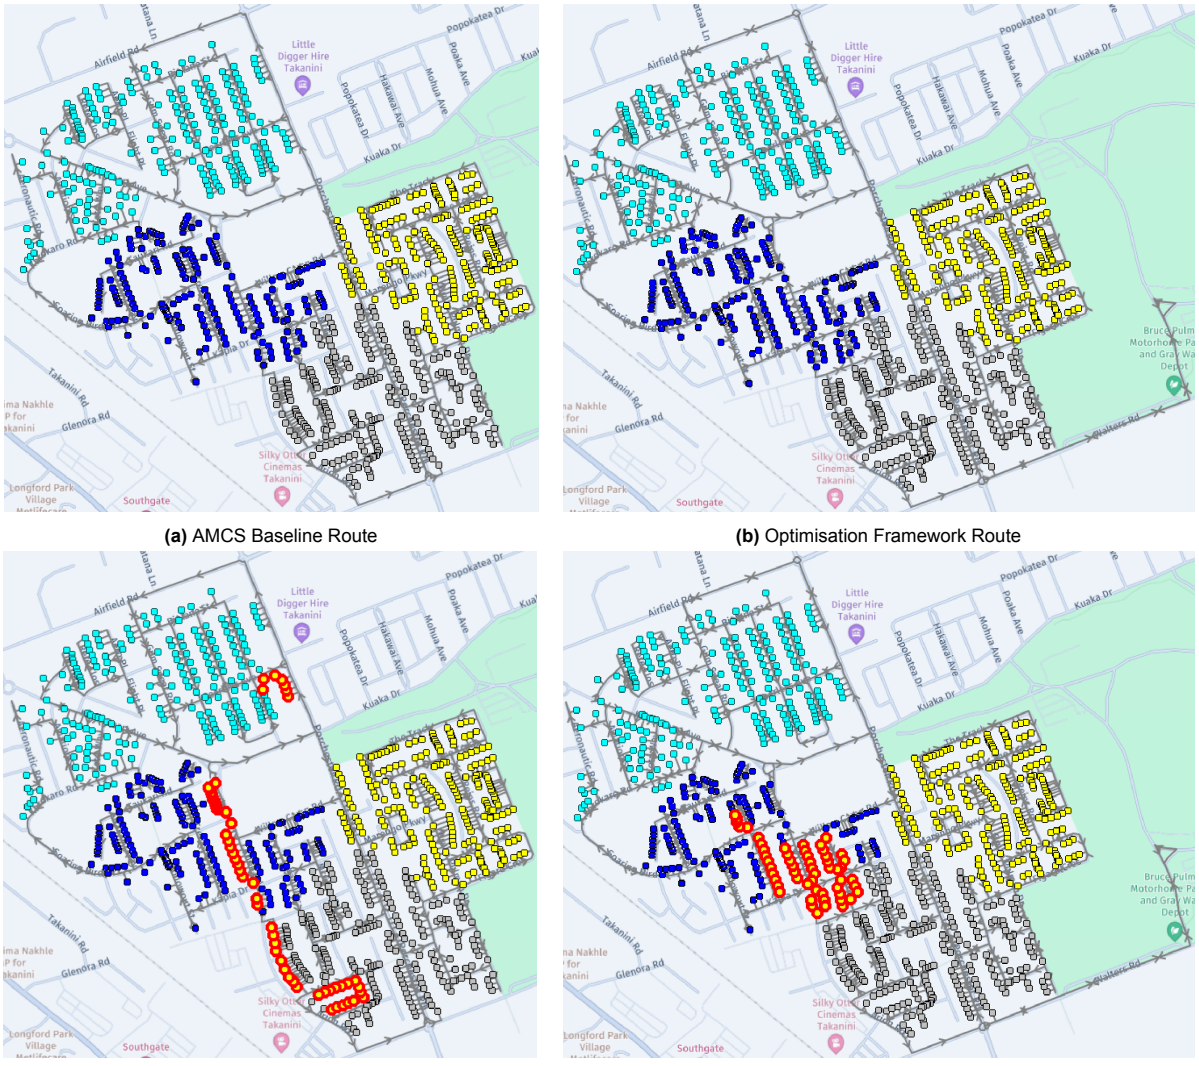

(a) AMCS Baseline Route

(b) Optimisation Framework Route

(c) AMCS Baseline Route with First 55 Orders Highlighted

(d) Optimisation Framework Route with First 55 Orders Highlighted

Figure 7.2: Comparison of Routing Solutions on the Road Network for Route 1 of Set 1 (New Zealand)

A closer inspection of a single neighbourhood in Figure 7.3 further illustrates the improvements in the side loader specific local compactness achieved by the optimisation framework. In the AMCS baseline route, several street segments are revisited after long gaps in the visitation sequence, as indicated by the order labels. In contrast, the optimisation framework services these street segments more consecutively, resulting in clearly reduced revisit spans within the neighbourhood. The street segments for which revisit behaviour improves most are highlighted in green in Figure 7.3b and contrasted with their counterparts in red in Figure 7.3a. Thus, the optimisation framework generally produces a more

compact local route structure by reducing revisit spans of opposite kerbsides for the street segments. However, not all local routing behaviour is improved. In particular, the detour highlighted in red in Figure 7.3b shows that the vehicle must turn around and traverse a larger part of the neighbourhood to move from order with sequence number 640 to 641. This is more clearly reflected in Figure 7.2b. This shows that local inefficiencies can still occur despite overall improvements.

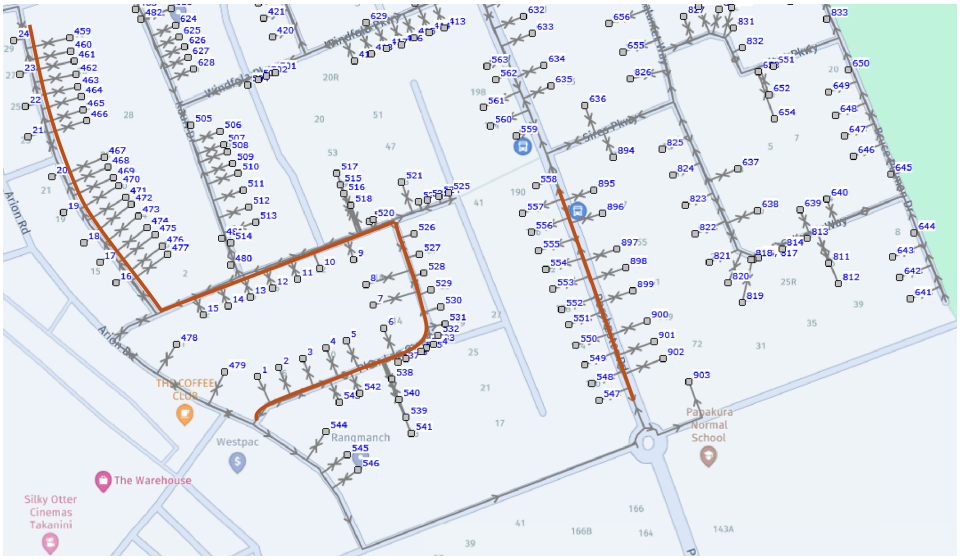

(a) AMCS Baseline Route with Highlighted Local Route Behaviour

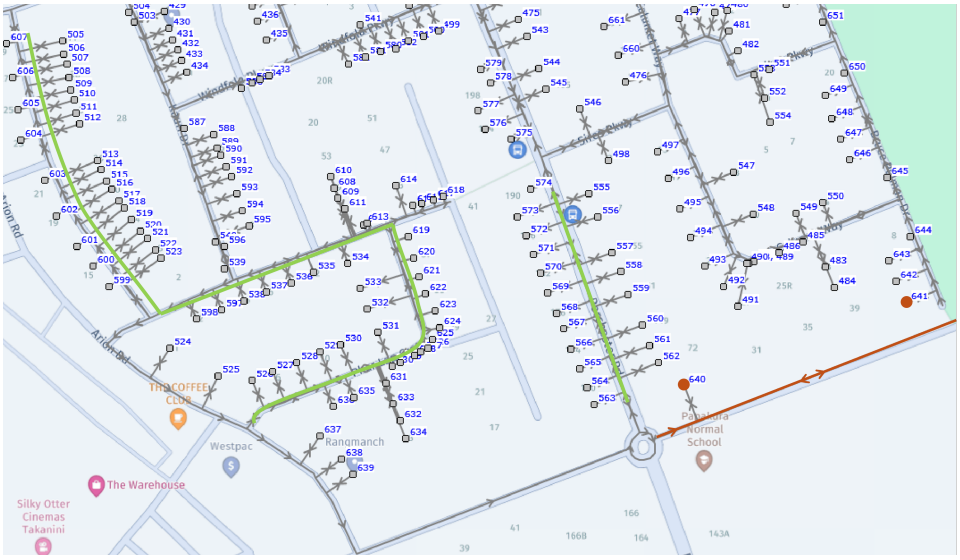

(b) Optimisation Framework Route with Highlighted Local Route Behaviour

Figure 7.3: Comparison of Routing Solutions on the Road Network for Neighbourhood 3 in Route 1 of Set 1 (New Zealand)

A second example of the route within a neighbourhood is shown in Figure 7.4. As before, coloured street segments highlight differences in revisit behaviour along the route. In the AMCS baseline route in Figure 7.4a, the street segment highlighted in red shows a large revisit span, while the adjacent street segment highlighted in green is serviced consecutively. In contrast, the optimisation framework route in Figure 7.4b services this same street segment with a substantially reduced revisit span, but does not visit the entire street as consecutively as the baseline route. This behaviour is not necessarily undesirable, as the segment is serviced consecutively with the adjacent street through the corner.

This example illustrates that improvements in visual attractiveness are not optimised for all objectives. While the optimisation framework reduces revisit spans on specific street segments, this may come at the expense of local jumps elsewhere. In addition, both routes show that the neighbourhood is re-entered after a longer interval, as indicated by the visitation sequence of the orders on the right-hand

side of the figures. These behaviours reflect the trade-offs inherent in the multi-objective formulation, in which improvements in one aspect of route structure can lead to less favourable behaviour in another.

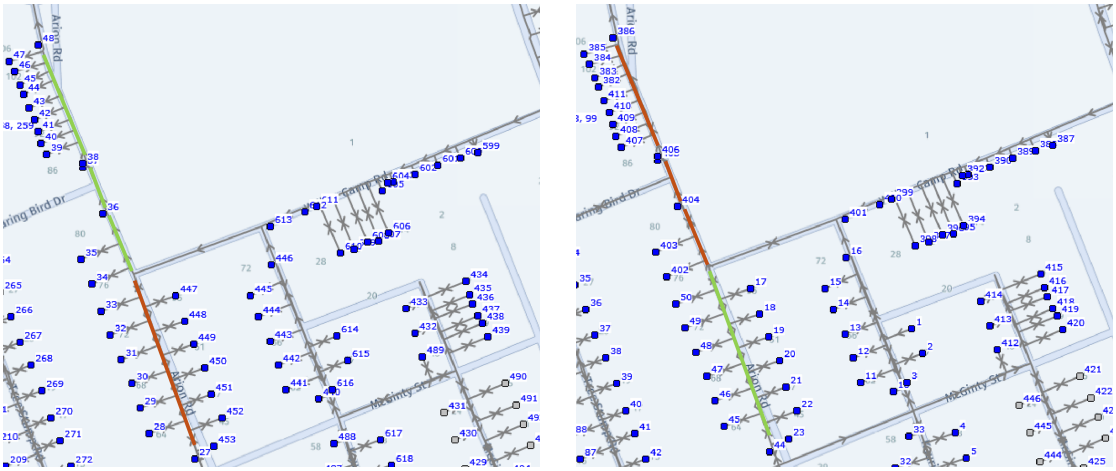

(a) AMCS Baseline Route with Highlighted Local Route Behaviour

(b) Optimisation Framework Route with Highlighted Local Route Behaviour

Figure 7.4: Comparison of Routing Solutions on the Road Network for Neighbourhood 2 in Route 1 of Set 1 (New Zealand)

An example where neighbourhood coherence is achieved in the abstract network but not fully reflected on the road network is shown in Figure 7.5. In both routes, the vehicle briefly crosses the light blue neighbourhood when servicing the dark blue neighbourhood. In the AMCS baseline route shown in Figure 7.5a, the transitions from the red-coloured orders 186 to 187 and from 194 to 195 require switching between neighbourhoods. As a result, they are penalised. In contrast, the transition from the orange-coloured orders 77 to 78 in the optimisation framework route shown in Figure 7.5b crosses the light blue neighbourhood on the road network. However, this move does not result in a neighbourhood switch in the abstract network and is therefore not penalised. This behaviour is not necessarily undesirable, as the route only passes along the outer boundary of a small neighbourhood rather than completely crossing it. However, this example highlights the sensitivity of the neighbourhood fragmentation metric to the underlying neighbourhood definition and illustrates a limitation of using link cluster representations with straight-line connections instead of the actual road network.

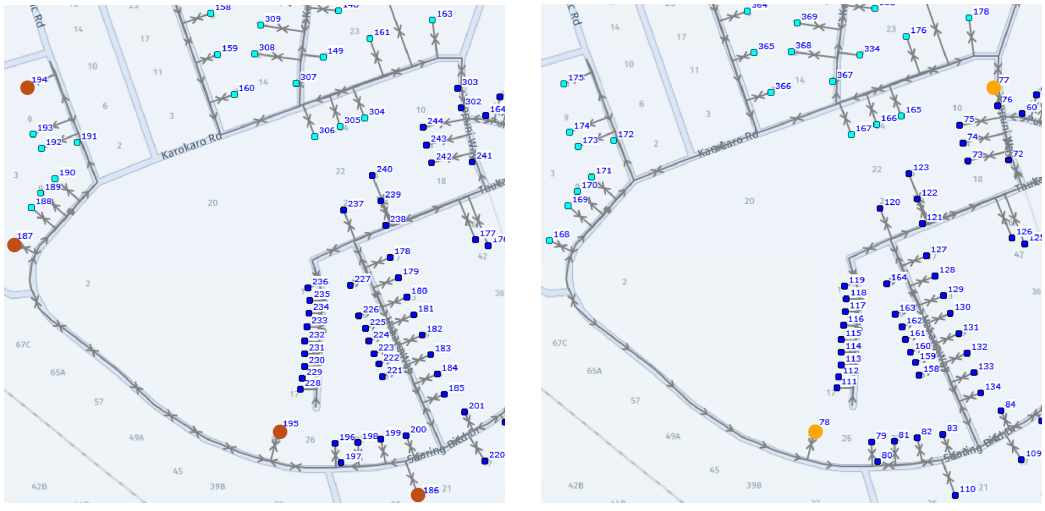

(a) AMCS Baseline Route Illustrating Penalised Neighbourhood Transitions

(b) Optimisation Framework Route Illustrating Unpenalised Neighbourhood Transitions

Figure 7.5: Example of Neighbourhood Boundary Effects on the Road Network for Route 1 of Set 1 (New Zealand)

A limitation of the underlying grouping of orders into link clusters is illustrated in Figure 7.6. In this

example, orders located along the same street are assigned to different neighbourhoods in both routing solutions. This behaviour partly arises from the use of k-means clustering to create neighbourhoods of link clusters, which does not enforce alignment with the road network. In addition, the underlying link cluster and revisit pair definitions contribute to this effect. Orders with sequence numbers 659–666 and 667–674 belong to different link clusters, even though they appear to be located on the same street segment. Additionally, these link clusters are not marked as revisit pairs. As a result, the neighbourhood construction procedure is allowed to separate them into different neighbourhoods. In this case, both the AMCS baseline and optimisation framework still service the street in a logical order. However, in other situations this separation could potentially lead to revisiting the same street multiple times. This example highlights the importance of both the neighbourhood definition and the grouping of orders into link clusters and revisit pairs when modelling visual attractiveness.

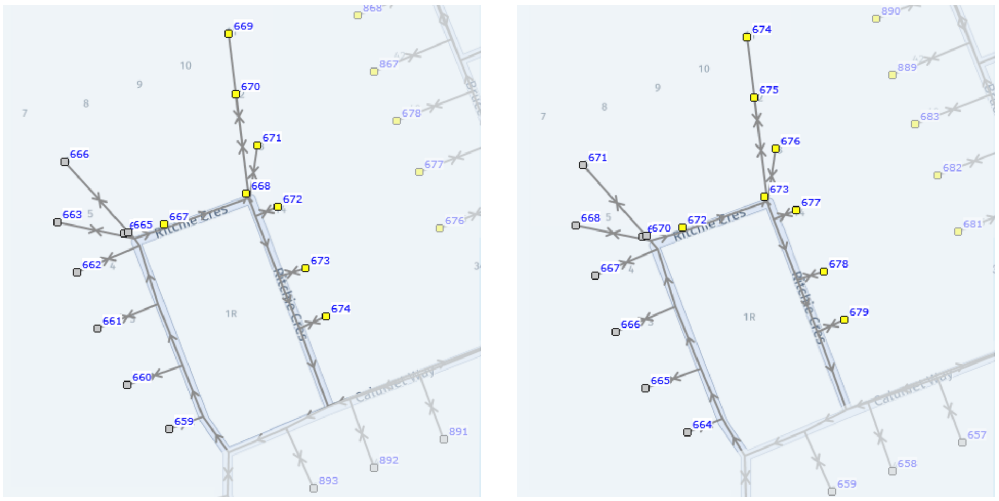

(a) AMCS Baseline Route Illustrating Orders Split Across Neighbourhoods

(b) Optimisation Framework Route Illustrating Orders Split Across Neighbourhoods

Figure 7.6: Comparison of Routing Solutions on the Road Network for Neighbourhood 4 in Route 1 of Set 1 (New Zealand)

Overall, the visual analysis for route 1 of set 1 confirms the numerical findings by illustrating how the optimisation framework reshapes the route structure relative to the AMCS baseline. Across both the abstract and road network representations, the route produced by the optimisation framework shows less neighbourhood fragmentation, reduced revisit spans along the same street segments, and fewer large intra-route crossings. This results in routes with improved regional compactness, side-loader-specific local compactness, and route simplicity, while accepting modest increases in the total directed Euclidean distance and travel time. This confirms that improvements in one metric may occur at the expense of another, as the framework balances multiple objectives through a weighted sum.

At the same time, the visual analysis highlights several limitations inherent to the chosen modelling abstractions and objective formulation. The practical performance of the neighbourhood fragmentation metric depends on the definition of neighbourhoods and their alignment with the underlying road network. For example, incorrect neighbourhood boundaries in the road network can lead to wrongly penalised or unpenalised neighbourhood switches. In addition, the grouping of orders into link clusters and the definition of revisit pairs influence local routing behaviour and could potentially lead to suboptimal sequencing of streets segments. Taken together, these observations show that while the optimisation framework substantially improves overall visual attractiveness compared to the AMCS baseline, its performance remains shaped by modelling choices and trade-offs inherent to multi-objective route optimisation. In particular, balancing partially conflicting metrics may occasionally result in locally unintuitive moves when improving one aspect of route structure requires sacrificing another.

7.4. Sensitivity Analysis of Preference Weights

This section investigates the sensitivity of the final routing outcomes to the choice of preference weights used in the objective function. While the final evaluation relies on a fixed set of weights, practical

use of the optimisation framework may require adjusting the relative importance of travel time and visual attractiveness. For example, some municipalities or customers may prefer visual attractiveness more over travel time efficiency than others. Therefore, this sensitivity analysis provides insight into the robustness of the chosen weights and illustrates how varying the relative importance of visual attractiveness against travel time affects routing outcomes.

The analysis is conducted by generating alternative routes through varying the preference weights using the optimisation framework under the experimental setup described in Section 7.1. To maintain interpretability and avoid redundancy, the analysis is restricted to a single case study, namely route 1 of set 1 (New Zealand). This case study is representative of the full set of case studies, as the relative changes in evaluation metrics closely match the median values reported in Section 7.2. Additionally, its descriptive statistics are consistent with those presented in Section 3.2. As the same case study was also used in the visual analysis in Section 7.3, this choice ensures clear interpretation of the results.

Table 7.2 summarises the raw route quality metrics obtained for different values of the preference weight scaling parameter λ. Lower values of λ prioritise travel time efficiency, whereas higher values place increasing emphasis on visual attractiveness. For each configuration, the table reports both the initial objective value F_0 of the travel-time-optimal route used to initialise the search procedure and the best-found objective value F returned by the optimisation framework. This allows for inspection of whether and to what extent the objective is improved for each value of the scaling parameter. For completeness, the table also includes the case $\lambda = 0$. In this case, the optimisation framework is not executed and the route corresponds directly to the travel-time-optimal solution obtained by solving the exact model that minimises travel time, as described in Section 7.1.

For $\lambda = 0.25$, the optimisation framework is executed but the objective value is not improved relative to the initial solution, indicating that the reported metrics correspond to the travel-time-optimal route. Increasing λ to 0.5 results in a modest objective improvement and a strong reduction in neighbourhood fragmentation and revisit span. Further increases in λ lead to progressively larger objective improvements and reductions in neighbourhood fragmentation, revisit span, and intra-route crossings. As expected, these improvements in visual attractiveness are accompanied by an increase in travel time. Runtime remains stable across all experiments, indicating that changes in the preference weights do not significantly affect computational effort for this case study.

Table 7.2: Raw Metrics for the Sensitivity Analysis of Preference Weights (Route 1, Set 1)

λ	TT (s)	NF	RS (s)	IC	LJ (m)	$\mathbf{F}$	$\mathbf{F_0}$	Runtime (s)
0.00	6112	48	55 205	70	11 065	2.000	2.000	300
0.25	6112	48	55 205	70	11 065	2.875	2.875	10 912
0.50	8414	0	4720	57	13 949	3.315	3.750	10 829
1.00	8892	2	876	38	12 321	3.816	5.500	10 829
2.00	9208	0	2980	19	12 199	4.495	9.000	10 813
4.00	10 112	0	372	12	11 059	5.678	16.000	10 799

$\lambda = 0$ corresponds to the initial solution obtained by solving the exact model that minimises travel time and is included for reference.

To visualise the trade-offs between travel time and visual attractiveness metrics, Figure 7.7 presents trade-off plots based on the values in Table 7.2. In each plot, the horizontal axis represents the total travel time, while the vertical axis shows the raw total value of a single visual attractiveness metric. Each point corresponds to a solution obtained for a specific value of λ. For clarity, the plots also include the case $\lambda = 0$, which represents the travel-time-optimal route obtained from the exact ILP model without executing the simulated annealing procedure.

The trade-off plots show a clear and consistent pattern. Increasing λ leads to strong reductions in neighbourhood fragmentation and revisit span at the cost of increased travel time. The most pronounced improvements in the neighbourhood fragmentation and revisit span metrics occur when increasing λ from 0.25 to 0.5, after which additional improvements become much smaller. Intra-route crossings follow a similar trend, but improve more gradually. This is consistent with their lower preference weight relative to neighbourhood fragmentation and revisit span.

In contrast, the local jumps metric does not always decrease with increasing λ and shows non-monotonic behaviour. This is consistent with earlier findings in Section 7.2 and Section 7.3 and reflects the distance-based nature of local jumps, which is more closely aligned with travel time than with the aspects captured by the other metrics. As a result, improvements in neighbourhood fragmentation, street segment revisit behaviour, and route simplicity do not necessarily translate into shorter local jumps. Under the chosen preference weights, the framework therefore prioritises improvements in neighbourhood fragmentation, revisit span, and intra-route crossings over reductions in the local jump metric.

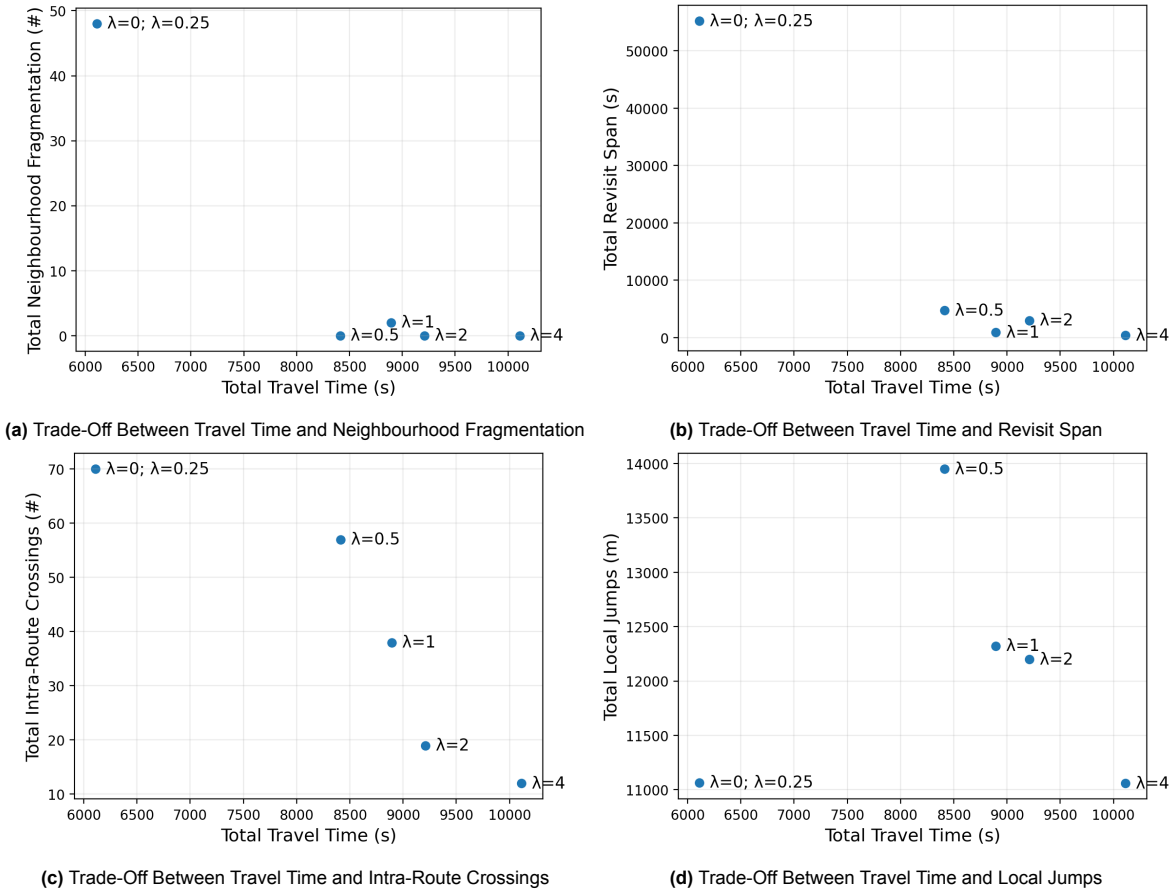

(a) Trade-Off Between Travel Time and Neighbourhood Fragmentation

(b) Trade-Off Between Travel Time and Revisit Span

(c) Trade-Off Between Travel Time and Intra-Route Crossings

(d) Trade-Off Between Travel Time and Local Jumps

Figure 7.7: Trade-Offs Between Travel Time and Visual Attractiveness Metrics for the Sensitivity Analysis (Route 1, Set 1)

Route visualisations are shown for $\lambda = 0.5$, 1.0, 2.0, and 4.0 in Figure 7.8. These configurations represent the first setting with objective improvement, the balanced default configuration, and two increasingly visually dominated settings. The route for $\lambda = 1.0$ shown in Figure 7.8b has already been analysed in Section 7.3 and is therefore treated as the baseline configuration in this sensitivity analysis.

For $\lambda = 0.5$ in Figure 7.8a, the route already exhibits some improvements in regional structure compared to the baseline. In particular, all neighbourhoods are serviced completely before moving on to the next. However, the route still contains many large revisit spans, long local jumps and intra-route crossings. As a result, local route structure within neighbourhoods remains fragmented, with several street segments serviced far apart in the visitation sequence. For example, the street segments shown at the bottom of Figure 7.8a, with visitation sequence numbers 149 to 152 and 169 to 172, are serviced further apart in the visitation sequence compared to the baseline route.

Compared to the baseline, increasing λ to 2.0 in Figure 7.8c results in more coherent servicing of neighbourhoods and fewer intra-route crossings due to the removal of several long arcs that previously crossed multiple other route segments. At the same time, the revisit span deteriorates, indicating that some revisit pairs are serviced further apart in the visitation sequence. Local compactness on the road network also decreases, as certain street segments are revisited multiple times in order to reduce neighbourhood fragmentation or local jumps. For example, link clusters with visitation numbers 141

to 144 and 190 to 191 illustrate repeated visits to the same street side to achieve reductions in these metrics. This behaviour explains the observed increase in travel time for this configuration.

For $\lambda = 4.0$ in Figure 7.8d, the route exhibits strong regional and local temporal compactness, with neighbourhoods largely serviced in a single block, revisit pairs visited close together, and successive link clusters in the route that are spatially close. In addition, the route contains very few intra-route crossings, due to more continuous local servicing of link clusters. However, achieving this level of visual attractiveness requires substantial detours, resulting in a substantial increase in travel time. Moreover, the visual attractiveness metrics are defined on an abstract network rather than the actual road network. Since travel time is the only metric that accounts for the road network through its underlying fastest-path computation, very large values of λ increase the risk of undesirable behaviour on the road network, such as repeated visits of the same streets or non-penalised movements through other neighbourhoods as previously identified in Figure 7.5.

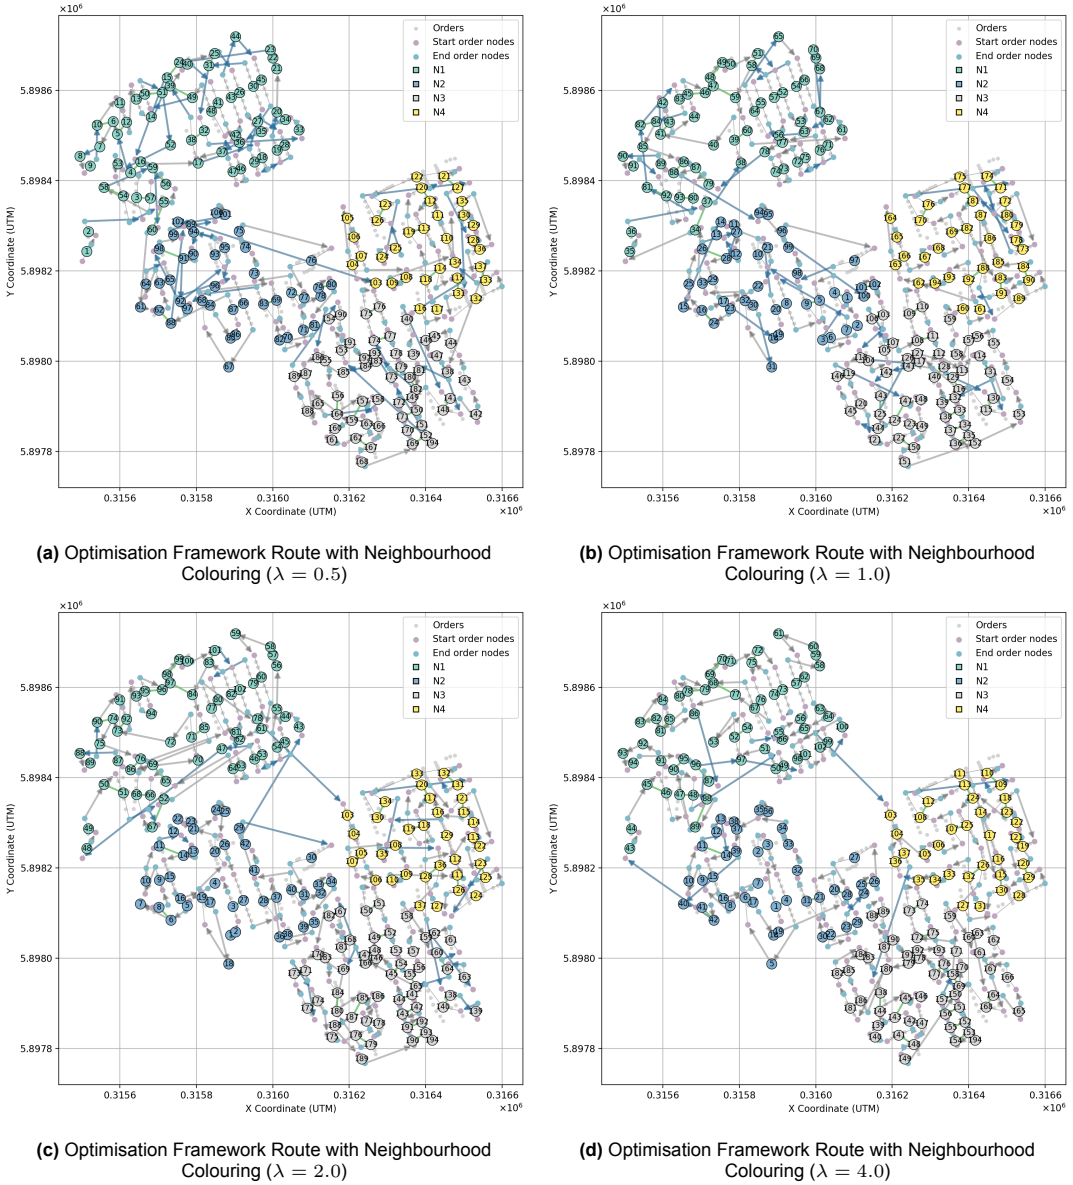

(a) Optimisation Framework Route with Neighbourhood Colouring ($\lambda = 0.5$)

(b) Optimisation Framework Route with Neighbourhood Colouring ($\lambda = 1.0$)

(c) Optimisation Framework Route with Neighbourhood Colouring ($\lambda = 2.0$)

(d) Optimisation Framework Route with Neighbourhood Colouring ($\lambda = 4.0$)

Figure 7.8: Comparison of Routing Solutions on the Abstract Network for the Sensitivity Analysis (Route 1, Set 1)

To better understand the transition from the initial solution optimal for travel time to the first improved solution, additional experiments were performed with intermediate values $\lambda \in \{0.30, 0.35, 0.40, 0.45\}$. These values allow for a more detailed analysis of how the optimisation framework behaves when

visual attractiveness receives only a small increase in importance relative to travel time. Table 7.3 reports the raw route quality metrics for these additional values of λ. The results confirm that $\lambda = 0.30$ still corresponds to the travel-time-optimal solution, as the optimisation framework does not improve the objective value relative to the initial route. This indicates that travel time remains dominant in the objective function for very small values of λ.

For larger values of λ within this range, the optimisation framework finds alternative routes that slightly improve the objective value relative to the initial solution. However, the changes in both the objective value and the individual metrics remain relatively small compared to the larger step observed beyond $\lambda = 0.50$ in the previous analysis. The runtime remains stable across all experiments.

Table 7.3: Raw Metrics for the Additional Sensitivity Analysis of Preference Weights (Route 1, Set 1)

λ	TT (s)	NF	RS (s)	IC	LJ (m)	$\mathbf{F}$	$\mathbf{F_0}$	Runtime (s)
0.00	6112	48	55 205	70	11 065	2.000	2.000	300
0.25	6112	48	55 205	70	11 065	2.875	2.875	10 912
0.30	6112	48	55 205	70	11 065	3.050	3.050	10 890
0.35	8183	8	3135	69	13 689	3.170	3.225	10 837
0.40	7758	12	8457	34	12 562	3.074	3.400	10 864
0.45	8396	2	3121	51	12 969	3.229	3.575	10 836
0.50	8414	0	4720	57	13 949	3.315	3.750	10 829

$\lambda = 0$ corresponds to the initial solution obtained by solving the exact model that minimises travel time and is included for reference.

The trade-off plots in Figure 7.9 show that the relationship between the metrics is less consistent within this narrow range of relatively low λ values. In contrast to the clearer trends observed in the initial sensitivity analysis shown in Figure 7.7, the solutions in this range do not follow a consistent pattern. For example, the configuration with $\lambda = 0.35$ achieves lower neighbourhood fragmentation and revisit span than the configuration with $\lambda = 0.40$. At the same time, the $\lambda = 0.40$ solution substantially reduces intra-route crossings while maintaining relatively low travel time and local jumps compared to the other configurations in this range. Similarly, the solution for $\lambda = 0.45$ has fewer intra-route crossings and lower travel time and local jumps than the solution for $\lambda = 0.50$.

This behaviour is likely explained by the interaction between the weighted objective and the stochastic local search procedure. At low values of λ, travel time still dominates the objective function, while the visual attractiveness metrics mainly act as secondary criteria. As a result, small variations in the search may lead to solutions that prioritise improvements in different metrics. Consequently, the resulting routes may show improvements in one metric while slightly worsening others. The local jumps metric continues to show non-monotonic behaviour. This is consistent with its lower preference weight and distance-based nature, which aligns it more closely with travel time than with the structural aspects captured by the other visual attractiveness metrics.

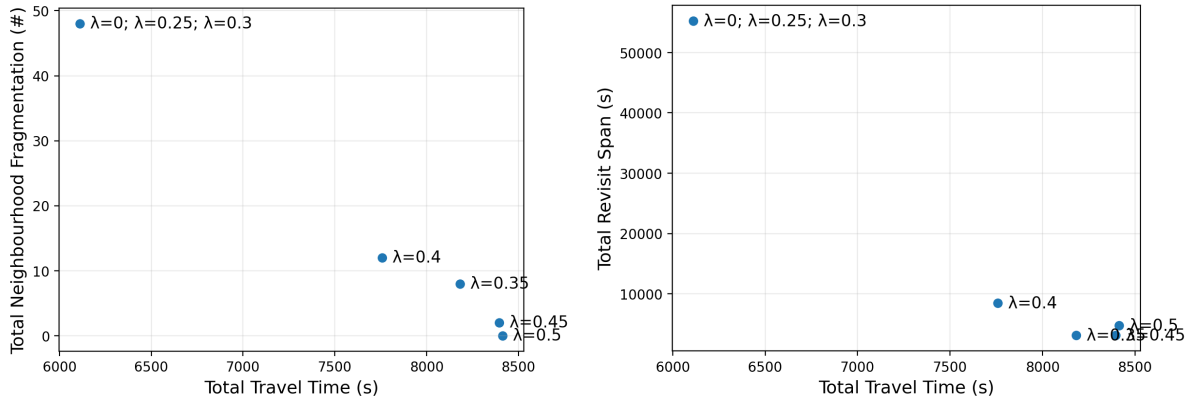

(a) Trade-off Between Travel Time and Neighbourhood Fragmentation

(b) Trade-off Between Travel Time and Revisit Span

Figure 7.9: Trade-Offs Between Travel Time and Visual Attractiveness Metrics for the Additional Sensitivity Analysis (Route 1, Set 1)

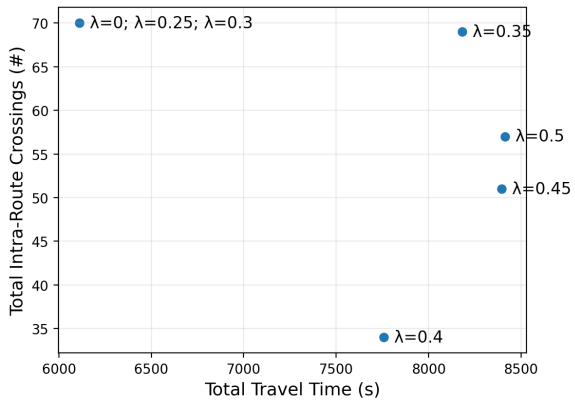

(c) Trade-off Between Travel Time and Intra-Route Crossings

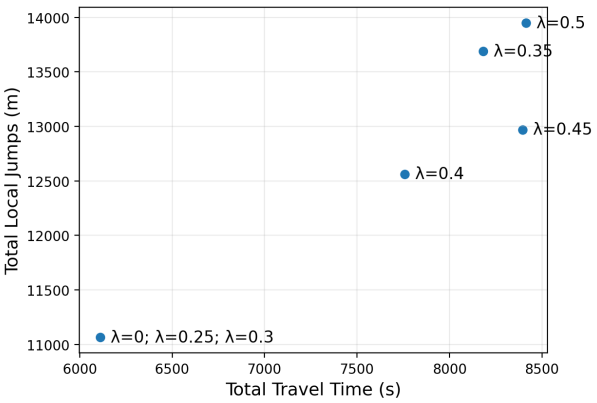

(d) Trade-off Between Travel Time and Local Jumps

Figure 7.9: Trade-Offs Between Travel Time and Visual Attractiveness Metrics for the Additional Sensitivity Analysis (Route 1, Set 1)

Taken together, the results show that moderate values of the preference weight scaling parameter λ, such as $\lambda = 1$, already achieve substantial improvements in visual attractiveness compared to the initial solution that is optimal for travel time. The additional analysis for intermediate values between $\lambda = 0.25$ and $\lambda = 0.5$ further shows that very small values of λ lead to less consistent behaviour. In this range, travel time dominates the weighted objective function and the visual attractiveness metrics mainly act as secondary criteria. As a result, the stochastic local search procedure may prioritise improvements in different metrics across runs, leading to solutions that do not follow a clear trend.

Increasing λ beyond this range yields progressively smaller improvements in visual attractiveness of the route on the abstract network, while causing a substantial increase in travel time. Consequently, this creates a higher risk of undesirable behaviour on the road network, as travel time is the only metric that accounts for the underlying road network.

These findings also illustrate the trade-offs inherent to multi-objective optimisation, where improvements in one metric may occur at the expense of another. Therefore, determining an appropriate balance between visual attractiveness and efficiency requires careful selection of preference weights and may depend on the characteristics and priorities of the operational context. Overall, the observations indicate that the selected preference weights provide a reasonable balance between travel time efficiency and visual attractiveness for this case study, while allowing the optimisation framework to be adapted to different operational preferences.

8

Discussion and Recommendations

This chapter discusses the main findings of the study in relation to the research question and the existing literature. It begins with a brief summary and interpretation of the results, highlighting the contributions to the literature and practice. This is followed by a discussion of the limitations of the modelling approach and optimisation framework, together with an explanation of why the results remain meaningful within the scope of this study. Finally, recommendations for future research and practice are presented to address the limitations.

8.1. Discussion of the Results

This section discusses the results presented in Chapter 7 in relation to the research objective and existing literature. First, the results are interpreted to assess how visual attractiveness can be quantified and integrated alongside travel time in a single-route optimisation framework for residential kerbside waste collection that accommodates double traversal. The results are compared with the currently operational routes of AMCS to evaluate whether this can be achieved without severely compromising travel time efficiency. Second, the implications of these findings are discussed from both research and practical perspectives.

8.1.1. Interpretation of the Results

This study investigates whether visual attractiveness can be quantified and incorporated into a single-route optimisation framework that accommodates double traversal while balancing travel time efficiency. To achieve this, several visual attractiveness metrics were proposed, analysed, and incorporated with travel time in a weighted sum objective and solved using SA. The relative importance of the different objectives is determined through preference weights, which guide the balance between efficiency and visual attractiveness. These metrics capture complementary aspects of visual attractiveness, including local compactness, regional compactness, and route complexity. Local compactness is captured by the local jumps and revisit span metrics, regional compactness is measured through the neighbourhood fragmentation metric, and complexity is expressed through the intra-route crossings metric.

The numerical and visual results indicate that visual attractiveness metrics influence route structure. In 11 of the 12 case studies, the optimisation framework produced a route that improved visual attractiveness while accepting modest increases in total travel time. The routes generated by the optimisation framework differ substantially from the travel-time-optimal routes produced by AMCS. Across the abstract and road network representations, the optimisation framework routes show less neighbourhood fragmentation, reduced revisit spans along the same street segments, and fewer large intra-route crossings. The reductions observed in these metrics for the case studies demonstrate that the optimisation framework produces routes that consistently improve regional compactness, local compactness specific to side loaders, and route simplicity. Improvements in route complexity are generally smaller than those in neighbourhood fragmentation and revisit span, consistent with its lower preference weight.

Both the numerical and visual results indicate that visual attractiveness metrics influence route structure. In 11 of the 12 case studies, the optimisation framework identified a route that improved visual

attractiveness while accepting modest increases in travel time. The routes generated by the optimisation framework differ substantially from travel-time-optimal routes produced by AMCS. Across the abstract and road network representations, the optimisation framework routes showed less neighbourhood fragmentation, reduced revisit spans on the same street segments, and fewer large intra-route crossings. These reductions demonstrate that the optimisation framework produces routes that consistently improve regional compactness, local compactness specific to side loaders, and route simplicity. Improvements in route complexity were generally smaller than those in neighbourhood fragmentation and revisit span, consistent with its lower preference weight.

In contrast, travel time increased in all case studies, reflecting the expected trade-off between operational efficiency and visual attractiveness. In addition, local distance-based compactness measured by the local jumps metric deteriorated. This reflects that local jump length and travel time are both distance-based measures derived from the spatial separation between link cluster endpoints. In contrast, the other metrics capture more structural or very specific aspects of the route. As a result, improvements in regional compactness, local compactness specific to side loaders, or route simplicity do not necessarily translate into shorter distances between link clusters on the route. This is consistent with the lower preference weight assigned to local jumps and highlights the trade-off between local jumps and other visual attractiveness metrics.

The sensitivity analysis of preference weights further illustrated the trade-offs between efficiency and visual attractiveness. The results confirm that the selected preference weights provide a reasonable compromise between travel time efficiency and visual attractiveness for the analysed case study. Increasing the relative importance of the visual attractiveness metrics using a scaling factor leads to reductions in neighbourhood fragmentation, revisit spans of opposite kerbsides, and intra-route crossings. However, these improvements become progressively smaller while travel time increases more substantially. In addition, the local jumps metric does not consistently improve and often deteriorates, reflecting its distance-based nature and closer alignment with travel time than with the other metrics. Additional experiments show that when visual attractiveness metrics receive only a very small weight relative to travel time, the resulting solutions become less consistent. In this situation, travel time dominates the weighted objective function and the visual attractiveness metrics mainly act as secondary criteria. As a result, the stochastic local search procedure may prioritise improvements in different metrics across runs, leading to solutions that do not follow a clear pattern.

Together, these results demonstrate that a residential kerbside waste collection route using a side loader can be optimised by balancing travel time efficiency with visual attractiveness. The optimisation framework incorporated visual attractiveness metrics with travel time in a weighted sum objective. In all case studies, this led to routes that are more regionally compact, service opposite kerbsides closer in time, and are less complex than travel-time-optimal routes. However, these improvements were accompanied by increases in travel time and directed Euclidean distance. Although trade-offs between metrics are inherent to multi-objective optimisation, the findings showed that the balance between travel time and visual attractiveness can be controlled to some extent. At the same time, improvements in one visual attractiveness metric may occur at the expense of another, which limits the ability to optimise all aspects simultaneously. Consequently, the weights may also require adjustment to the operational context and client preferences.

8.1.2. Scientific Contributions

The findings of this research align with the literature on visual attractiveness in vehicle routing, which identifies compactness, non-overlapping, and low complexity as key components of visually attractive solutions (Rossit et al., 2019). In existing studies, compactness and non-overlapping are predominantly analysed for multiple routes, where visual attractiveness concerns the spatial separation of route service territories (Kim et al., 2006; Matis, 2008; Poot et al., 2002; Sahoo et al., 2005; Tang & Miller-Hooks, 2006). In contrast, low complexity evaluates the internal structure of a route and is therefore applicable to an individual route (Poot et al., 2002; Rossit et al., 2019). A wide range of quantitative metrics has been developed to capture these components in optimisation models.

This study extends visual attractiveness concepts and metrics from vehicle routing problems to optimisation of a single side loader route. For a single route, the territorial separation between routes is irrelevant. Instead, the non-overlapping component appears as repeated revisits of areas or streets

within the same route. Consequently, non-overlapping is interpreted as an aspect of compactness rather than a separate component. Based on this reformulation of compactness, this study introduces metrics specifically designed for the optimisation of a single side loader route. These metrics adapt concepts from vehicle routing problems to capture regional and local compactness within a single route, including a revisit span metric that reflects behaviour specific to side loaders. To the best of my knowledge, such metrics have not been proposed for optimising a single side loader route. Moreover, existing waste collection studies rarely consider side loaders and those that do typically exclude visual attractiveness measures (Angelelli & Speranza, 2002). Thus, this study introduces novel metrics that extend the visual attractiveness literature by reformulating compactness for a single route and introducing metrics that capture regional and local compactness, including one specific to side loaders.

Low complexity and travel time remain directly applicable to a single route, as both evaluate properties of an individual route. The intra-route crossings metric provides a computationally efficient proxy for route complexity and has been used to penalise illogical route structures (Poot et al., 2002). Travel time remains the primary measure of operational efficiency, as it accounts for traffic conditions. Companies have shown that time-based routing increases margins (Broughton, 2024). Although neither travel time nor intra-route crossings are novel metrics, their integration with compactness metrics in a weighted optimisation objective is another contribution of this study. The results reveal clear trade-offs. Improvements in route compactness or simplicity may come at the expense of travel time and directed Euclidean distance. By jointly optimising travel time, compactness, and route complexity, this research demonstrates how visual attractiveness can be balanced with efficiency in individual side loader routes.

8.1.3. Practical Contributions

The practical relevance of this research is growing as municipalities increasingly adopt automated side loaders for residential waste collection (Blink Afvalbeheer, 2024; DAR NV, 2024; Gemeente Midden-Delfland, 2021). Compared to rear-loading trucks, side loaders offer labour and safety advantages by reducing manual handling and requiring only one operator (Archive Market Research, 2024). At the same time, the double traversal of streets associated with side loaders introduces routing challenges that are not well captured by routing methods that optimise for traditional efficiency metrics.

Visually attractive routes increase the acceptance of routing plans of waste collection operators (Poot et al., 2002; Rossit et al., 2019). Prior studies show that operators naturally prefer routes that are compact and largely free of crossings, even when these are not strictly optimal in distance or time (Van Rooij et al., 2003; Vickers et al., 2006). In this study, regional compactness, captured by the neighbourhood fragmentation metric, improves operational robustness to disruptions. If a route cannot be completed due to disruptions, such as vehicle breakdowns or delays, the unfinished work remains concentrated in a limited area rather than being scattered across neighbourhoods. In addition, routes with fewer self-crossings are easier for drivers to interpret and execute, reducing the likelihood of navigation errors.

Visual attractiveness is also relevant for residents. Regional and local compactness improves service predictability by ensuring that nearby streets and opposite kerbsides are serviced within a short time window. When collection vehicles move through an area in a logical and continuous pattern, residents can anticipate when the vehicle will pass. This reduces situations in which large collection trucks unexpectedly re-enter previously serviced streets or areas. This contributes to safer neighbourhoods, particularly in suburban residential areas where large side loaders operate on streets shared with pedestrians and children. In addition, servicing one kerbside of a street and returning several hours later to service the opposite side may create confusion or the perception that collections were missed. Therefore, more predictable vehicle movements contribute to a better experience for residents.

For AMCS, this research shows that visual attractiveness metrics can be quantified and optimised without requiring data on driving paths and geometry. The use of a weighted objective allows the optimisation framework to be tuned to different preferences, enabling AMCS to balance travel time, regional and local compactness, and route complexity based on requirements of their clients. This flexibility supports practical use, as municipalities or operators may differ in their tolerance for travel time increases in exchange for more visually attractive routes. Consequently, the optimisation framework provides AMCS with a configurable method to optimise a single side loader route.

8.2. Research Limitations

In this section, several limitations are acknowledged to clarify what can and cannot be concluded from this research. These limitations relate to the data and network representation, the definition of visual attractiveness metrics, and the design choices made for the optimisation framework. For each limitation, it is also discussed why the results remain valid for addressing the research questions.

8.2.1. Data and Network Representation

A primary limitation of this research is the absence of data on driving paths and road geometry. As a result, optimisation and evaluation are performed at the level of link clusters rather than on the true road network. Except for travel time and revisit span, the visual attractiveness metrics do not explicitly account for actual driving paths. Therefore, improvements in metric values do not always translate directly into visually attractive routes on the underlying road network. As shown in Section 7.3, some neighbourhood switches in the real road network are not counted in the abstract network. Similarly, reducing intra-route crossings based on straight-line connections does not always eliminate crossings in the actual road network. These examples highlight the limitations of the abstract network. Nevertheless, modelling at the link cluster level keeps the problem computationally manageable and preserves the sequencing decisions needed to study visual attractiveness in optimisation.

A second limitation concerns the grouping of orders into link clusters. As discussed in Section 7.3, orders that appear to lie on the same street segment are not always assigned to the same link cluster. In some cases, this may affect the neighbourhood fragmentation metric, especially when these link clusters are placed in different neighbourhoods. Moreover, information on which link clusters belong to the same street was not available. Although defining link clusters as street segments between intersections offers routing flexibility, it makes it harder to penalise servicing the same street in multiple passes. Incorporating data on which link clusters belong to which streets could improve the route.

Another data limitation is that only 12 case study datasets from two regions were available. While this limits the generalisability of the findings to other geographic contexts, the descriptive analysis in Chapter 3 shows reasonable variation between the datasets in terms of size and spatial structure. Moreover, the use of real operational datasets allows direct comparison with existing routes. This is particularly relevant for this study, as there are, to the best of the author's knowledge, no publicly available benchmark datasets for side loader routes. This enabled studying the behaviour of the proposed visual attractiveness metrics and optimisation framework in a realistic operational setting.

Finally, the optimisation of the exact single-objective models used to analyse individual metrics in Chapter 4 was only possible for small instances. Therefore, a smaller test set had to be extracted. Although this limits the generalisability of the conclusions, the test set was derived from real-world data and carefully selected to represent a geographically coherent neighbourhood, as described in Section 3.3. Therefore, it remains representative of side loader routing and is sufficient to analyse metric behaviour, which was the primary goal of these models.

8.2.2. Visual Attractiveness Metrics

Several assumptions were made to define the candidate visual attractiveness metrics in Chapter 4. The intra-route crossings metric relies on geometric approximations, as road geometry data was unavailable. These approximations do not align perfectly with crossings on the actual road network. Similarly, spread-based metrics were computed over fixed segments of the route. Alternative formulations, such as adaptive route segments, were not considered within the scope of this research.

In addition, neighbourhoods were derived using k-means clustering, since no spatial neighbourhoods were provided. This approach may misrepresent neighbourhood boundaries in irregular networks. As demonstrated in Section 7.3, the practical performance of the neighbourhood fragmentation metric is sensitive to the definition of neighbourhoods and their alignment with the road network. Inaccurate or misaligned boundaries can result in neighbourhood switches being incorrectly or not penalised, which can affect the visual attractiveness of routes on the road network.

Despite these simplifications, the metric definitions remain valid for the research objective. The goal of this study was not to identify optimal metric definitions, but to investigate whether such metrics can be incorporated into single-route optimisation and how they influence route structure. The chosen

definitions are computationally efficient, interpretable, and sufficient to demonstrate the effects on route behaviour.

8.2.3. Optimisation Framework

An important limitation of the optimisation framework is that it relies on a fixed set of parameters. The values of several parameters described in Chapter 5, including the operator probabilities, stopping conditions, preference weights, revisit tolerance, and neighbourhood definitions, were chosen based on educated estimates, as tuning would be too computationally extensive. Although the parameters that affect search behaviour were tuned within predefined ranges in Chapter 6, the results remain dependent on the selected ranges. As a result, alternative value ranges or tuning of other parameters may result in improved performance of the optimisation framework.

In addition, the optimisation framework requires an initial route from which the search process starts. In this study, (near-)optimal solutions for travel time were used for this purpose. However, the sensitivity of the framework to the initial solution was not explicitly investigated. Different starting solutions could influence convergence behaviour and route quality.

Moreover, the choice of optimisation method was partly arbitrary, although it was guided by the literature reviewed in Chapter 2. For a specialised problem such as optimising a single side loader route, alternative methods may achieve faster convergence or better results. Similarly, the selection of metrics involves subjective design choices. Although these choices were informed by the results in Chapter 4, alternative combinations of metrics could also have been explored.

Furthermore, the optimisation framework relies on a relatively small set of simple neighbourhood operators. While this design choice follows standard SA implementations and ensures broad exploration of the solution space, only a limited set of operators and configurations was considered in this study. In particular, the framework applies relatively small moves through relocate, swap, and Or-Opt operators and does not include a broader range of operators that explore larger neighbourhood moves. As a result, the search process may not fully exploit beneficial transformations of the route. A more extensive investigation of alternative neighbourhood operators and their combinations was beyond the scope of this study but could improve the effectiveness of the search procedure.

Another practical limitation concerns runtime. Especially for larger case studies, the optimisation framework requires considerable computation time. While this is acceptable for research and strategic planning purposes, this may limit its applicability in realistic operational environments.

Finally, the weighted sum objective and the adopted sensitivity analysis limit the extent to which trade-offs between individual objectives can be fully explored. Although varying the preference weights of visual attractiveness metrics with a scaling factor provides insight into the balance between travel time and visual attractiveness, it does not characterise the full Pareto-optimal solution set. Additionally, trade-offs between the visual attractiveness metrics themselves were not examined.

Despite these limitations, the chosen optimisation method and its parameter settings are appropriate for the research objective. The framework was designed to balance solution quality and computational effort by allowing exploration and exploitation of the solution space. Using a travel-time-optimal route as the initial solution ensures a feasible and efficient starting point, while the weighted-sum objective guides the search performed by the neighbourhood operators towards improved routes. The selected parameter and weight configurations produced clear improvements in visual attractiveness across the case studies. Although not exhaustive, the optimisation framework demonstrates that visual attractiveness objectives can be meaningfully incorporated into single-route optimisation under realistic operating conditions for side loaders. As such, it provides a solid foundation for further extensions.

8.3. Recommendations for Future Research and Practice

Based on the discussion of the results and limitations of this study, several recommendations for future research and practical implementation are formulated in the following subsections.

8.3.1. Recommendations for Research

Future research can extend this work in several directions, including the refinement of visual attractiveness metrics, further development of the optimisation framework, and exploration of alternative optimisation methods. The recommendations are presented in order of priority, reflecting their expected impact on optimisation performance and the effort required for further investigation.

The most important direction concerns the further development of visual attractiveness metrics for the optimisation of a single route. The metrics proposed in this thesis are not intended to represent an exhaustive set of possible measures. Other metrics may capture additional aspects of visual attractiveness. In particular, the use of road geometry and driving path data could enable more realistic metrics of route complexity, such as penalties based on turns or street curvatures. This would also allow intra-route crossings to be evaluated directly on the road network. Such metrics may better reflect how drivers and residents perceive route quality. Moreover, more sophisticated clustering methods that account for the road network could improve neighbourhood definitions. This would allow neighbourhoods to better reflect how drivers and residents perceive service areas and improve the practical usefulness of the neighbourhood fragmentation metric. In addition, alternative combinations of visual attractiveness metrics could also be explored. The results indicated that the local jumps metric may conflict with other compactness measures, suggesting that the current set of metrics is not necessarily optimal. Future research could therefore investigate whether alternative metric sets lead to more desirable route structures. In such analyses, it is important to explore the trade-offs between the visual attractiveness metrics themselves and between visual attractiveness and travel time in more detail. In this study, trade-offs among the visual attractiveness metrics were not examined explicitly, and the analysis of the trade-off between travel time and visual attractiveness could be extended to better approximate the Pareto-optimal solution set.

A second priority concerns the further development of the optimisation framework itself. The influence of the initial solution on the performance and convergence behaviour of the framework was not analysed and could therefore be examined in future work. Similarly, the performance of the framework should be evaluated on additional datasets, as this study considered only 12 case studies from two regions. Additionally, more extensive parameter tuning, including operator probability distributions, neighbourhood sizes, revisit tolerances, and other search behaviour parameters, may further improve performance. Moreover, the current implementation relies on only three simple local search operators, namely relocate, swap, and Or-Opt. Future research could investigate whether introducing larger moves or a broader set of operators improves solution quality or runtime. Examples include extended Or-Opt moves, k-opt exchanges, or more advanced operators such as Lin–Kernighan moves. In contrast, highly targeted operators that modify specific parts of the route structure may restrict exploration and bias the search toward particular route structures, potentially reducing the ability of the algorithm to escape local optima. In addition, operators that rely on symmetric route structures, such as 2-opt, are less suitable for asymmetric routing problems.

Improving the computational efficiency of the optimisation framework is also an important topic for future research. As illustrated in Subsection 6.2.2, the best-so-far objective value stabilises before the algorithm reaches the final temperature, suggesting that good solutions may be obtained earlier in the search process. Therefore, future research could investigate alternative stopping criteria, such as increasing the final temperature or introducing time-based stopping criteria. Similarly, the exploration phase could potentially be shortened by lowering the initial temperature. However, this should be carefully evaluated, as it can affect the algorithm's ability to escape local optima. In addition, runtime may be reduced by parallel computation. Since each iteration evaluates multiple candidate routes independently, the generation and evaluation of these candidates could be parallelised across multiple processing cores.

A final future research direction concerns the investigation of alternative optimisation methods. The choice of SA as the solution method was guided by the literature, but remains subjective to some extent. Therefore, alternative optimisation methods could be investigated to determine whether other metaheuristics or hybrid methods achieve improved solution quality or runtime. In addition, multi-objective optimisation methods that better approximate the Pareto frontier could provide more insight into trade-offs between the metrics. Furthermore, other multi-objective optimisation methods may be preferred when objectives are defined with strict priorities or have minimum and maximum service requirements.

8.3.2. Recommendations for AMCS

Several recommendations can be formulated for the practical implementation of visual attractiveness within AMCS routing systems. The recommendations are presented in order of priority, reflecting their expected impact on operational performance and implementation efforts.

The most important recommendation concerns the prioritisation of certain metrics within the optimisation framework. In particular, the neighbourhood fragmentation metric should receive the highest priority, as it enforces regional compactness and ensures that service activity remains concentrated within a limited area before moving on to another. This improves operational robustness to disruptions. When a route cannot be completed due to vehicle breakdowns or delays, the unfinished work remains confined to a small region rather than scattered across multiple neighbourhoods. Closely related, the revisit span metric should also receive a relatively high priority. This metric captures behaviour that is specific to side loader operations by encouraging the servicing of opposite kerbsides within a short time interval. In practice, both metrics improve service predictability and safety for residents, since nearby streets and both kerbsides of a street are serviced within a shorter time window. In contrast, the local jumps and intra-route crossings metrics appear less critical in practical implementations. The local jumps metric tends to conflict with other compactness measures and both metrics rely on straight-line approximations rather than the actual road network. Therefore, AMCS may consider assigning these metrics low preference weights or excluding them from the optimisation objective in certain applications.

At the same time, the preference weights should remain configurable and calibrated to the operational context. Routes with different spatial structures or service requirements may require different weight combinations. The weighted sum objective provides flexibility in this regard, allowing routing software providers such as AMCS to balance travel time, regional and local compactness, and route complexity according to the preferences of municipalities or operators. Since optimisation is performed on an abstract network, routes generated by the framework should also be visually inspected on the actual road network when calibrating to the operational context.

In addition, neighbourhood definitions should be refined beyond the simple k-means clustering method. In many cases, neighbourhoods can be defined based on existing administrative boundaries or postcodes that better reflect how drivers and residents perceive service areas. Such definitions improve the practical usefulness and performance of the neighbourhood fragmentation metric. When predefined neighbourhoods are not available, clustering methods such as k-means can still be applied to approximate neighbourhoods. Routing software providers such as AMCS may already implement clustering methods that could be adapted for this purpose.

Furthermore, attention should be given to how orders are grouped into street segments, referred to in this study as link clusters. The results indicate that link clusters do not always correspond to complete street segments, which can affect the neighbourhood definitions and logical sequencing of orders. Improving the grouping of orders into link clusters could improve the effectiveness of the neighbourhood fragmentation and revisit span metrics. A related extension would be connecting the link clusters to complete streets and introducing an additional metric that encourages the consecutive servicing of entire streets rather than individual kerbside segments. While the revisit span metric ensures that opposite kerbsides of the same street segments are serviced within a short time interval, it does not necessarily ensure that all of the street segments of a street are serviced together.

A research direction for the longer term concerns the modelling of household service locations. This study assumes fixed service points along kerbsides, requiring double traversal of streets by side loaders. In practice, municipalities may allow or encourage residents to place waste at designated collection areas or on a specific kerbside. Incorporating flexible service locations into the optimisation framework could create opportunities to reduce double traversal, thereby improving travel time efficiency and visual attractiveness of routes. This extension would require careful consideration of street characteristics and coordination between policy and routing decisions in the model, but may substantially improve efficiency and visual attractiveness in suitable contexts.

9

Conclusion

The purpose of this research was to develop an optimisation framework for residential kerbside waste collection that produces a single route, or household service sequence, that balances operational efficiency and visual attractiveness. The framework explicitly accounted for operating conditions specific to side loaders, including double traversal of streets, traffic directions, and U-turn restrictions. Double traversal is required because side loaders can only collect waste from one side of the street per pass. Consequently, routes optimised solely for efficiency are often visually unattractive in side loader operations, as they result in fragmented service of areas, long time intervals between servicing opposite kerbsides, and complex route shapes. This reduces operational clarity for drivers and predictability for residents. Against this background, this thesis addressed the following research question:

> *How can a residential kerbside waste collection route be designed to accommodate double traversal, while optimising visual attractiveness and travel time under relevant operating conditions?*

This research demonstrated that a residential kerbside waste collection route accommodating double traversal can be designed by modelling both operational efficiency and visual attractiveness within an optimisation framework. The problem was formulated as a node-based optimisation model designed to construct a single route. Double traversal and realistic operating conditions were incorporated through a directed network representation in which travel times were computed as the fastest paths on the underlying road network, respecting traffic directions and U-turn restrictions. Visual attractiveness was quantified using a set of complementary metrics that capture regional compactness, local compactness, and route complexity. These metrics were integrated with travel time in a weighted sum objective and optimised using simulated annealing. Empirical evaluation across the analysed case studies demonstrated that the selected parameters and preference weights provide a balanced configuration. This resulted in routes that are more visually attractive compared to the travel-time-optimal baseline routes generated by AMCS. The sensitivity analysis showed that the trade-off between travel time and visual attractiveness can be managed by uniformly scaling the weights of the visual attractiveness metrics.

The following findings support and clarify this conclusion. The problem was formulated as a single-route optimisation model in which each street segment is treated as a node to be serviced. This allows the route to be represented as an ordered sequence of street segments and ensures compatibility with the routing systems for which it is designed. By modelling travel time using the fastest paths on the underlying road network, the framework accounted for the double traversal of streets required by side loaders. This abstract network also captured realistic operating conditions, such as U-turn prohibitions on regular streets, U-turn time penalties on dead-heading streets, and traffic directions.

Furthermore, this research showed that visual attractiveness can be integrated into an optimisation framework for a waste collection route using a side loader. To quantify visual attractiveness, candidate performance metrics were developed and evaluated. These metrics were developed for a single side loader route by adapting metrics from the literature on visual attractiveness for multiple routes. Through a correlation analysis and single-objective optimisation experiments, it was demonstrated that visual at-

tractiveness is a multi-dimensional concept, capturing regional compactness, local compactness, and route complexity. Based on these analyses, a subset of metrics was selected to represent complementary aspects of route structure. Regional compactness was captured by neighbourhood fragmentation, local compactness by local jumps, side-loader-specific local compactness by revisit span, and route complexity by intra-route crossings. These metrics were formulated on the abstract network, where connections between street segments are approximated by straight-line links between segment endpoints rather than geometric paths on the road network.

The multi-objective optimisation framework integrates these metrics with travel time in a weighted sum objective to balance efficiency and visual attractiveness. This multi-objective optimisation method was chosen because no strict priorities or service requirements were imposed on the objectives. Given the scale and complexity of realistic municipal networks, simulated annealing was identified as an appropriate optimisation method. The framework proved to be capable of handling larger networks while balancing exploration and exploitation of the solution space to avoid getting stuck in local optima. The weighted sum objective is well suited to simulated annealing, as it provides a single scalar objective that guides the search process.

Empirical evaluation on 12 real municipal datasets demonstrated that the optimisation framework significantly reshapes route structure. These case studies were provided by the waste collection software provider AMCS, whose routing system generates baseline routes that are optimised for travel time. In 11 of the 12 case studies, routes were identified that improved visual attractiveness compared to these baseline routes. Neighbourhood fragmentation, revisit spans between opposite kerbsides, and intra-route crossings were consistently reduced, with median decreases of approximately 90% for neighbourhood fragmentation and revisit span and nearly 50% for intra-route crossings. These results indicate that optimising for visual attractiveness metrics formulated on the abstract network can meaningfully influence route design on the underlying road network. They further suggest that detailed geometric modelling of driving paths is not required to produce visually attractive routes in practice. Moreover, the results confirm that visually attractive routes do not emerge from travel time optimisation but require explicit modelling.

At the same time, the results highlight trade-offs inherent to multi-objective optimisation. Improvements in compactness and route simplicity were achieved at the cost of increased travel time, with a median increase of approximately 40%, reflecting the expected trade-off between efficiency and visual attractiveness. In addition, improving these metrics often led to a deterioration in local compactness, as measured by local jumps in terms of directed Euclidean distance. This metric increased by a median of approximately 19%, reflecting its closer alignment with travel time than with the other metrics. Overall, these findings show that improving one aspect of visual attractiveness may come at the expense of another, limiting the ability to optimise all metrics simultaneously.

The sensitivity analysis further confirms this behaviour. Increasing the relative importance of the visual attractiveness metrics leads to reductions in neighbourhood fragmentation, revisit spans between opposite kerbsides, and intra-route crossings. However, these improvements become progressively smaller, while travel time and directed Euclidean distance increase more substantially. Additional experiments also show that when the visual attractiveness metrics receive only a very small weight relative to travel time, the resulting solutions become less consistent. Travel time dominates the objective and the stochastic search procedure may prioritise improvements in different metrics. Overall, these findings indicate that the selected preference weights provide a reasonable configuration for the case study analysed. More broadly, the results show that the trade-off between efficiency and visual attractiveness can be controlled through calibration of the preference weights. Nevertheless, careful calibration remains necessary to achieve an appropriate balance for networks with different structural characteristics.

In addition to the empirical findings, it is important to reflect on the modelling assumptions and choices in this study. This research was guided by the expectation that visual attractiveness concepts developed for optimisation problems with multiple routes could be adapted to a single side loader route. The results largely confirm this assumption, while also showing that visual attractiveness is a subjective and multi-dimensional concept that cannot be captured by a single metric. Trade-offs inevitably arise when optimising for different visual attractiveness metrics. Moreover, the modelling choices in this thesis prioritised interpretability and applicability to the data available from AMCS. Although this required several simplifications, the optimisation framework proved capable of generating more visually attrac-

tive routes on the road network across multiple case studies. It also provided insights into how visual attractiveness objectives interact with travel time in realistic side loader operations.

Based on these findings and reflection, the contributions of this thesis can be described from both a scientific and a practical perspective. From a scientific perspective, this study extends existing visual attractiveness concepts for multiple routes to a single route setting for a side loader. This context has received limited attention in existing studies. Whereas compactness and non-overlapping are typically analysed for multiple routes using service areas, this research reformulates these concepts for a single route. In particular, compactness is evaluated for segments of a single route rather than across route territories. Moreover, territorial separation between routes is no longer relevant. Instead, non-overlapping shows as repeated revisits of the same streets or areas and is therefore treated as an aspect of compactness. Furthermore, the introduction of a revisit span metric captures behaviour specific to side loader operations. By reformulating visual attractiveness concepts and introducing compactness metrics that capture both regional and local compactness for the optimisation of a single side loader route, this research contributes novel metrics to the visual attractiveness literature.

This study also shows how these metrics can be combined with travel time and intra-route crossings in a weighted sum objective. Although neither travel time nor intra-route crossings are novel metrics individually, their integration with compactness objectives for a single side loader route provides new insights into how visually attractive routes can be generated and balanced with efficiency. The results further clarify the trade-off between travel time and visual attractiveness, contributing to a better understanding of how these objectives can be balanced in the optimisation of a single side loader route.

From a practical perspective, the findings are particularly relevant as municipalities increasingly adopt automated side-loading vehicles. The double traversal required by side loaders introduces challenges that are not addressed by traditional routing methods focused on efficiency. The optimisation framework demonstrates that more visually attractive routes can be generated compared to current operational routes. For collection crews, compact routes with few self-crossings are easier to interpret and execute, improving operational clarity and reducing the likelihood of navigation errors. In addition, regional compactness improves operational robustness to disruptions, as unfinished work remains concentrated within a limited area. For residents, regional and local compactness improves service predictability by ensuring that nearby streets and opposite kerbsides are serviced within a short time window. This reduces situations in which collection vehicles unexpectedly return to previously serviced areas and contributes to safer neighbourhoods, particularly in suburban areas where large side loaders operate on streets shared with pedestrians and children. Finally, the optimisation framework provides technology providers such as AMCS with a configurable method to incorporate visual attractiveness objectives into routing methods. The weighted sum objective allows the relative importance of travel time and visual attractiveness to be adjusted according to client preferences, supporting a flexible routing solution in real-world operations.

Overall, this research demonstrates that visually attractive routing is not only a subjective concept, but can be formalised, quantified, and optimised within a structured framework. For residential waste collection with side loaders, this means designing routes that concentrate service within one area before moving to another, minimise long time gaps between servicing nearby streets or opposite kerbsides, and follow a clear spatial pattern without unnecessary zigzagging or crossings. By combining regional compactness, local compactness, and route simplicity with travel time in a weighted sum objective, the optimisation framework generates routes that are more predictable, easier to execute, and more robust to operational disruptions. At the same time, double traversal and realistic operating conditions are explicitly incorporated through a directed network representation in which travel times are computed as the fastest paths on the underlying road network, respecting traffic directions and U-turn restrictions.

However, improving visual attractiveness comes at the cost of increased travel time. In addition, visual attractiveness metrics may conflict, which means that not all structural aspects can be improved simultaneously. These trade-offs are inherent to multi-objective optimisation. Nevertheless, the results show that these trade-offs can be managed through the use of a weighted sum in the optimisation framework. As a result, residential kerbside waste collection routes accommodating double traversal can be designed to balance efficiency with visual attractiveness under realistic operating conditions. As municipalities increasingly adopt side loaders, incorporating visual attractiveness into routing decisions offers a promising direction for improving operational clarity for crews and service predictability

for residents.

Based on these findings, there are several directions for further research and practical implementation. Future research could refine the proposed visual attractiveness metrics by incorporating detailed road geometry and driving path data, exploring alternative metric combinations, and analysing the trade-offs between all metrics in more detail. In addition, the optimisation framework could be further developed through improved neighbourhood operators, broader parameter tuning on additional datasets, and enhanced computational efficiency. Moreover, alternative optimisation methods could be explored for improved solution quality or runtime. From a practical perspective, routing systems such as those developed by AMCS may benefit from prioritising compactness-related metrics, particularly neighbourhood fragmentation and revisit span, as these improve operational robustness and promote more predictable and safer collection operations. Additional recommendations include calibrating preference weights to client requirements and operational context, validating routes on the actual road network, using administrative neighbourhood definitions, and improving the grouping of orders into link clusters and their connection to complete streets. Finally, modelling flexible service locations, such as shared collection points or designated kerbside, could reduce double traversal and therefore improve both efficiency and visual attractiveness. Together, these directions provide opportunities to extend and refine the visual attractiveness metrics and optimisation framework.

References

AMCS Group. (2025a). *About amcs*. Retrieved April 10, 2025, from https://www.amcsgroup.com/

AMCS Group. (2025b). *Amcs routeplanner brochure* [Cover image adapted in this thesis]. Retrieved March 10, 2025, from https://www.amcsgroup.com/nl/kennisbank/brochures/amcs-routeplanner/

Angelelli, E., & Speranza, M. G. (2002). The application of a vehicle routing model to a waste collection problem: Two case studies. *Journal of the Operational Research Society*, *53*(9), 944–952. https://doi.org/10.1057/palgrave.jors.2601402

Aransay, J., Galilea, V., & Hernández, I. (2020). A tool to improve the efficiency of waste collection: Development and application to a case study. *Proceedings of 8th Transport Research Arena (TRA 2020)*.

Archive Market Research. (2024). *Automated side loader garbage truck market report*. Retrieved August 25, 2025, from https://www.archivemarketresearch.com/reports/automated-side-loader-garbage-truck-600836

Bautista, J., Fernández, E., & Pereira, J. (2008). Solving an urban waste collection problem using ants heuristics. *Computers & Operations Research*, *35*(9), 3020–3033. https://doi.org/10.1016/j.cor.2007.01.029

Beliën, J., De Boeck, L., & Van Ackere, J. (2012). Municipal solid waste collection and management problems: A literature review. *Transportation Science*, *48*(1). https://doi.org/10.1287/trsc.1120.0448

Benitez-Bravo, R., Gomez-González, R., Rivas-García, P., Botello-Álvarez, J. E., Huerta-Guevara, O. F., García-León, A. M., & Rueda-Avellaneda, J. F. (2021). Optimization of municipal solid waste collection routes in a latin-american context. *Journal of the Air & Waste Management Association*, *71*(11), 1415–1427. https://doi.org/10.1080/10962247.2021.1957040

Berg, M. D., Cheong, O., Kreveld, M., & Overmars, M. (2008). *Computational geometry: Algorithms and applications* (3rd ed.). Springer. https://doi.org/10.1007/978-3-540-77974-2

Blink Afvalbeheer. (2024). *Afvalinzameling met zijlader*. Blink Afvalbeheer. Retrieved January 23, 2026, from https://mijnblink.nl/afvalinzameling-zijlader

Bonomo, F., Durán, G., Larumbe, F., & Marenco, J. (2012). A method for optimizing waste collection using mathematical programming: A buenos aires case study. *Waste Management & Research*, *30*(3), 311–24. https://doi.org/10.1177/0734242X11402870

Boryczka, U., & Szwarc, K. (2019). The adaptation of the harmony search algorithm to the atsp with the evaluation of the influence of the pitch adjustment place on the quality of results. *Journal of Information and Telecommunication*, *3*(1), 2–18. https://doi.org/10.1080/24751839.2018.1503149

Boryczka, U., & Szwarc, K. (2020). An effective hybrid harmony search for the asymmetric travelling salesman problem. *Engineering Optimization*, *52*(2), 218–234. https://doi.org/10.1080/0305215X.2019.1579804

Brest, J., & Žerovnik, J. (2005). A heuristic for the asymmetric traveling salesman problem. *Proceedings of the 6th Metaheuristics International Conference (MIC 2005)*, 145–150. https://www.researchgate.net/publication/228640882_A_heuristic_for_the_asymmetric_traveling_salesman_problem

Broughton, K. (2024). Waste management boosts profit margins by rerouting and replacing trash trucks. *The Wall Street Journal*.

Buriol, L. S., França, P. M., & Moscato, P. (2004). A new memetic algorithm for the asymmetric traveling salesman problem. *Journal of Heuristics*, *10*(5), 483–506. https://doi.org/10.1023/B:HEUR.0000045320.33792.2f

Burke, E. K., Cowling, P. I., & Keuthen, R. (2001). Effective local and guided variable neighbourhood search methods for the asymmetric travelling salesman problem. In E. Boers (Ed.), *Applications*

of evolutionary computing (pp. 203–212, Vol. 2037). Springer. https://doi.org/10.1007/3-540-45365-2_21

Choi, I.-C., Kim, S.-I., & Kim, H.-S. (2003). A genetic algorithm with a mixed region search for the asymmetric traveling salesman problem. *Computers & Operations Research*, *30*(5), 773–786. https://doi.org/10.1016/S0305-0548(02)00050-3

Cirasella, J., Johnson, D. S., McGeoch, L. A., & Zhang, W. (2001). The asymmetric traveling salesman problem: Algorithms, instance generators, and tests. *Algorithm Engineering and Experimentation*, *2153*, 32–59. https://doi.org/10.1007/3-540-44808-X_3

Clark, P. J., & Evans, F. C. (1954). Distance to nearest neighbor as a measure of spatial relationships in populations. *Ecology*, *35*(4), 445–453. https://doi.org/10.2307/1931034

Cortinhal, M. J., Mourão, M. C., & Nunes, A. C. (2016). Local search heuristics for sectoring routing in a household waste collection context. *European Journal of Operational Research*, *255*(1), 68–79. https://doi.org/10.1016/j.ejor.2016.04.013

Croes, G. A. (1958). A method for solving traveling-salesman problems. *Operations Research*, *6*(6), 791–812. https://doi.org/10.1287/opre.6.6.791

Dantzig, G., Fulkerson, R., & Johnson, S. (1954). Solution of a large-scale traveling-salesman problem. *Journal of the Operations Research Society of America*, *2*(4), 393–410. https://doi.org/10.1287/opre.2.4.393

Dantzig, G., & Ramser, J. (1959). The truck dispatching problem. *Management Science*, *6*(1), 80–91. https://doi.org/10.1287/mnsc.6.1.80

DAR NV. (2024). *Afvalinzameling met de zijlader*. DAR NV. Retrieved January 23, 2026, from https://www.dar.nl/zijlader/

Das, S., & Bhattacharyya, B. K. (2015). Optimization of municipal solid waste collection and transportation routes. *Waste Management*, *43*, 9–18. https://doi.org/10.1016/j.wasman.2015.06.033

Dorigo, M., Maniezzo, V., & Colorni, A. (1996). Ant system: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, *26*(1), 29–41. https://doi.org/10.1109/3477.484436

Freisleben, B., & Merz, P. (1996). A genetic local search algorithm for solving symmetric and asymmetric traveling salesman problems. *Proceedings of the IEEE International Conference on Evolutionary Computation*, 616–621. https://doi.org/10.1109/ICEC.1996.542672

Frieze, A. M., Galbiati, G., & Maffioli, F. (1982). On the worst-case performance of some algorithms for the asymmetric traveling salesman problem. *Networks*, *12*(1), 23–39. https://doi.org/10.1002/net.3230120103

Gambardella, L. M., & Dorigo, M. (1996). Solving symmetric and asymmetric tsps by ant colonies. *Proceedings of IEEE International Conference on Evolutionary Computation*, 622–627. https://doi.org/10.1109/ICEC.1996.542672

Gemeente Midden-Delfland. (2021). *Afvalinzameling*. Retrieved September 8, 2025, from https://www.middendelfland.nl/afvalinzameling

Glover, F. (1986). Future paths for integer programming and links to artificial intelligence. *Computers & Operations Research*, *13*(5), 533–549. https://doi.org/10.1016/0305-0548(86)90048-1

Golden, B. L., Assad, A. A., & Wasil, E. A. (2002). Routing vehicles in the real world: Applications in the solid waste, beverage, food, dairy, and newspaper industries. In P. Toth & D. Vigo (Eds.), *The vehicle routing problem* (pp. 245–286). SIAM. https://doi.org/10.1137/1.9780898718515.ch10

Goodrich, M. T., & Tamassia, R. (2014). *Algorithm design and applications*. John Wiley & Sons.

Han, H., & Ponce-Cueto, E. (2015). Waste collection vehicle routing problem: A literature review. *Promet - Traffic & Transportation*, *27*(4), 345–358. https://doi.org/10.7307/ptt.v27i4.1616

Hess, C., Dragomir, A. G., Doerner, K. F., & Vigo, D. (2024). Waste collection routing: A survey on problems and methods. *Central European Journal of Operations Research*, *32*, 399–434. https://doi.org/10.1007/s10100-023-00892-y

Holland, J. H. (1975). *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence*. University of Michigan Press.

Johnson, D. S., & McGeoch, L. A. (1997). The traveling salesman problem: A case study in local optimization. In E. H. L. Aarts & J. K. Lenstra (Eds.), *Local search in combinatorial optimization* (pp. 215–310). John Wiley & Sons.

Kanellakis, P.-C., & Papadimitriou, C. H. (1980). Local search for the asymmetric traveling salesman problem. *Operations Research*, *28*(5), 1086–1099. https://doi.org/10.1287/opre.28.6.1086

Karadimas, N. V., Papatzelou, K., & Loumos, V. G. (2007a). Genetic algorithms for municipal solid waste collection and routing optimization. In C. Boukis, A. Pnevmatikakis, & L. Polymenakos (Eds.), *Artificial intelligence and innovations 2007: From theory to applications* (pp. 223–231, Vol. 247). Springer. https://doi.org/10.1007/978-0-387-74161-1_24

Karadimas, N. V., Papatzelou, K., & Loumos, V. G. (2007b). Optimal solid waste collection routes identified by the ant colony system algorithm. *Waste Management & Research*, *25*(2), 139–147. https://doi.org/10.1177/0734242X07071312

Karp, R. M. (1979). A patching algorithm for the nonsymmetric traveling-salesman problem. *SIAM Journal on Computing*, *8*(4), 45–72. https://doi.org/10.1137/0208045

Kendall, M. G. (1938). A new measure of rank correlation. *Biometrika*, *30*(1/2), 81–93.

Kim, B.-I., Kim, S., & Sahoo, S. (2006). Waste collection vehicle routing problem with time windows. *Computers & Operations Research*, *33*(12), 3624–3642. https://doi.org/10.1016/j.cor.2005.02.045

Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, *220*(4598), 671–680. https://doi.org/10.1126/science.220.4598.671

Laporte, G. (1992). The traveling salesman problem: An overview of exact and approximate algorithms. *European Journal of Operational Research*, *59*(2), 231–247. https://doi.org/10.1016/0377-2217(92)90138-Y

Matis, P. (2008). Decision support system for solving the street routing problem. *Transport*, *23*(3), 230–235. https://doi.org/10.3846/1648-4142.2008.23.230-235

Mekamcha, K., Souier, M., Bessenouci, H. N., & Bennekrouf, M. (2021). Two metaheuristics approaches for solving the traveling salesman problem: An algerian waste collection case. *Operational Research*, *21*, 1641–1661. https://doi.org/10.1007/s12351-019-00529-6

Miller, C., Tucker, A., & Zemlin, R. (1960). Integer programming formulation of traveling salesman problems. *Journal of the ACM (JACM)*, *7*(4), 326–329.

Mladenović, N., & Hansen, P. (1997). Variable neighborhood search. *Computers & Operations Research*, *24*(11), 1097–1100. https://doi.org/10.1016/S0305-0548(97)00031-2

Nagata, Y., & Soler, D. (2012). A new genetic algorithm for the asymmetric traveling salesman problem. *Expert Systems with Applications*, *39*(10), 8947–8953. https://doi.org/10.1016/j.eswa.2012.02.029

Odili, J. B., Noraziah, A., & Zarina, M. (2021). A comparative performance analysis of computational intelligence techniques to solve the asymmetric travelling salesman problem. *Computational Intelligence and Neuroscience*, *2021*, 1–13. https://doi.org/10.1155/2021/6625438

Öncan, T., Altınel, I. K., & Laporte, G. (2009). A comparative analysis of several asymmetric traveling salesman problem formulations. *Computers & Operations Research*, *36*(3), 637–654. https://doi.org/10.1016/j.cor.2007.11.008

Pearson, K. (1896). Vii. mathematical contributions to the theory of evolution.-iii. regression, heredity, and panmixia. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, *187*, 253–318. https://doi.org/10.1098/rsta.1896.0007

Poot, A., Kant, G., & Wagelmans, A. P. (2002). A savings based method for real-life vehicle routing problems. *Journal of the Operational Research Society*, *53*(1), 57–68. https://doi.org/10.1057/palgrave.jors.2601265

Rambandara, R. D. S. S., Prabodanie, R., Karunarathne, E., & Rajapaksha, R. (2022). Improving the efficiency of urban waste collection using optimization: A case study. *Process Integration and Optimization for Sustainability*, *6*(3), 809–818. https://doi.org/10.1007/s41660-022-00232-8

Roberti, R., & Toth, P. (2012). Models and algorithms for the asymmetric traveling salesman problem: An experimental comparison. *EURO Journal on Transportation and Logistic*, *1*(1), 113–133. https://doi.org/10.1007/s13676-012-0010-0

Rosenkrantz, D. J., Stearns, R. E., & II, P. M. L. (1977). An analysis of several heuristics for the traveling salesman problem. *SIAM Journal on Computing*, *6*(3), 563–581. https://doi.org/10.1137/0206041

Rossit, D., Vigo, D., Tohmé, F., & Frutos, M. (2019). Visual attractiveness in routing problems: A review. *Computers & Operations Research*, *103*, 13–34. https://doi.org/10.1016/j.cor.2018.10.012

Sahoo, S., Kim, S., Kim, B.-I., Kraas, B., & Popov Jr, A. (2005). Routing optimization for waste management. *Interfaces*, *35*(1), 24–36. https://doi.org/10.1287/inte.1040.0109

Scott, D. W. (1992). *Multivariate density estimation: Theory, practice, and visualization*. John Wiley & Sons. https://doi.org/10.1002/9780470316849

Sniezek, J., & Bodin, L. (2006). Using mixed integer programming for solving the capacitated arc routing problem with vehicle/site dependencies with an application to the routing of residential sanitation collection vehicles. *Annals of Operations Research*, *144*, 33–58. https://doi.org/10.1007/s10479-006-0006-y

Spearman, C. (1904). The proof and measurement of association between two things. *The American Journal of Psychology*, *15*(1), 72–101.

Szwarc, K., & Boryczka, U. (2017). A comparative study of different variants of a memetic algorithm for atsp. *Computational Collective Intelligence: 9th International Conference, ICCCI 2017, Nicosia, Cyprus, September 27–29, 2017, Proceedings, Part II*, 76–86. https://doi.org/10.1007/978-3-319-67077-5_8

Talbi, E.-G. (2009). *Metaheuristics: From design to implementation*. John Wiley & Sons.

Tang, H., & Miller-Hooks, E. (2006). Interactive heuristic for practical vehicle routing problem with solution shape constraints. *Transportation Research Record: Journal of the Transportation Research Board*, *1964*(1), 9–18. https://doi.org/10.1177/0361198106196400102

United Nations Environment Programme. (2024). *Global waste management outlook 2024: Beyond an age of waste – turning rubbish into a resource* (tech. rep. No. DTI/2619/NA). United Nations Environment Programme (UNEP). https://wedocs.unep.org/20.500.11822/44939

Van Rooij, I., Stege, U., & Schactman, A. (2003). Convex hull and tour crossings in the euclidean traveling salesperson problem: Implications for human performance studies. *Memory & Cognition*, *31*(2), 215–220. https://doi.org/10.3758/BF03194380

Vickers, D., Lee, M. D., Dry, M., Hughes, P., & McMahon, J. A. (2006). The aesthetic appeal of minimal structures: Judging the attractiveness of solutions to traveling salesperson problems. *Perception and Psychophysics*, *68*(1), 32–42. https://doi.org/10.3758/BF03193653

Viotti, P., Polettini, A., Pomi, R., & Innocenti, C. (2003). Genetic algorithms as a promising tool for optimisation of the msw collection routes. *Waste Management & Research*, *21*(4), 292–298. https://doi.org/10.1177/0734242X0302100402

Zhang, W. (1993). Truncated branch-and-bound: A case study on the asymmetric tsp. *Proceedings of the AAAI 1993 Spring Symposium on AI and NP-Hard Problems*, 160–166.

A

Scientific Article

The scientific article of this study is presented on the next page.

Optimisation of Residential Waste Collection: Balancing Travel Time and Visual Attractiveness in Side Loader Routes

Amber van der Helm[1]

[1]Delft University of Technology, The Netherlands
Corresponding author: amber@vdhelm.com

Abstract. Residential waste collection with side loaders requires double traversal because vehicles service only one kerbside per pass. Consequently, routes optimised for efficiency often result in fragmented service of areas, long time intervals between visiting opposite kerbsides, and complex route shapes that are unintuitive to execute. Therefore, study developed a single-route optimisation framework that balances travel time efficiency with visual attractiveness under side loader operating constraints, including traffic directions, double traversal, and U-turn restrictions. The problem was formulated as an asymmetric travelling salesman problem (ATSP) on a directed network, where nodes represent directed kerbside street segments and travel times are computed as the fastest paths on the underlying road network. Visual attractiveness was quantified using complementary metrics for regional compactness, local compactness, and route complexity. These metrics were defined on an abstract network using straight-line connections to maintain computational tractability. They were combined with travel time in a weighted sum objective and optimised using simulated annealing. This enables a configurable trade-off between visual attractiveness and efficiency. The framework was evaluated on 12 municipal case studies provided by AMCS. Compared to the travel-time-optimal routes generated by AMCS, the framework consistently reduced neighbourhood fragmentation, revisit spans between opposite kerbsides, and intra-route crossings, while moderately increasing total travel time and directed Euclidean distance. A sensitivity analysis showed that uniformly scaling the visual attractiveness weights controls the trade-off between visual attractiveness and efficiency. However, very small scaling factors produced less consistent improvements in visual attractiveness, whereas larger factors yielded diminishing improvements while travel time increased substantially. Overall, the results suggest that visual attractiveness can be formalised on an abstract network and integrated into a configurable optimisation framework to produce visually attractive side loader routes on the road network.

Key words: asymmetric travelling salesman problem; multi-objective optimisation; residential waste collection; side loader; simulated annealing; visual attractiveness

1. Introduction

Global municipal solid waste generation is increasing rapidly as a result of urbanisation, economic growth, and changing consumption patterns. In 2023, global municipal solid waste production reached approximately 2.1 billion tonnes and is projected to approach 3.8 billion tonnes by 2050 [1]. Waste collection is the most resource-intensive activity of municipal solid waste management, accounting for up to 80% of total costs [2]. Therefore, improving waste collection operations is a priority for municipalities and waste collection operators seeking to reduce operational costs, traffic congestion, and environmental impact.

Within municipal solid waste collection systems, residential kerbside collection typically serves the largest number of stops and requires frequent service [3]. In many suburban areas, municipalities are increasingly transitioning from traditional rear-loading trucks to automated side-loading vehicles. Side loaders offer significant safety and labour advantages, as they can be operated by a single driver and eliminate most manual handling [4]. However, they introduce a fundamental routing complication. Because side loaders collect waste from only one side of the street, both kerbsides must be serviced in separate passes [5]. This so-called double traversal effectively doubles parts of the service network and forces vehicles to re-enter already serviced streets or neighbourhoods, often later in the route. As a result, a single route can become spatially fragmented and operationally difficult to follow. These challenges are amplified by traffic directions and street geometry, as heavy collection vehicles often cannot make U-turns or zigzag between opposite kerbsides.

In practice, route quality is not determined solely by operational efficiency. Prior studies show that operators prefer routes that are compact and free of crossings, even when they are not optimal in distance or time [6], [7]. This notion is commonly referred to as visual attractiveness in the vehicle routing literature [8], [9], [10]. Visually attractive routes are easier for drivers to execute, enhance operational robustness to disruptions, and improve public acceptance and safety by providing predictable and logical service patterns.

In the literature, visual attractiveness has predominantly been studied in vehicle routing problems, where a set of routes services an area together. In this context, three main components are commonly distinguished, namely compactness, non-overlapping, and low complexity [11]. Compactness refers to the spatial concentration of stops assigned to the same route. Non-overlapping focuses on the spatial separation between different routes,

ensuring that service territories do not interfere. Both concepts are defined by comparing multiple routes and their territories. In contrast, low complexity evaluates the geometric simplicity of an individual route, favouring smooth shapes without unnecessary crossings or zigzagging. Thus, it is directly applicable to a single route.

However, for a single route collecting residential waste with side loaders, unattractive behaviour shows differently. A single route lacks compactness when service is spatially fragmented, while non-overlapping is no longer defined between routes and is implicitly captured by compactness. This occurs when a route moves to new areas before completing the current one, when nearby streets are not serviced consecutively, and when large time intervals separate visits to opposite kerbsides of the same street. Additionally, low complexity remains relevant, as intuitive route shapes can be disrupted by unnecessary crossings or detours. These effects are caused by double traversal, traffic directions, and U-turn restrictions, which force the vehicle to revisit already serviced areas and streets, often much later. This makes even a single route vulnerable to unattractive behaviour. However, existing measures are mostly defined for multiple routes where visual attractiveness is assessed by comparing the shapes and territories of different routes. Consequently, current metrics do not capture the spatial fragmentation that arises when optimising a single waste collection route for a side loader, revealing a gap in the literature.

In addition, efficiency remains critical to ensure that routes are feasible in practice. While distance is often used as a proxy for operational cost and emissions [12], [13], this study quantifies efficiency using total travel time on a directed road network. Operators increasingly focus on travel time, since the shortest route is not necessarily the fastest and travel time reflects the ability to complete a route within a work shift [14]. However, side loader routes that minimise efficiency do not necessarily display visually attractive behaviour [10], [11]. Fast routes may cross neighbourhoods or postpone revisiting a street to service the other kerbside, while more attractive routes may increase total travel time. This creates a trade-off between efficiency and visual attractiveness, particularly in the presence of operational constraints such as double traversal, traffic directions and U-turn restrictions.

Therefore, the goal of this study is to develop a single-route optimisation framework for residential kerbside waste collection that optimises visual attractiveness alongside travel time efficiency under realistic operating conditions relevant to side loaders. To achieve this, the concepts of visual attractiveness are redefined for this context and a set of complementary visual attractiveness metrics is formulated for these concepts. These metrics are incorporated into an asymmetric travelling salesman problem (ATSP) formulation and combined with travel time in a weighted objective function. Double traversal, traffic directions, and U-turn restrictions are incorporated through a directed network representation in which travel times are computed as the fastest paths on the underlying road network. The optimisation framework is based on simulated annealing and evaluated on real municipal case studies in collaboration with waste management software provider AMCS [15].

The remainder of this article is structured as follows. section 2 reviews related work on waste collection routing, visual attractiveness in vehicle routing, and optimisation methods for the ATSP. section 3 describes the case studies and network representation. section 4 introduces the proposed visual attractiveness metrics for a single route and presents the mathematical formulation. section 5 describes the optimisation framework. section 6 reports results on real case studies and presents a sensitivity analysis of the weights in the objective function. Finally, section 7 concludes with scientific and practical implications and directions for future research.

2. Literature Review

Municipal solid waste collection is commonly studied as a routing problem in which vehicles serve collection points or streets and transport waste to disposal facilities. Unlike classical distribution routing, waste collection involves reverse logistics rather than forward deliveries from a central depot to customers [5], [16]. This section reviews the modelling approaches, visual attractiveness metrics, and optimisation methods relevant for this study.

2.1. Modelling Approaches for Waste Collection

A key modelling decision in waste collection routing concerns whether demand is represented at nodes or along arcs. This distinction reflects structural differences between industrial, commercial, and residential collection systems [3]. Industrial and commercial collection typically involve fixed container locations with relatively low stop density. In contrast, residential kerbside collection requires vehicles to traverse many street segments and serve a large number of households, resulting in higher operational complexity [10]. In node-based formulations, vehicles must visit a set of locations, such as containers or aggregated service points. In arc-based formulations, vehicles are required to traverse street segments that carry demand. Both modelling perspectives appear in the literature. Node-based models are commonly used in practice due to their compatibility with standard routing data structures and the wide availability of solution methods [5], [16]. In contrast, arc-based formulations naturally reflect kerbside service, particularly when both sides of a street must be serviced.

In addition to this distinction, modelling choices are shaped by operational constraints arising from vehicle limitations and the road network. These constraints determine which movements are allowed and, for side loaders, whether streets must be traversed twice to complete service. One-way streets, turn restrictions, and prohibited U-turns are typically incorporated by removing or penalising directed arcs in the graph [12], [17]. Mixed-graph formulations may further distinguish between different road types or allow parallel arcs to represent service on both kerbsides [18]. However, explicit modelling of side loader operations remains limited in the literature [19].

In this study, the focus lies on residential kerbside waste collection for a single route using a side loader. Because each kerbside must be serviced separately, demand is inherently directional. Although an asymmetric arc routing formulation would be a natural representation, the problem is reformulated as an ATSP. Each kerbside of a street segment that requires collection is represented as a node, while the deadhead travel travel between two nodes is represented by directed edges. This transformation makes it possible to model realistic directional travel times, align with AMCS data and software systems, and use established optimisation methods for ATSP problems. A similar transformation has been adopted in urban waste collection studies [12].

2.2. Visual Attractiveness: Concepts, Metrics, and Practical Considerations

The concept of visual attractiveness was introduced in vehicle routing research to capture the observation that operators may reject cost-optimal plans if they appear counter-intuitive or when routes cross [8]. This is motivated by both practical and human factors [11]. In practical contexts such as waste collection, compact and non-overlapping routes reduce confusion and simplify operations. In some cases, such as gated communities, compactness and non-overlapping are even required to preserve clear territorial boundaries. Experimental studies on the travelling salesman problem (TSP) further show that people naturally prefer clustered, crossing-free routes [6], [7]. This helps explain why such routes are perceived as more logical even when they do not strictly minimise travel time. Visual attractiveness has primarily been studied in vehicle routing problems, where a set of routes partitions the service area. In this setting, it is typically characterised by three interrelated components, namely compactness, non-overlapping, and low complexity [11].

Compactness reflects the degree to which the stops assigned to a route are spatially concentrated. Routes that serve geographically clustered stops are perceived as more logical and easier to manage. Compactness is typically operationalised through geographical compactness and proximity measures [11]. Geographical compactness is quantified either by distance-to-centre metrics, which sum the distances from each stop to a route centre [8], [10], [20], or by pairwise-distance measures, which aggregate distances between all stops assigned to the same route [8], [21]. Distances are often computed using Euclidean or Manhattan metrics, although travel times or distances can also be used. Route centres are typically defined as geometric centroids, the stop closest to that centroid, or as a medoid minimising the total distance to other stops. In addition, proximity metrics evaluate whether stops are assigned to the geographically nearest route centre, penalising allocations to non-nearest routes [20].

Non-overlapping refers to the spatial separation between different routes. For multiple routes, visually attractive solutions avoid spatial crossings and route overlap, thereby creating clearly defined service territories. This clarifies which vehicle is responsible for which area and makes routes easier to follow. This is typically assessed using metrics such as convex hull overlap, which counts the number of stops that fall within the convex hulls of multiple routes [8], [9], [10], or inter-route crossings, which count how often the route paths geometrically intersect [8]. This inter-route crossings metric is commonly computed using straight-line connections between stops rather than the actual street network, providing a computationally efficient geometric proxy for visual interference between routes. Both compactness and non-overlapping are inherently multi-route concepts, as they require comparison between distinct service territories.

In contrast, low complexity concerns the geometric structure of an individual route. Routes with many sharp turns, self-crossings, or irregular zigzag patterns tend to appear confusing and less intuitive. Metrics such as bending energy or intra-route crossings capture this aspect [8], [11]. Bending energy measures the turn angles along a route. Similar to inter-route crossings, intra-route crossings are often evaluated using straight-line connections. Unlike compactness and non-overlapping, low complexity is applicable to both multi-route and single-route optimisation problems, as it evaluates the internal structure of a single route.

Visual attractiveness metrics have also been applied in waste collection contexts. Some studies address compactness and overlap through clustering-first approaches, using distance-to-centre compactness and convex hull overlap as guiding criteria [9], [10]. Others incorporate spatial considerations more directly in optimisation models. For instance, by limiting the geographic spread in the objective function or enforcing sector connectivity to reduce overlap of routes [18], [22]. These applications demonstrate that visual attractiveness can be operationalised through route construction or dedicated objective terms and constraints.

From a practical perspective, empirical evidence suggests that many visual attractiveness metrics are strongly

correlated. In particular, compactness and overlap measures often capture similar structural properties, indicating that a limited set of complementary metrics may suffice to effectively represent visual attractiveness [11]. Moreover, metrics differ in data and computational requirements. Compactness and convex hull measures can be computed from stop coordinates alone, whereas crossing- and turn-based metrics require information on route geometry. Although approximations using stop coordinates for crossing-based metrics are computationally efficient, they provide only proxies for operational conflicts. In addition, the bending energy cannot be approximated in this way. These considerations highlight the importance of developing metrics that are both informative and computationally tractable. While most existing formulations target multi-route problems, their underlying principles remain relevant. In this study, these principles are adapted to define visual attractiveness for a single route using a side loader.

2.3. Multi-Objective Optimisation Methods

Operational efficiency in waste collection routing is commonly quantified using travel cost, distance, or time. Distance minimisation is frequently used as a proxy for fuel use and emissions [12], [13], while cost objectives aggregate labour, vehicle, and operational factors into monetary terms [5]. In practice, travel time is often the primary feasibility measure, as it directly determines whether a route can be completed within a shift. The fastest route can differ substantially from the shortest route due to speed limits and traffic conditions. Industry examples further illustrate that minimising travel time can improve margins by allowing more collection per crew per day [14]. For these reasons, travel time is used as an efficiency objective in this study.

Routes optimised for efficiency objectives often differ substantially from those optimised for visual attractiveness [11]. This motivates a multi-objective optimisation method to jointly optimise operational efficiency and visual attractiveness. Multi-objective optimisation can be addressed using scalarisation methods, criteria-based methods, or dominance-based approaches [23]. In scalarisation, multiple objectives are combined into a single function, most commonly through a weighted-sum formulation that enables continuous trade-offs via preference weights. An alternative is the ϵ-constraint method, which optimises one objective while bounding others and can be used to construct a Pareto frontier. However, it typically requires solving multiple instances. Criteria-based methods, such as lexicographic optimisation, impose a priority order between objectives and are useful when preferences are strict but offer limited flexibility in exploring trade-offs. Dominance-based methods aim to approximate sets of non-dominated solutions, but usually require population-based metaheuristics and higher computational effort. In waste collection routing, weighted sum approaches are commonly used because they are simple to implement and compatible with a wide range of solution methods. For example, two studies combined travel time, compactness, and workload balance using weighted objectives within their clustering algorithm [9], [10], while other work integrates distance and visual attractiveness in a savings-based heuristic [8].

2.4. Optimisation Methods for the Asymmetric Travelling Salesman Problem

The TSP is one of the most widely studied problems in combinatorial optimisation. In its asymmetric variant, the ATSP, arc costs depend on direction, reflecting one-way streets, turn restrictions, and other network constraints. Such asymmetry is common in real-world routing applications, including waste collection [24]. Although many studies address the TSP, optimisation methods are not always directly transferable to the ATSP, as asymmetry invalidates key assumptions underlying classical TSP moves such as route segment reversals. Optimisation methods for the ATSP are commonly discussed in three groups, namely exact methods, heuristics, and metaheuristics. The methods reviewed here primarily target distance or time minimisation, as most visual attractiveness work is studied in multi-route settings.

Exact methods provide provably optimal solutions but scale poorly with problem size because the ATSP is NP-hard [25], [26]. Mixed-integer programming formulations solved via branch-and-bound or branch-and-cut are common approaches. In waste collection, exact solvers have been applied to small and medium instances with reported distance reductions [17], [27], [28]. Notably, one study modelled the ATSP as a directed Hamiltonian path and transformed it into a symmetric TSP [17].

Heuristic methods construct feasible routes quickly without guaranteeing optimality and are commonly divided into construction and improvement heuristics [24]. Construction heuristics generate an initial route using simple rules and are primarily intended to provide a good starting solution for refinement by improvement heuristics or metaheuristics. Greedy construction methods such as nearest neighbour build a route by repeatedly selecting the closest unvisited node, offering high computational efficiency but often limited solution quality. For example, a study combined nearest neighbour with hill climbing [29]. Insertion heuristics instead grow a partial route by inserting nodes at positions that minimise additional cost, typically producing more structured routes

than nearest neighbour. One study proposed randomised arbitrary insertion, which repeatedly removes and reinserts nodes to diversify tours [30]. In other waste collection work nearest, farthest, and arbitrary insertion were combined with local improvement [28]. Assignment-based construction heuristics follow a different principle. They first solve a relaxed assignment problem to obtain a cycle cover and then merge cycles into a single tour, often producing stronger initial solutions than purely greedy approaches [24].

Improvement heuristics refine an existing feasible route through local modifications and are typically implemented as local search procedures [24]. Starting from an initial solution, they explore a neighbourhood of candidate routes obtained by small structural changes and iteratively accept improving moves until a local optimum is reached. Common neighbourhood operators include node-based moves, such as relocate, swap, and Or-Opt, which adjust the position of individual stops. Additionally, arc-exchange moves (e.g., directed k-opt variants) reconnect route segments to reduce travel cost. These operators are computationally efficient and are frequently embedded within broader metaheuristics. For example, one study incorporated relocate and swap operators within a hill-climbing framework [29], while other work applied directed 3-opt moves with "don't look bits" in genetic algorithm metaheuristics to improve ATSP solutions [31], [32]. Similar strategies have been adopted in waste collection routing. A municipal waste collection study embedded node swap and subtour-reversal operators within a genetic algorithm [33], whereas other research applied simulated annealing and tabu search using swap-based neighbourhood moves [34]. These studies show that simple local operators can significantly enhance solution quality within metaheuristics.

Although heuristics are computationally efficient, they often rely on greedy or problem-specific rules and can become trapped in local optima during the search process. To address these limitations, metaheuristics provide general stochastic search frameworks that balance exploration of new regions of the solution space with exploitation of promising routes [23]. These methods are typically classified into single-solution and population-based approaches.

Single-solution methods iteratively improve one route at a time using neighbourhood operators and controlled acceptance criteria. Examples include tabu search, simulated annealing, and variable neighbourhood search. For instance, a study applied variable neighbourhood search to the ATSP using a specialised local search that combined a 3-opt operator and a problem-specific heuristic [35]. In addition, they replaced random shaking with a guided shake that disrupted promising parts of the route to encourage exploration of new regions of the solution space. Another study employed simulated annealing and tabu search with swap-based neighbourhood moves to improve municipal collection routes [34]. In contrast, population-based methods evolve multiple candidate solutions simultaneously through recombination and diversification mechanisms. Genetic algorithms and ant colony optimisation are among the most widely used approaches. In waste collection contexts, these methods have demonstrated performance improvements. For example, one study developed a genetic algorithm to reduce route costs for municipal waste collection in Rome, combining rank selection, two-point crossover, and swap-based mutation to reduce operational costs [36]. Similarly, other work implemented a genetic algorithm using GIS-based data for Athens and found swap mutations particularly effective in reducing route distance [33]. For the same case study, an ant colony optimisation variant minimising route distance achieved stronger improvements [37].

3. Data Description and Network Representation

The formulation and evaluation of the optimisation framework is based on 12 municipal case studies using side loaders obtained from AMCS. This section describes their data structure and network representation. The evaluation is based on two suburban residential waste collection regions, referred to as Set 1 and Set 2. Set 1 comprises five routes in a suburban area in New Zealand, where vehicles operate on the left side of the road. Set 2 consists of seven routes in a suburban area in the United States, where driving occurs on the right side.

Each route forms a case study. Collection orders represent individual service locations along street kerbsides. Each route includes one or two terminals that serve as start and end points, typically located at depots outside the service area. Vehicle capacities and assignments of orders to a route are assumed fixed. Consequently, the optimisation focuses only on determining an efficient and visually attractive service sequence of orders.

To reduce problem size, collection orders are aggregated into directed link clusters. A link cluster represents a kerbside street segment between two intersections and may contain one or multiple orders. Each link cluster must be serviced in the permitted traffic direction, from its first to its last order. Street segments requiring service on both kerbsides are represented by two directed link clusters, while street segments serviced on only one kerbside correspond to a single link cluster. An example of this aggregation is shown in Figure 1, where each link cluster is represented by its midpoint, computed as the average coordinates of its associated orders. The size of the 12 case studies ranges from 138 to 398 link clusters.

The routing problem is formulated as an ATSP defined over link clusters and terminals. Each node represents

either a terminal or a directed link cluster. The route begins at a start terminal, visits every link cluster exactly once, and ends at the end terminal. Directed arcs are defined from the start terminal to all link clusters, between every pair of link clusters, and from each link cluster to the end terminal. If only one terminal is present, separate start and end nodes are defined at the same location.

Arc travel times are derived from the underlying road network using a fastest-path algorithm that respects traffic directions and prohibits U-turns unless unavoidable against a penalty of 300 seconds. For each ordered pair of link clusters (c, d), travel time corresponds to the fastest feasible path from the last collection order of link cluster c to the first collection order of link cluster d. Thus, arcs represent aggregated travel times rather than individual street segments, preserving realistic network constraints while maintaining computational tractability.

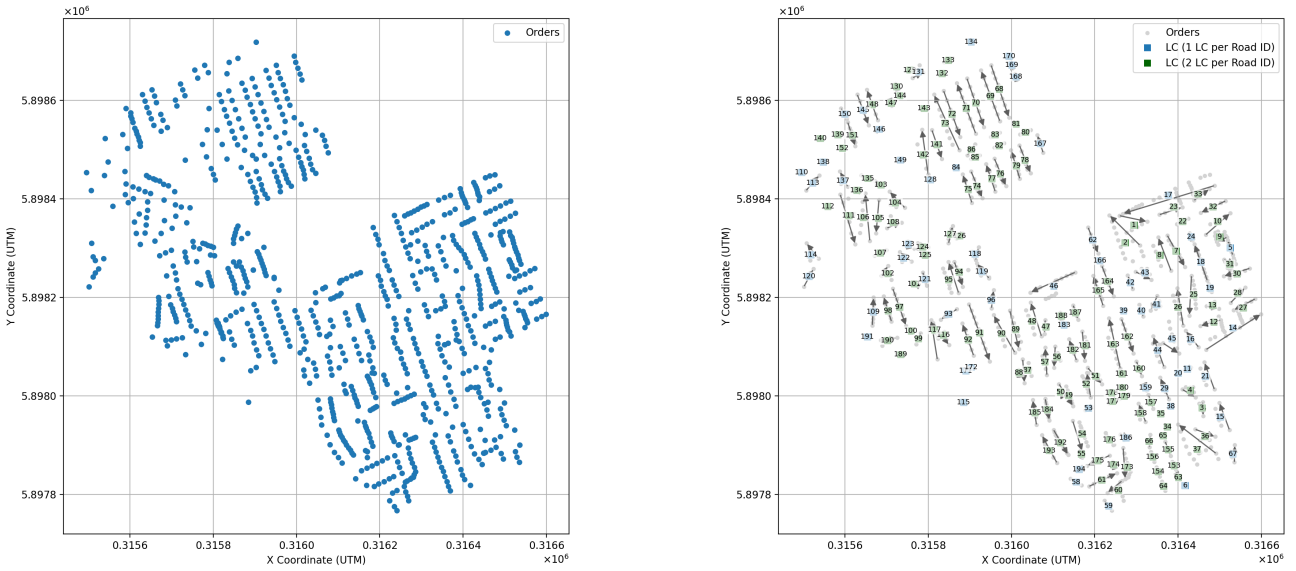

(a) Collection Orders

(b) Aggregation of Collection Orders to Link Clusters

Figure 1: Dataset Representation on the Abstract Network for Route 1 in Set 1 (New Zealand)

4. Metric Formulation and Mathematical Model

This section defines the proposed visual attractiveness metrics for a single route using a side loader. Additionally, it presents the corresponding Mixed Integer Linear Programming (MILP) formulation that combines them with travel time in a weighted sum objective function.

4.1. Visual Attractiveness for a Single Side Loader Route

Visual attractiveness for multiple routes reflects compact, non-overlapping, and simple route structures. For a single route using a side loader, these concepts translate differently. Territorial separation is irrelevant and overlap instead shows as revisiting previously serviced areas within the same route. Therefore, non-overlapping is implicitly captured by compactness. Low complexity remains directly applicable as it concerns the internal structure of a single route. Therefore, visual attractiveness consists of two components in this study:

1. **Compactness:** Routes should concentrate service within a limited area before moving on and nearby link clusters should be served together, avoiding detours and returns to previously serviced streets or areas. Compactness is considered at two scales: *regional compactness*, encouraging one area to be fully serviced before moving to another, and *local compactness*, limiting long moves between consecutive link clusters.
2. **Low complexity:** Routes should have a simple and intuitive shape without excessive zigzag movements and crossing over themselves.

These components provide several practical benefits in residential waste collection. Regional compactness improves operational robustness by concentrating service activity within a limited area, so that disruptions such as breakdowns or delays leave unfinished work confined to a small region. Both regional and local compactness also improve service predictability for residents by ensuring that nearby streets and opposite kerbsides are serviced within a short time window. This reduces situations in which vehicles unexpectedly re-enter previously serviced areas and contributes to safer neighbourhood operations, particularly in suburban streets shared with pedestrians

and children. Finally, routes with low complexity are easier for drivers to interpret and execute, improving operational clarity and reducing navigation errors.

Although these visual attractiveness components are qualitative, they can be approximated using quantitative metrics. This study developed four complementary metrics that capture regional compactness, distance-based and side-loader-specific local compactness, and route complexity. All metrics are evaluated only on link clusters, so terminals and terminal arcs are excluded.

- **Neighbourhood fragmentation (NF)** captures regional compactness as the total penalty associated with re-entries into the same neighbourhood, encouraging contiguous servicing of neighbourhoods. Neighbourhoods are defined by clustering link clusters into spatial groups based on undirected Euclidean distances between their midpoints. Re-entries are penalised more heavily the later they occur in the sequence of neighbourhood visits.
- **Local jumps (LJ)** captures local compactness as the total directed Euclidean distance between consecutive link clusters along the route, encouraging shorter average movements between successive link clusters.
- **Revisit span (RS)** captures local compactness specific to side loaders as the total penalised time between servicing both kerbsides of the same street segment. Only the excess time beyond a predefined tolerance is penalised, so short revisit intervals are permitted.
- **Intra-route crossings (IC)** captures low complexity as the total number of geometric self-intersections along the route. Crossings are detected using straight-line connections between link cluster endpoints. This gives a computationally efficient approximation of route complexity without requiring driving path data and road network geometry [8].

For all metrics, lower values indicate a better score. NF is based on undirected Euclidean distances between link cluster midpoints, whereas LJ uses directed Euclidean distance between link cluster endpoints. Travel time (TT) is used to evaluate operational efficiency and revisit spans. However, travel time or distance are not used as a proxy for spatial compactness, as road network constraints may distort spatial proximity. For example, two link clusters on opposite kerbsides may be close in Euclidean distance but far in travel time or distance.

4.2. Mathematical Model

The problem is formulated on a directed graph $G = (V, A)$, where V contains the terminal and link cluster nodes and A contains all feasible directed arcs. Table 1 summarises the notations used. Link clusters have fixed service times st_c, while terminals have $st_{z_s} = st_{z_e} = 0$. Neighbourhood labels $\nu(c)$ required for neighbourhood fragmentation are obtained using the k-means clustering procedure described in Appendix A, as spatial neighbourhoods are not given in the data. Since each link cluster represents a kerbside street segment between two intersections, revisit spans are evaluated across street segments rather than entire streets. Let R denote the set of revisit pairs, where both link clusters correspond to opposite kerbsides of the same street segment. Set $\mathcal{X}$ contains all ordered pairs of directed link cluster arcs $\big((c,d),(p,q)\big)$ whose straight lines between link cluster endpoints intersect, excluding arcs incident to terminals. Crossings are detected geometrically using the standard orientation test from computational geometry [38]. Two route segments intersect if their endpoints lie on opposite sides of each other's supporting lines.

Table 1: Mathematical Notation

Symbol	Description	Domain	Unit	Metrics[a]
Sets and indices				
C	Set of link cluster nodes ($\lvert C\rvert = m$)	$c \in C$	–	All
Z	Set of start and end terminal nodes ($Z = \{z_s, z_e\}$)	$z \in Z$	–	All
V	Set of all nodes ($V = C \cup Z;\ \lvert V\rvert = n$)	$i \in V$	–	All
A	Set of directed arcs ($A = \{(z_s,d) : d \in C\} \cup \{(c,d) : c,d \in C, c \neq d\} \cup \{(c,z_e) : c \in C\}$)	$(i,j) \in A$	–	All
H	Set of neighbourhoods	$h \in H$	–	NF
R	Set of unordered link cluster pairs on the same street segment	$(c,d) \in R$	–	RS
$\mathcal{X}$	Set of pairs of directed link cluster arcs whose straight lines between endpoints intersect	$\big((c,d),(p,q)\big) \in \mathcal{X}$	–	IC
Parameters				

Continued on next page

Table 1: Mathematical Notation *(Continued)*

Symbol	Description	Domain	Unit	Metrics[a]
tt_{ij}	Travel time from node i to node j	$tt_{ij} \in \mathbb{R}_{\geq 0}$	$[s]$	TT, RS
st_i	Service time for node i	$st_i \in \mathbb{R}_{\geq 0}$	$[s]$	RS
d_{cd}	Undirected Euclidean distance between link cluster midpoints c and d	$d_{cd} \in \mathbb{R}_{\geq 0}$	$[m]$	NF
dd_{cd}	Directed Euclidean distance from the end order node of link cluster c to the start order node of link cluster d	$dd_{cd} \in \mathbb{R}_{\geq 0}$	$[m]$	LJ
$\nu(c)$	Neighbourhood label of link cluster c	$\nu(c) \in H$	–	NF
δ_{cd}	Neighbourhood transition indicator, equals 1 if $\nu(c) \neq \nu(d)$; 0 otherwise	$\delta_{cd} \in \{0, 1\}$	–	NF
E_i	Earliest feasible arrival time at node i	$E_i \in \mathbb{R}_{\geq 0}$	$[s]$	RS
L_i	Latest feasible arrival time at node i	$L_i \in \mathbb{R}_{\geq 0}$	$[s]$	RS
M_{ij}^{LB}	Arc-specific big-M for lower bound time propagation	$M_{ij}^{\mathrm{LB}} \in \mathbb{R}_{\geq 0}$	$[s]$	RS
M_{ij}^{UB}	Arc-specific big-M for upper bound time propagation	$M_{ij}^{\mathrm{UB}} \in \mathbb{R}_{\geq 0}$	$[s]$	RS
Δ_{RS}	Revisit span tolerance	$\Delta_{RS} \in \mathbb{R}_{\geq 0}$	$[s]$	RS
Decision variables				
x_{ij}	Equals 1 if node j is visited immediately after node i; 0 otherwise	$x_{ij} \in \{0, 1\}$	–	All
u_i	Order position of node i (MTZ formulation)	$u_i \in \{0, \ldots, n-1\}$	–	All
q_c	Neighbourhood transition index for link cluster c	$q_c \in \{0, \ldots, \lvert C \rvert\}$	–	NF
F_h	First neighbourhood transition index for neighbourhood h	$F_h \in \{0, \ldots, \lvert C \rvert\}$	–	NF
L_h	Last neighbourhood transition index for neighbourhood h	$L_h \in \{0, \ldots, \lvert C \rvert\}$	–	NF
t_i	Arrival time at node i	$t_i \in [E_i, L_i]$	$[s]$	RS
r_{cd}	Excess revisit span between two link clusters c and d on the same street segment above revisit tolerance	$r_{cd} \in \mathbb{R}_{\geq 0}$	$[s]$	RS
$\chi_{cd,pq}$	Equals 1 if route uses link cluster arcs (c, d) and (p, q) and the corresponding straight lines between endpoints intersect; 0 otherwise	$\chi_{cd,pq} \in \{0, 1\}$	–	IC

Note. $\mathbb{R}_{\geq 0}$ denotes the set of non-negative real numbers, $\mathbb{Z}_{\geq 0}$ denotes the set of non-negative integers, $[a, b]$ denotes the closed interval of real numbers from a to b, and $\{a, \ldots, b\}$ denotes the discrete set of consecutive integers from a to b.
[a] TT: travel time; NF: neighbourhood fragmentation; LJ: local jumps; RS: revisit span; IC: intra-route crossings.

The five objective components are defined in (a)–(e). Equation (a) minimises the total travel time of the route, including movements from the start terminal to link clusters, between link clusters, and from link clusters to the end terminal Since all nodes are visited exactly once, the total service time is constant and is omitted from the metric. Additionally, equation (b) penalises neighbourhood fragmentation by minimising the total span between the first and last occurrences of each neighbourhood along the route. Equation (c) minimises the total directed Euclidean distance between consecutive link clusters, penalising long moves on average. Moreover, equation (d) minimises the total excess revisit time between revisit pairs beyond an allowed revisit tolerance. Finally, equation (e) minimises the total number of geometric self-intersections of the route.

$$\text{Travel Time (TT):} \quad \sum_{d \in C} tt_{z_s d}\, x_{z_s d} + \sum_{c \in C} \sum_{\substack{d \in C \\ d \neq c}} tt_{cd}\, x_{cd} + \sum_{c \in C} tt_{cz_e}\, x_{cz_e} \tag{a}$$

$$\text{Neighbourhood Fragmentation (NF):} \quad \sum_{h \in H} (L_h - F_h) \tag{b}$$

$$\text{Local Jumps (LJ):} \quad \sum_{c \in C} \sum_{\substack{d \in C \\ d \neq c}} dd_{cd}\, x_{cd} \tag{c}$$

$$\text{Revisit Span (RS):} \quad \sum_{(c,d) \in R} r_{cd} \tag{d}$$

$$\text{Intra-Route Crossings (IC):} \quad \sum_{\big((c,d),(p,q)\big) \in \mathcal{X}} \chi_{cd,pq} \tag{e}$$

Metrics (a)–(e) are combined using a weighted sum objective, where w_k denotes the associated weight:

$$\min \; F = w_{\mathrm{TT}} TT + w_{\mathrm{NF}} NF + w_{\mathrm{IC}} IC + w_{\mathrm{LJ}} LJ + w_{\mathrm{RS}} RS \tag{0}$$

The routing constraints ensure that the route forms a path from the start terminal z_s to the end terminal z_e, visiting each link cluster exactly once. constraints (1) and (2) impose the degree conditions for the link clusters, each requiring exactly one incoming and one outgoing arc. Constraints (3) and (4) enforce the degrees for the start and end terminals, requiring one outgoing and incoming arc, respectively.

$$\sum_{\substack{j \in V \\ j \neq c, j \neq z_s}} x_{cj} = 1 \qquad \forall c \in C \tag{1}$$

$$\sum_{\substack{i \in V \\ i \neq c, i \neq z_e}} x_{id} = 1 \qquad \forall d \in C \tag{2}$$

$$\sum_{d \in C} x_{z_s d} = 1 \tag{3}$$

$$\sum_{c \in C} x_{c z_e} = 1 \tag{4}$$

Subtour elimination is enforced using the MTZ formulation, which assigns each node $i \in V$ an order position u_i. Constraints (5) and (6) fix the positions of the start and end terminals to reduce symmetry, while constraint (7) ensures consistent ordering and prevents disconnected subtours.

$$u_{z_s} = 0 \tag{5}$$

$$u_{z_e} = n - 1 \tag{6}$$

$$u_i - u_j + n\, x_{ij} \leq n - 1 \qquad \forall i, j \in V \setminus \{z_s\},\ i \neq j \tag{7}$$

Neighbourhood fragmentation penalises leaving a neighbourhood and returning later. Instead of penalising each neighbourhood switch directly, the model propagates a neighbourhood transition index q_c along the route. Constraints (8)–(9) increase this index only when a selected arc (c, d) involves a neighbourhood change ($\delta_{cd} = 1$). Otherwise, they are relaxed using a big-M constant $m = |C|$. For each neighbourhood h, constraints (10)–(11) define F_h and L_h as the minimum and maximum transition indices of its link clusters. If a neighbourhood is served in one block, the difference $L_h - F_h = 0$. If the routes leaves and re-enters the neighbourhoods, this span increases. Later re-entries incur larger penalties due to the cumulative index.

$$q_d \geq q_c + \delta_{cd} - m(1 - x_{cd}) \qquad \forall c, d \in C,\ c \neq d \tag{8}$$

$$q_d \leq q_c + \delta_{cd} + m(1 - x_{cd}) \qquad \forall c, d \in C,\ c \neq d \tag{9}$$

$$F_h \leq q_c \qquad \forall h \in H,\ c \in C : \nu(c) = h \tag{10}$$

$$L_h \geq q_c \qquad \forall h \in H,\ c \in C : \nu(c) = h \tag{11}$$

Revisit span penalises long time gaps between servicing both kerbsides of the same street segment. Arrival time variables t_i represent service start times at nodes and are propagated along selected arcs ($x_{ij} = 1$) using travel and service times through constraints (12)–(13). For non-selected arcs, these constraints are relaxed using arc-specific big-M coefficients M_{ij}^{LB} and M_{ij}^{UB}, yielding a tighter linear relaxation than a global big-M. They are defined as $M_{ij}^{\mathrm{LB}} = \max(0, L_i + st_i + tt_{ij} - E_j)$ and $M_{ij}^{\mathrm{UB}} = \max(0, L_j - E_i - st_i - tt_{ij})$. These values represent the minimum slack required to deactivate the respective constraint when the arc is not selected ($x_{ij} = 0$). They are derived from the earliest and latest feasible arrival times E_i and L_i, while negative values are capped at zero. These precomputed bounds E_i and L_i are derived from the minimum and maximum feasible total travel time. For each link cluster c, the earliest arrival time is $E_c = tt_{z_s c}$, while the latest feasible arrival time is $L_c = T_{\max}^{\mathrm{path}} - (st_c + tt_{c z_e})$. Here, $T_{\max}^{\mathrm{path}}$ is the maximum feasible total travel time. For the start terminal, arrival times are fixed to zero ($Ez_s = L_{z_s} = 0$). For the end terminal, they are bounded by the minimum and maximum feasible travel times (($E_{z_e} = T_{\min}^{\mathrm{path}}, L_{z_e} = T_{\max}^{\mathrm{path}}$). Here, $T_{\min}^{\mathrm{path}}$ is the minimum feasible total travel time. For each revisit pair $(c, d) \in R$, constraints (14)–(15) define r_{cd} as the excess absolute difference in arrival times beyond the tolerance Δ_{RS}, so only long revisit intervals are penalised.

$$t_j \ \geq\ t_i + st_i + tt_{ij} - M_{ij}^{LB}\,(1 - x_{ij}) \qquad \forall i \in V, j \in V \setminus \{z_s\}, i \neq j \tag{12}$$

$$t_j \ \leq\ t_i + st_i + tt_{ij} + M_{ij}^{UB}\,(1 - x_{ij}) \qquad \forall i \in V, j \in V \setminus \{z_s\}, i \neq j \tag{13}$$

$$r_{cd} \ \geq\ t_c - t_d - \Delta_{RS} \qquad \forall (c, d) \in R \tag{14}$$

$$r_{cd} \ \geq\ t_d - t_c - \Delta_{RS} \qquad \forall (c, d) \in R \tag{15}$$

Intra-route crossings penalise geometric self-intersections of the route. For each pair of link cluster arcs whose straight-line endpoint connections intersect, decision variable $\chi_{cd,pq}$ is activated by constraint (16) only if both

arcs (c, d) and (p, q) are used in the route. The objective term then counts the number of realised crossings.

$$\chi_{cd,pq} \geq x_{cd} + x_{pq} - 1 \qquad \forall\big((c,d),(p,q)\big) \in \mathcal{X} \tag{16}$$

Finally, constraints (17)–(22) define the domains of all decision variables. Binary and integer variables represent discrete routing and ordering decisions, while continuous variables capture arrival times and revisit spans.

$$x_{ij} \in \{0,1\} \qquad \forall (i,j) \in A \tag{17}$$
$$u_i \in \{0,\ldots,n-1\} \qquad \forall i \in V \tag{18}$$
$$F_h \in \{0,\ldots,|C|\}, \quad L_h \in \{0,\ldots,|C|\} \qquad \forall h \in H \tag{19}$$
$$E_i \leq t_i \leq L_i \qquad \forall i \in V \tag{20}$$
$$r_{cd} \geq 0 \qquad \forall (c,d) \in R \tag{21}$$
$$\chi_{cd,pq} \in \{0,1\} \qquad \forall\big((c,d),(p,q)\big) \in \mathcal{X} \tag{22}$$

5. Optimisation Framework

Since the routing problem with competing objectives is NP-hard, exact optimisation becomes computationally infeasible for instances with hundreds of link clusters. Therefore, an optimisation framework based on simulated annealing (SA) is developed. SA is well suited to problems with multiple metrics combined in a weighted sum objective function, where improving one metric may deteriorate others. Greedy or targeted local search would get stuck in local optima. SA addresses this limitation by allowing the probabilistic acceptance of worsening moves, enabling the search to escape local optima and explore the solution space more effectively. Over time, the search gradually shifts from exploration to exploitation as the acceptance of worsening moves becomes less likely. An overview of the algorithm steps is provided in Figure 2, with a detailed pseudocode in Appendix B.

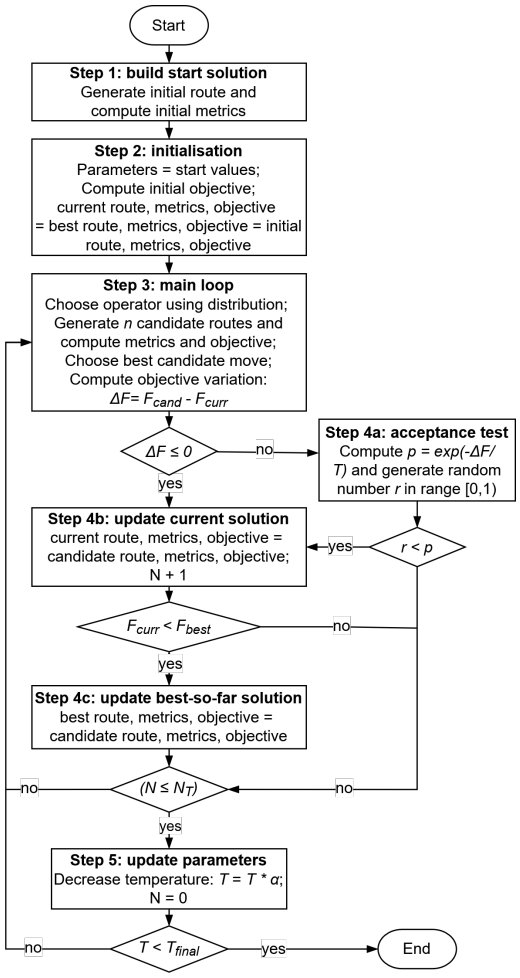

Figure 2: Overview of the SA Optimisation Framework

5.1. Initialisation and Objective Evaluation

The optimisation is initialised with a route that is (near-)optimal with respect to travel time, obtained by solving the exact travel time minimisation model. This model minimises the travel time objective (a) subject to constraints (1)–(7) and constraints (17)–(18). For this initial route, the five objective components (i.e., TT, NF, RS, IC, LJ) are computed. Route quality is assessed using the weighted objective function in equation (0). Let $\hat{w}_k$ denote preference weights for the associated metrics. To ensure comparability across metrics, the preference weights are normalised using the initial metric values $m_{k,0}$ as

$$w_k = \frac{\hat{w}_k}{m_{k,0}}.$$

At the start of the search, both the current solution and the best-so-far solution are set equal to the initial route, including its link cluster visitation sequence, metric values, and objective value. Additionally, the preference weights and all parameters are set to their predefined values, as specified in the following subsections.

5.2. Acceptance Rule, Cooling Schedule, and Stopping Criterion

At the start of the search, the current temperature T is set to the initial temperature T_0. During the optimisation loop, the framework iteratively applies random local search operators to generate alternative routes. In each iteration, an operator is selected according to a predefined probability distribution and used to generate a set of $n = 10$ candidate routes. All candidates are evaluated using the weighted objective function and the best candidate is selected for consideration (i.e., with lowest objective value). Sampling multiple candidates reduces sensitivity to single random moves while preserving the stochastic exploration. Let F_{cur} denote the objective of the current route and F_{cand} that of the selected candidate. The change in objective function value is

$$\Delta F \;=\; F_{\text{cand}} - F_{\text{cur}}.$$

If $\Delta F \leq 0$, the candidate is accepted deterministically. Otherwise, it is accepted with probability

$$p \;=\; \exp\!\left(-\frac{\Delta F}{T}\right),$$

where T is the current temperature. Accepted solutions replace the current solution, while a separate best-so-far solution is maintained and updated whenever an accepted solution improves the best objective value encountered. The temperature follows a geometric cooling schedule $T \leftarrow \alpha T$, where α denotes the cooling factor. For each temperature level T, N_T iterations are performed before cooling. This allows the search to sufficiently explore the neighbourhood at a given temperature before gradually shifting towards more exploitative behaviour. The procedure terminates when the current temperature T falls below the final temperature T_{final}.

5.3. Local Search Operators

At each iteration, a neighbourhood operator is sampled from a fixed probability distribution and applied to the current route to generate candidate solutions. These operators modify the visitation sequence of link clusters while preserving the feasibility of the route. The framework relies on simple random operators, avoiding the need for highly specialised problem-specific rules. Although candidate moves are generated randomly, their evaluation and acceptance are guided by the objective function. This ensures that the search favours routes with improved overall quality. Three local search operators are used:

1. **Relocate:** remove one link cluster from its current position and insert it at another position in the route.
2. **Swap:** exchange the positions of two link clusters in the route.
3. **Or-Opt:** relocate a block of three consecutive link clusters to a different position in the route.

These operators are widely used in routing heuristics and provide simple but effective modifications of the route [23]. They modify the route without reversing parts of the route, which makes them suitable for asymmetric routing problems such as the ATSP, where arc costs depend on direction. In contrast, operators such as 2-opt reverse route segments and therefore perform poorly in asymmetric settings. Larger arc-exchange operators, such as extended k-opt moves, were not considered because they substantially increase computational effort. The operator probability distribution is fixed at 0.4 for relocate and 0.3 for both swap and Or-Opt. Relocate is slightly favoured as it corrects small inefficiencies, while swap and Or-Opt enable slightly larger adjustments.

5.4. Parameter Configuration

All computational experiments follow a fixed parameter configuration. For neighbourhood fragmentation, k-means clustering is applied with a target neighbourhood size of 50 link clusters. The revisit span tolerance is fixed at 600 seconds. The simulated annealing parameters are set as follows. The initial temperature is $T_0 = 0.5$, the cooling factor $\alpha = 0.99$, and the final temperature $T_{\text{final}} = 10^{-4}$. For each temperature level, N_T iterations are performed, where N_T depends on instance size. Consequently, $N_T = 50$ for 100–200 link clusters, $N_T = 70$ for 200–300, and $N_T = 90$ for 300–400. At each iteration, $n = 10$ candidate solutions are sampled. The relocate, swap, and Or-Opt operators are selected with probabilities 0.4, 0.3, and 0.3, respectively, and the Or-Opt block length is fixed to three. The random seed is fixed to 42 to ensure reproducibility.

Preference weights are fixed to $\hat{w}_{\text{TT}} = 2.0$, $\hat{w}_{\text{NF}} = 1.5$, $\hat{w}_{\text{RS}} = 1.0$, $\hat{w}_{\text{IC}} = 0.5$, and $\hat{w}_{\text{LJ}} = 0.5$. Travel time receives the highest weight, as it is the only metric directly derived from the directed road network and therefore ensures operational feasibility. Neighbourhood fragmentation is weighted strongly because it enforces regional compactness by grouping link clusters into contiguous service areas. Revisit span and local jumps improve compactness within neighbourhoods, with revisit span weighted higher due to its relevance for side loader operations. Intra-route crossings and local jumps rely on straight-line approximations and do not fully reflect road network constraints. Therefore, they are assigned low weights and mainly act as tie-breaking criteria. The selected values represent a practical baseline rather than universally optimal preferences.

6. Results

The optimisation framework is evaluated on the 12 real case studies introduced in section 3. First, numerical results across all case studies are presented and illustrated through a visual inspection for one representative case study. These results are compared against the current operational routes of AMCS. Second, a sensitivity analysis for the same case study examines how uniformly scaling the visual attractiveness preference weights affects the trade-off between travel time and visual attractiveness. All experiments were implemented in Python. Initial routes were obtained by solving the exact model that minimises travel time using the Gurobi Optimiser (version 12) with a time limit of 300 seconds. Optimal solutions were found for four instances, while the remaining case studies terminated with small optimality gaps, ranging from 0.008 to 8.6%. All computations were performed on a computer with a 3.6 GHz 8-core processor and 32 GB RAM.

6.1. Results of Case Studies

Table 2 reports the numerical comparison between the routes of AMCS optimal for travel time and the routes generated by the optimisation framework across all 12 case studies. For each instance, relative changes in the five route quality metrics are computed with respect to the AMCS baseline. For a given metric M, the relative change is defined as

$$\Delta M_{\%} = \frac{M_{\text{SA}} - M_{\text{AMCS}}}{M_{\text{AMCS}}} \times 100.$$

Negative values indicate an improvement relative to the AMCS solution, while positive values indicate deterioration. In addition, the objective value change is computed relative to the initial solution used to initialise the SA search procedure, with objective value $F_0 = 5.5$, as

$$\Delta F_{\%} = \frac{F_{\text{SA}} - F_0}{F_0} \times 100.$$

Summary statistics, including the median, IQR, minimum, and maximum, are provided to capture overall performance across instances. These measures are preferred over the mean and standard deviation, which can be disproportionately influenced by individual instances in such a limited set of case studies. Runtime values for both approaches are included to indicate computational effort. Runtime differences across case studies primarily reflect differences in instance size and the number of simulated annealing iterations per temperature level.

For 11 of the 12 case studies, the optimisation framework improves the objective value relative to the initial solution, with a median reduction of -30.4% and small variability (IQR 11.8%). In one case study (Route 4 of Set 1), the optimisation framework does not improve on the initial route, which means the reported values correspond to this initial solution.

Across the improved case studies, neighbourhood fragmentation and revisit span substantially decrease with low variability relative to the AMCS baseline (medians -90.2% and -92.0%, both IQR 6.5). This indicates improved regional compactness and side loader specific local compactness, meaning the routes have reduced repeated returns to the same neighbourhood and decreased time intervals between servicing opposite kerbsides.

Intra-route crossings also decrease, but with smaller and less consistent gains (median −47.7%, IQR 37.9). This is consistent with its lower preference weight in the objective function.

These improvements come at the cost of increased travel time (median +40.5%), reflecting the trade-off made between efficiency and visual attractiveness. The objective reductions indicate that, under the chosen preference weights, the visual attractiveness improvements outweigh this increase. Moreover, the local jumps metric increases (median +19.2%), indicating that local jumps and travel time are both distance-based measures between the link cluster endpoints. In contrast, the other metrics capture more structural or specific aspects of the route. As a result, improvements in regional compactness, side-loader-specific local compactness, or route simplicity do not translate into shorter distances between successive link clusters. This highlights a trade-off between the local jumps metric and the other visual attractiveness metrics and is consistent with its relatively low preference weight. Overall, the optimisation framework produces more visually attractive routes than the AMCS baseline, while accepting moderate increases in total travel time and directed Euclidean distance.

Table 2: Relative Changes in Metrics and Objective Value Across Case Studies

Case Study	**ΔTT (%)**	**ΔNF (%)**	**ΔRS (%)**	**ΔIC (%)**	**ΔLJ (%)**	**ΔF (%)**	**Runtime AMCS (s)**	**Runtime SA (s)**
Set 1 – Route 1	45.2	−95.7	−98.2	−55.8	6.9	−30.5	93	7681
Set 1 – Route 2	71.7	−83.3	−83.8	−19.8	34.5	−10.2	193	49 278
Set 1 – Route 3	65.8	−91.5	−92.2	−34.6	46.5	−22.7	133	39 301
Set 1 – Route 4	−1.3	21.7	22.6	−15.4	−1.0	0.0	361	49 812
Set 1 – Route 5	57.8	−88.0	−92.9	−23.9	44.9	−19.5	1156	37 591
Set 2 – Route 1	40.6	−93.3	−91.9	−72.1	21.7	−30.9	5	11 757
Set 2 – Route 2	40.5	−85.5	−94.0	−39.5	15.7	−33.8	32	59 602
Set 2 – Route 3	37.3	−94.1	−82.7	−62.5	24.3	−32.0	4	6624
Set 2 – Route 4	36.0	−92.3	−90.0	−73.1	10.2	−34.5	4	4001
Set 2 – Route 5	39.7	−88.9	−94.4	−69.7	−0.3	−30.2	13	15 777
Set 2 – Route 6	41.4	−87.5	−88.6	−38.2	19.5	−28.0	34	19 582
Set 2 – Route 7	36.4	−98.2	−94.0	−70.0	18.9	−34.7	10	11 627
Median	40.5	−90.2	−92.0	−47.7	19.2	−30.4	–	–
IQR	11.2	6.5	6.5	37.9	17.5	11.8	–	–
Min	−1.3	−98.2	−98.2	−73.1	−1.0	−34.7	–	–
Max	71.7	21.7	22.6	−15.4	46.5	0.0	–	–

Note. All metric changes ΔTT, ΔNF, ΔRS, ΔIC, and ΔLJ are reported relative to the AMCS baseline route for each case study. The objective change $\Delta\mathbf{F}$ is computed relative to the initial solution with objective value $F_0 = 5.5$.

While the numerical results quantify improvements across all instances, visual attractiveness is inherently spatial. To illustrate how the optimisation framework reshapes route structure, a detailed visual comparison is provided for one representative case study, namely Route 1 in Set 1 (New Zealand). This instance is selected as its relative metric changes closely match the median values observed across the 12 case studies.

Figure 3 compares the AMCS baseline route and the route generated by the optimisation framework on the abstract network. Link clusters are shown at their midpoints and labelled by visitation sequence, with grey arrows indicating travel direction. Revisit pairs are marked with green lines, crossing arcs in blue, and terminal nodes and arcs are omitted. The first two subfigures colour link clusters by neighbourhood assignment, whereas the other two subfigures use a colour gradient to represent the visitation sequence.

In the AMCS baseline, frequent transitions between neighbourhoods are visible, leading to repeated returns to previously serviced areas and long cross-neighbourhood movements. In addition, opposite kerbsides of the same street segment are often serviced far apart in the visitation sequence, and several long arcs intersect other parts of the route, indicating higher route complexity.

In contrast, the optimisation framework route shows a more attractive structure. Link clusters within the same neighbourhood are visited more consecutively, reducing neighbourhood fragmentation. Revisit spans between opposite kerbsides are shorter, and fewer long arcs intersect the route, indicating reduced intra-route crossings. These differences visually confirm the strong numerical reductions in neighbourhood fragmentation, revisit span, and intra-route crossings. At the same time, the optimised route contains slightly longer moves between some successive link clusters, consistent with the increase in total directed Euclidean distance. This reflects the trade-off in the weighted objective. Improvements in regional compactness, side-loader-specific local compactness, and route simplicity are achieved at the expense of distance-based local compactness. Note that changes in travel

time are not directly visible in these figures, as arcs represent directed Euclidean distances rather than actual driving paths.

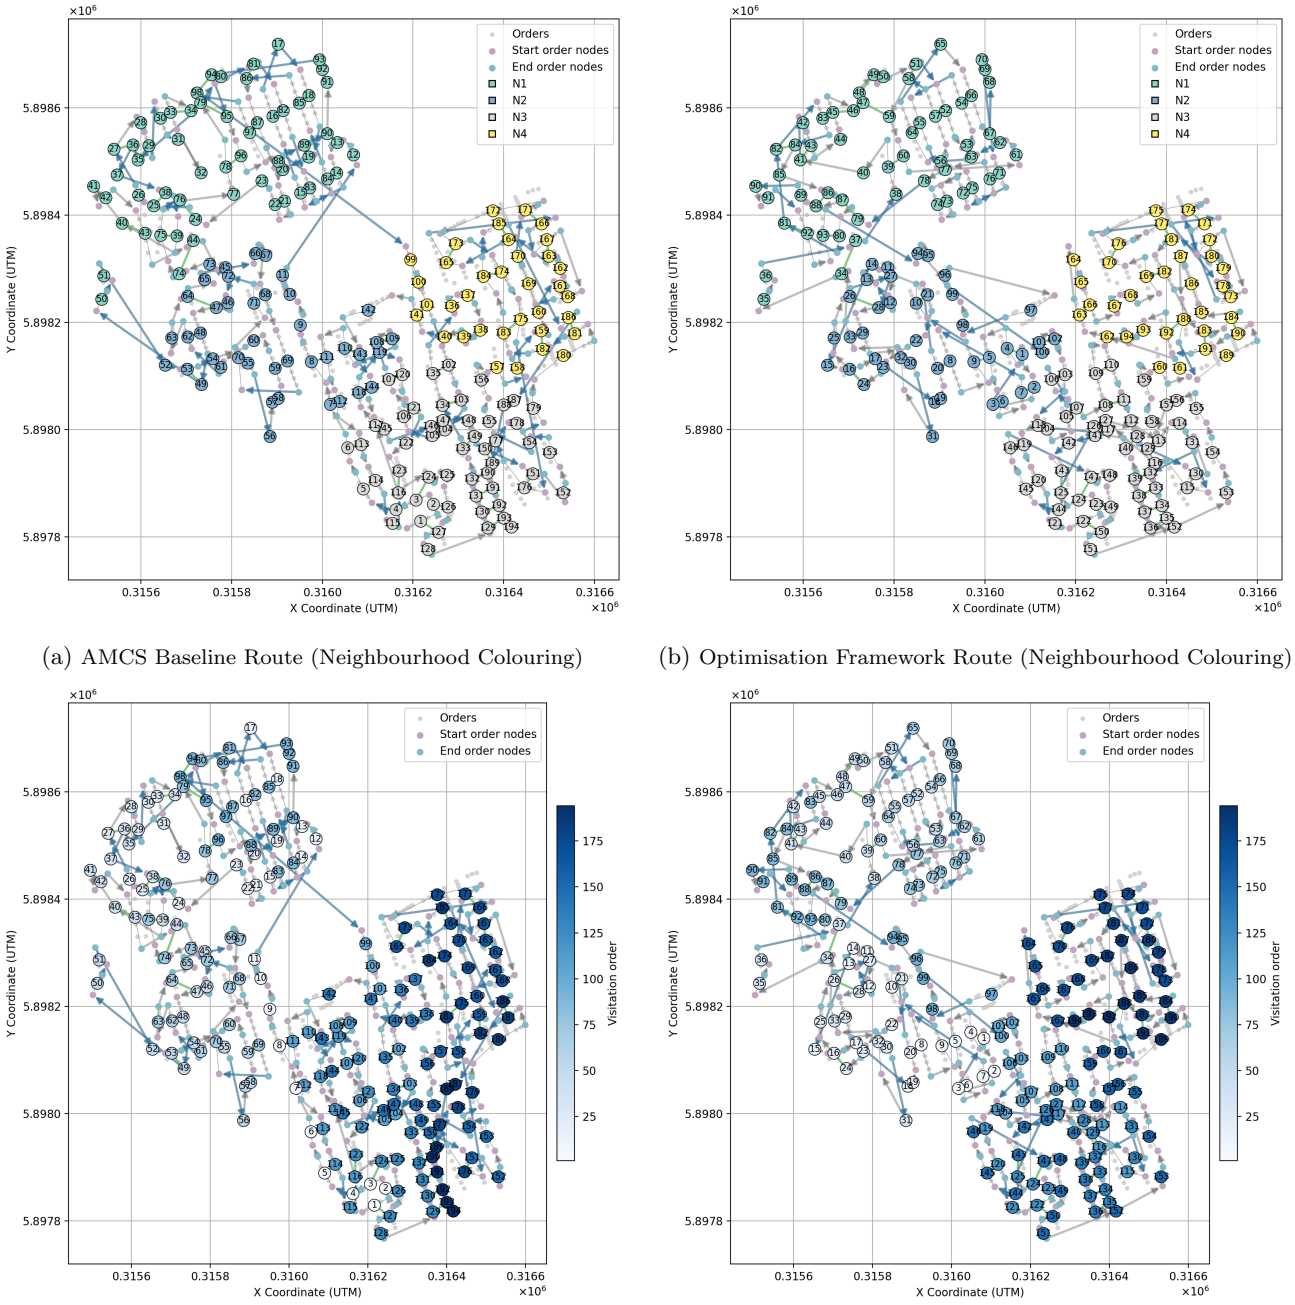

(a) AMCS Baseline Route (Neighbourhood Colouring)

(b) Optimisation Framework Route (Neighbourhood Colouring)

(c) AMCS Baseline Route (Visitation Order Colouring)

(d) Optimisation Framework Route (Visitation Order Colouring)

Figure 3: Comparison of Routing Solutions on the Abstract Network for Route 1 in Set 1 (New Zealand)

To further examine routing behaviour on the actual road network, Figure 4 highlights the first 55 serviced orders for both the AMCS baseline and the optimisation framework. Unlike the abstract network, these figures show the true driving paths and sequence of collection orders on the road network, with nodes coloured by neighbourhood. Within this initial segment, the AMCS baseline traverses multiple neighbourhoods and spans a large part of the service area early in the route. In contrast, the optimisation framework services orders more compactly, largely remaining within a single neighbourhood before transitioning to others and servicing opposite kerbsides closer in time. This visually illustrates the improved regional compactness and reduced revisit span achieved by the optimisation framework. Intra-route crossings are less apparent in the road network visualisation, as this metric is defined on straight-line connections in the abstract network rather than on the actual driving paths.

(a) AMCS Baseline Route (Highlighted First 55 Orders) (b) Optimisation Framework Route (Highlighted First 55 Orders)

Figure 4: Comparison of Routing Solutions on the Road Network for Route 1 in Set 1 (New Zealand)

An example where neighbourhood coherence is achieved in the abstract network but is not fully reflected in the road network is shown in Figure 5. In both routes, the vehicle crosses the light blue neighbourhood while servicing the adjacent dark blue neighbourhood. In the AMCS baseline route, the moves from the red-coloured orders 186 to 187 and from 194 to 195 require switching between neighbourhoods and are therefore penalised. In contrast, the move from the orange-coloured orders 77 to 78 in the optimisation framework route crosses the light blue neighbourhood on the road network. However, this move does not result in a neighbourhood switch in the abstract network and is therefore not penalised.

This example highlights the sensitivity of the neighbourhood fragmentation metric to the chosen neighbourhood definition and the use of straight-line link cluster representations in the metrics. Although this boundary crossing are not necessarily undesirable in practice, as it only passes along the edge of a neighbourhood, it illustrates how metric evaluation depends on modelling abstractions rather than the exact road network geometry.

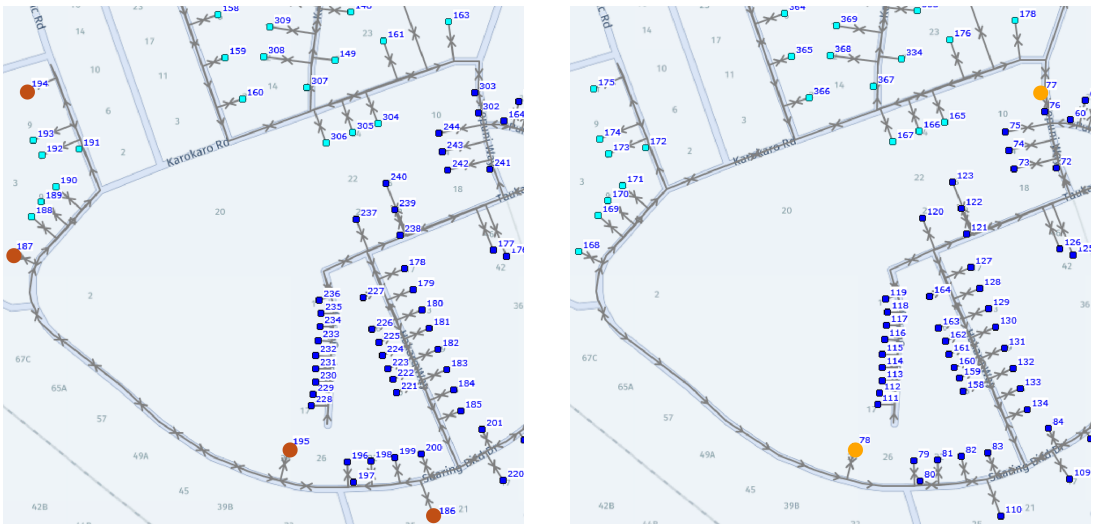

(a) AMCS Baseline Route Illustrating Penalised Neighbourhood Transitions

(b) Optimisation Framework Route Illustrating Unpenalised Neighbourhood Transitions

Figure 5: Example of Neighbourhood Boundary Effects on the Road Network for Route 1 of Set 1 (New Zealand)

6.2. Sensitivity Analysis of Preference Weights

This section investigates the sensitivity of the final routing outcomes to the choice of preference weights used in the objective function. While the final evaluation relies on a fixed and validated set of weights, practical use of the optimisation framework may require adjusting the relative importance of travel time and visual attractiveness. For example, some municipalities or customers may prefer visual attractiveness more over efficiency than others.

Therefore, this sensitivity analysis shows how varying the relative importance of visual attractiveness against travel time affects routes.

The travel time weight is kept fixed, while all visual attractiveness weights are multiplied by a scaling factor $\lambda \in \{0, 0.25, 0.5, 1.0, 2.0, 4.0\}$. The analysis is conducted for the same case study (Route 1, Set 1) used in the visual analysis. The case $\lambda = 0$ represents the initial solution obtained by solving the exact model that minimises travel time without executing the search procedure and is included to explicitly indicate the travel-time-optimal baseline. Figure 6 illustrates the trade-offs between travel time and visual attractiveness metrics. In each plot, the horizontal axis represents the total travel time, while the vertical axis shows the raw total value of a visual attractiveness metric. Each point corresponds to a solution obtained for a specific value of λ.

For $\lambda = 0.25$, the objective value is not improved relative to the initial solution, indicating that the reported values correspond to the initial solution optimal for travel time. Increasing λ results in a higher total travel time. In addition, it leads to substantial reductions in neighbourhood fragmentation and revisit span, particularly when moving from $\lambda = 0.25$ to $\lambda = 0.5$. Beyond this range, further improvements become marginal, while travel time increases more. Intra-route crossings decrease more gradually, which is consistent with its lower weight. In contrast, the local jumps metric shows non-monotonic behaviour. This is consistent with earlier findings and reflects that it is more closely aligned with travel time than with the other metrics. As a result, improvements in neighbourhood fragmentation, street segment revisit behaviour, and route simplicity do not translate into shorter directed Euclidean distance and travel time. Under the chosen preference weights, the optimisation framework therefore prioritises improvements in neighbourhood fragmentation, revisit span, and intra-route crossings over reductions in the local jump metric.

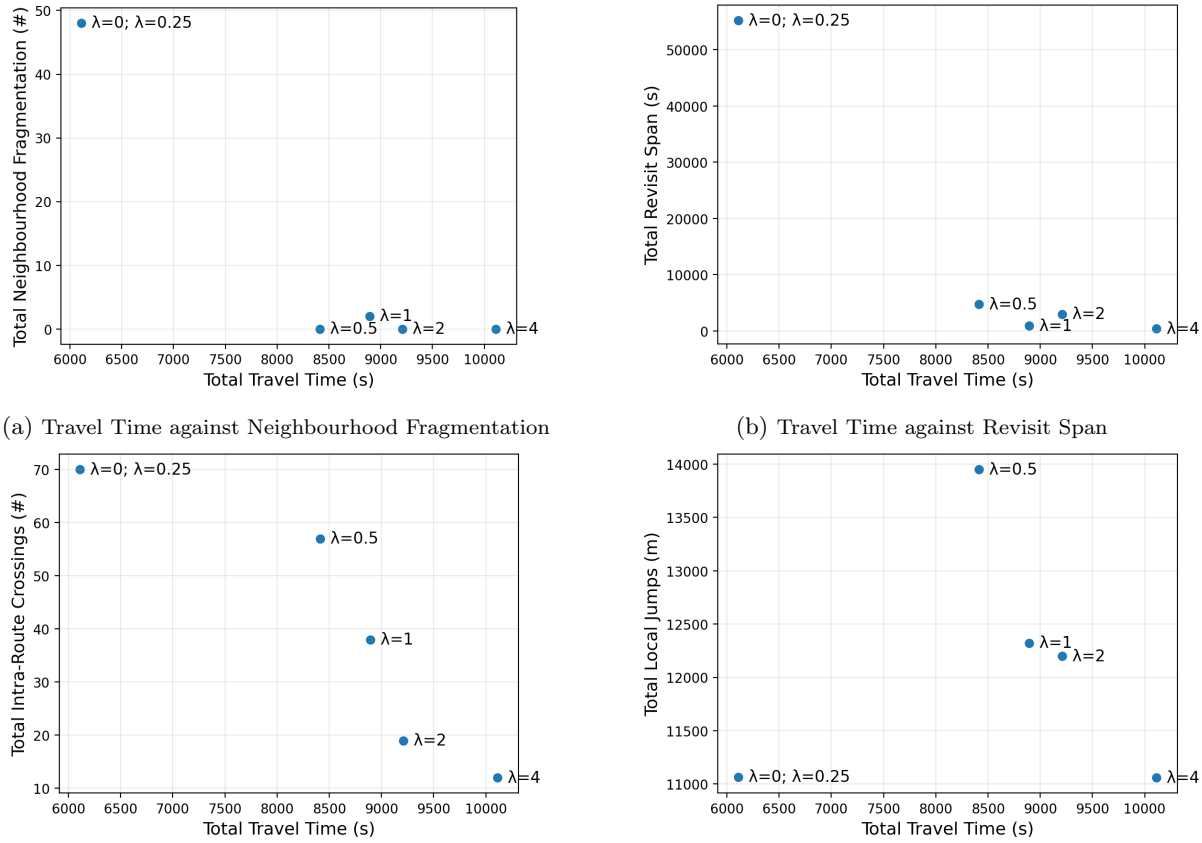

(a) Travel Time against Neighbourhood Fragmentation

(b) Travel Time against Revisit Span

(c) Travel Time against Intra-Route Crossings

(d) Travel Time against Local Jumps

Figure 6: Trade-Offs Between Travel Time and Visual Attractiveness Metrics (Route 1, Set 1)

To better understand the transition between the initial solution that is optimal for travel time and the first improved solution, additional experiments were performed for intermediate values $\lambda \in \{0.30, 0.35, 0.40, 0.45\}$. These values allow for a more detailed analysis of how the optimisation framework behaves when visual attractiveness receives only a small increase in importance relative to travel time. The corresponding trade-off plots are shown in Figure 7.

The results confirm that $\lambda = 0.30$ still corresponds to the initial solution, indicating that travel time remains dominant in the objective function for very small values of λ. For larger values in this range, the optimisation

framework identifies alternative routes that slightly improve the objective value relative to the initial solution. However, the changes in both the objective value and the individual metrics remain relatively small compared to the larger improvements observed when increasing λ to 0.50.

Within this narrow range, the relationship between the metrics is less consistent. In contrast to the clearer trends observed in Figure 6, the solutions do not follow a consistent monotonic pattern. For example, the configuration with $\lambda = 0.35$ achieves lower neighbourhood fragmentation and revisit span than the configuration with $\lambda = 0.40$, while the $\lambda = 0.40$ solution substantially reduces intra-route crossings while maintaining relatively low travel time and local jumps. Similarly, the solution for $\lambda = 0.45$ has fewer intra-route crossings and lower travel time and local jumps than the solution for $\lambda = 0.50$.

This behaviour can be explained by the interaction between the weighted objective and the stochastic local search procedure. At low values of λ, travel time still dominates the objective function and the visual attractiveness metrics mainly act as secondary criteria. As a result, small variations in the search may lead to solutions that prioritise improvements in different metrics.

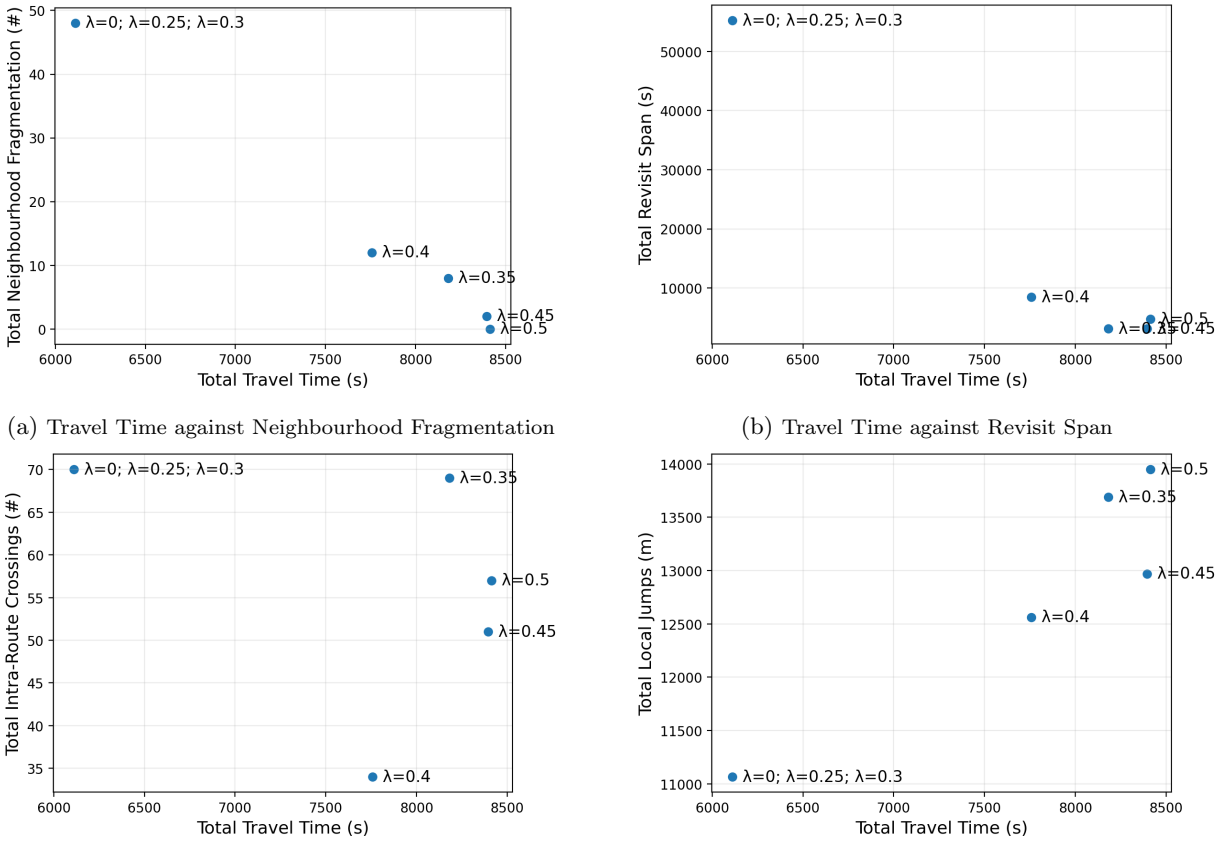

(a) Travel Time against Neighbourhood Fragmentation

(b) Travel Time against Revisit Span

(c) Travel Time against Intra-Route Crossings

(d) Travel Time against Local Jumps

Figure 7: Trade-Offs Between Travel Time and Visual Attractiveness Metrics for Intermediate Values of λ Between 0.25 and 0.5 (Route 1, Set 1)

To complement the trade-off plots, Figure 8 compares the route structures for $\lambda = 1.0$ and $\lambda = 4.0$. The route for $\lambda = 1.0$ corresponds to the default setting analysed in the previous subsection and represents a balanced compromise between efficiency and visual attractiveness. For $\lambda = 4.0$, neighbourhoods are mostly serviced in single blocks, revisit spans between opposite kerbsides are minimal, and intra-route crossings are nearly eliminated. The route appears structured on the abstract network. However, achieving this level of visual attractiveness results in a large increase in travel time. Moreover, very large values of λ may induce undesirable behaviour on the road network, as visual attractiveness metrics are evaluated on an abstract network. Since travel time is the only metric that accounts for the underlying road network through its fastest-path computation, heavily down-weighting it may lead to repeated street visits or unpenalised neighbourhood crossings in practice.

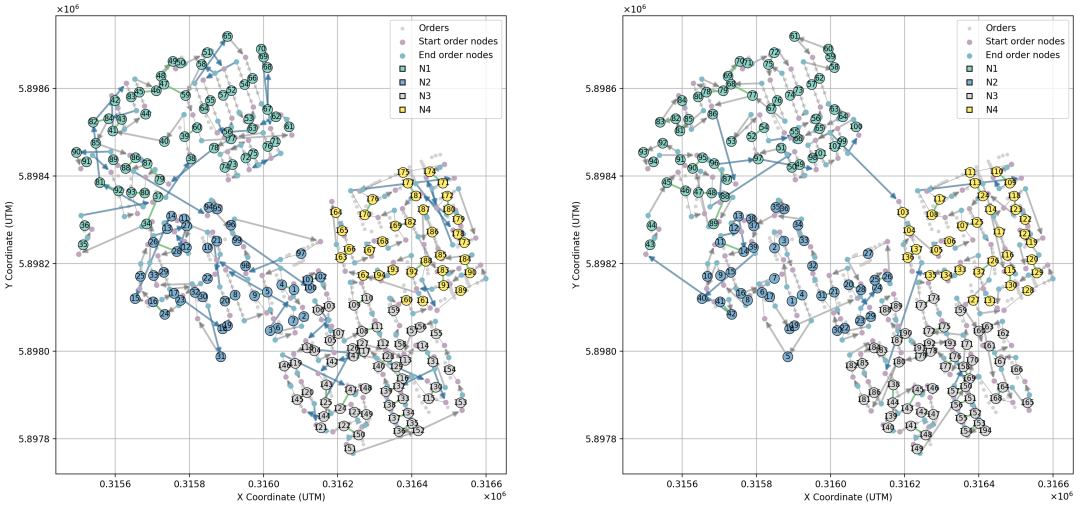

(a) Optimisation Framework Route ($\lambda = 1.0$) (b) Optimisation Framework Route ($\lambda = 4.0$)

Figure 8: Comparison of Routing Solutions on the Abstract Network for Different Values of λ (Route 1, Set 1)

Taken together, the results show that moderate values of λ, such as $\lambda = 1$, already achieve substantial improvements in visual attractiveness compared to the initial solution that is optimal for travel time. The additional analysis for intermediate values between $\lambda = 0.25$ and $\lambda = 0.5$ further shows that very small values of λ lead to less consistent behaviour. In this range, travel time dominates the weighted objective function and the visual attractiveness metrics mainly act as secondary criteria. Consequently, the stochastic search procedure may prioritise improvements in different metrics, resulting in solutions that do not follow a clear trend.

Increasing λ beyond this range yields diminishing improvements in visual attractiveness while causing substantial increases in travel time. This highlights the trade-offs inherent to multi-objective optimisation, where improvements in one metric may occur at the expense of another. Overall, the selected preference weights provide a reasonable balance between travel time efficiency and visual attractiveness for this case study, while allowing adjustment to different operational preferences.

7. Conclusions

This study developed an optimisation framework for residential kerbside waste collection that generates a single route balancing operational efficiency and visual attractiveness. The framework explicitly accounts for operating conditions specific to side loaders, including double traversal of streets, traffic directions, and U-turn restrictions. Double traversal is required because vehicles can collect waste from only one side of the street per pass. Consequently, routes optimised for efficiency are often visually unattractive in side loader operations, as they result in fragmented service of neighbourhoods, long time intervals between servicing opposite kerbsides, and complex route shapes. This reduces operational clarity for drivers and predictability for residents.

The problem was formulated as a node-based optimisation model designed to construct a single route, also known as an ATSP. Double traversal and realistic operating conditions were incorporated through an abstract network representation in which travel times were computed as the fastest paths on the underlying road network, respecting traffic directions and U-turn restrictions. Visual attractiveness was quantified using a set of complementary metrics that capture regional compactness, local compactness, and route complexity. These metrics were integrated with travel time in a weighted sum objective and optimised using simulated annealing.

Empirical evaluation on 12 real municipal case studies demonstrated that the optimisation framework produces more visually attractive routes compared to the travel-time-optimal routes generated by AMCS. Neighbourhood fragmentation, revisit spans between opposite kerbsides, and intra-route crossings were reduced, with median decreases of approximately 90% for neighbourhood fragmentation and revisit span and nearly 50% for intra-route crossings. These improvements were achieved at the cost of a median increase of approximately 40% in total travel time and 20% in total directed Euclidean distance. Thus, the results show that the visual attractiveness metrics may conflict, meaning that improvements in one aspect of visual attractiveness may occur at the expense of another. In addition, a clear trade-off emerges between efficiency and visual attractiveness. The sensitivity analysis confirmed that this trade-off can be controlled through uniform scaling of the visual attractiveness preference weights. However, very small weights lead to less consistent behaviour, as travel time dominates the

objective function and the stochastic search procedure may prioritise different improvements. Conversely, over-emphasising visual attractiveness produces diminishing improvements in visual attractiveness while substantially increasing travel time.

From a scientific perspective, these findings extend existing visual attractiveness concepts from vehicle routing problems to the optimisation of a single side loader route. Compactness and non-overlapping, traditionally defined across multiple routes, are reformulated for a single route setting. Specifically, compactness is evaluated along a single route rather than across route territories. Moreover, territorial separation between routes is no longer relevant. Instead, non-overlapping shows as repeated revisits of the same streets or areas and is therefore treated as an aspect of compactness. Based on this reformulation, the study introduces compactness metrics that evaluate regional and local compactness along a single route, including a revisit span metric that captures behaviour specific to side loaders. The study further shows how these metrics can be integrated with travel time in a weighted sum. Although travel time and intra-route crossings are not novel metrics, their combination with the proposed compactness metrics provides new insight into how efficiency and visual attractiveness can be balanced in the optimisation of a single side loader route.

From a practical perspective, these findings are relevant as municipalities increasingly adopt automated side loaders. The double traversal required by side loaders introduces challenges that are not addressed by traditional routing methods. The optimisation framework demonstrates that more visually attractive routes can be generated compared to current operational routes. For collection crews, compact routes with few self-crossings are easier to interpret and execute, improving operational clarity. In addition, regional compactness enhances robustness to disruptions by concentrating unfinished work within a limited area. For residents, regional and local compactness improve service predictability by ensuring that nearby streets and opposite kerbsides are serviced within a short time window. This reduces situations in which collection vehicles unexpectedly return to previously serviced areas and contributes to safer neighbourhoods. Finally, the weighted sum objective provides routing software providers with a configurable method to balance travel time and visual attractiveness according to client preferences.

Several directions for future research emerge from this study. A first priority concerns the further development of visual attractiveness metrics for optimising a single side loader route. In particular, incorporating detailed road geometry and driving path data could enable more realistic measures of route complexity and improved neighbourhood definitions aligned with the road network. Future work could also investigate alternative combinations of visual attractiveness metrics and analyse the trade-offs between the metrics more explicitly, for example, by approximating the Pareto frontier. A second direction concerns further development of the optimisation framework itself. This includes evaluating the influence of the initial solution, broader parameter tuning on additional datasets, and investigating larger or alternative local search operators. Computational efficiency may also be improved by introducing earlier stopping criteria or by parallelising the evaluation of candidate routes within each iteration. Finally, alternative metaheuristics or hybrid optimisation methods could be explored to assess whether they achieve improved solution quality or runtime.

Finally, several practical recommendations arise for the implementation of visual attractiveness in routing systems such as those developed by AMCS. The most important recommendation is to prioritise compactness-related metrics, particularly neighbourhood fragmentation and revisit span, as these improve operational robustness and service predictability in side loader operations. Furthermore, preference weights should remain configurable so that routing solutions can be calibrated to the priorities of clients. In addition, neighbourhood definitions should reflect administrative or operational boundaries rather than rely on simple clustering approaches. Improving the grouping of orders into street segments (i.e., link clusters) may further enhance the effectiveness of compactness metrics. A related extension would be linking link clusters to streets and introducing a metric that encourages consecutive service of entire streets rather than kerbside street segments. Finally, the modelling of flexible service locations could be explored, such as shared collection points or designated kerbsides. This may reduce double traversal and therefore improve efficiency and visual attractiveness.

References

[1] United Nations Environment Programme, "Global waste management outlook 2024: Beyond an age of waste – turning rubbish into a resource," United Nations Environment Programme (UNEP), Tech. Rep. DTI/2619/NA, 2024. [Online]. Available: https://wedocs.unep.org/20.500.11822/44939.

[2] S. Das and B. K. Bhattacharyya, "Optimization of municipal solid waste collection and transportation routes," *Waste Management*, vol. 43, pp. 9–18, 2015. DOI: 10.1016/j.wasman.2015.06.033.

[3] B. L. Golden, A. A. Assad, and E. A. Wasil, "Routing vehicles in the real world: Applications in the solid waste, beverage, food, dairy, and newspaper industries," in *The Vehicle Routing Problem*, P. Toth and D. Vigo, Eds., SIAM, 2002, pp. 245–286. DOI: 10.1137/1.9780898718515.ch10.

[4] Archive Market Research, "Automated side loader garbage truck market report," *Market Report*, 2024. [Online]. Available: https://www.archivemarketresearch.com.

[5] J. Beliën, L. De Boeck, and J. Van Ackere, "Municipal solid waste collection and management problems: A literature review," *Transportation Science*, vol. 48, no. 1, 2012. DOI: 10.1287/trsc.1120.0448.

[6] I. Van Rooij, U. Stege, and A. Schactman, "Convex hull and tour crossings in the euclidean traveling salesperson problem: Implications for human performance studies," *Memory & Cognition*, vol. 31, no. 2, pp. 215–220, 2003. DOI: 10.3758/BF03194380.

[7] D. Vickers, M. D. Lee, M. Dry, P. Hughes, and J. A. McMahon, "The aesthetic appeal of minimal structures: Judging the attractiveness of solutions to traveling salesperson problems," *Perception and Psychophysics*, vol. 68, no. 1, pp. 32–42, 2006. DOI: 10.3758/BF03193653.

[8] A. Poot, G. Kant, and A. P. Wagelmans, "A savings based method for real-life vehicle routing problems," *Journal of the Operational Research Society*, vol. 53, no. 1, pp. 57–68, 2002. DOI: 10.1057/palgrave.jors.2601265.

[9] S. Sahoo, S. Kim, B.-I. Kim, B. Kraas, and A. Popov Jr, "Routing optimization for waste management," *Interfaces*, vol. 35, no. 1, pp. 24–36, 2005. DOI: 10.1287/inte.1040.0109.

[10] B.-I. Kim, S. Kim, and S. Sahoo, "Waste collection vehicle routing problem with time windows," *Computers & Operations Research*, vol. 33, no. 12, pp. 3624–3642, 2006. DOI: 10.1016/j.cor.2005.02.045.

[11] D. Rossit, D. Vigo, F. Tohmé, and M. Frutos, "Visual attractiveness in routing problems: A review," *Computers & Operations Research*, vol. 103, pp. 13–34, 2019. DOI: 10.1016/j.cor.2018.10.012.

[12] J. Bautista, E. Fernández, and J. Pereira, "Solving an urban waste collection problem using ants heuristics," *Computers & Operations Research*, vol. 35, no. 9, pp. 3020–3033, 2008. DOI: 10.1016/j.cor.2007.01.029.

[13] R. D. S. S. Rambandara, R. Prabodanie, E. Karunarathne, and R. Rajapaksha, "Improving the efficiency of urban waste collection using optimization: A case study," *Process Integration and Optimization for Sustainability*, vol. 6, no. 3, pp. 809–818, 2022. DOI: 10.1007/s41660-022-00232-8.

[14] K. Broughton, "Waste management boosts profit margins by rerouting and replacing trash trucks," *The Wall Street Journal*, 2024.

[15] AMCS Group, *About amcs*, 2025. [Online]. Available: https://www.amcsgroup.com/.

[16] H. Han and E. Ponce-Cueto, "Waste collection vehicle routing problem: A literature review," *Promet - Traffic & Transportation*, vol. 27, no. 4, pp. 345–358, 2015. DOI: 10.7307/ptt.v27i4.1616.

[17] F. Bonomo, G. Durán, F. Larumbe, and J. Marenco, "A method for optimizing waste collection using mathematical programming: A buenos aires case study," *Waste Management & Research*, vol. 30, no. 3, pp. 311–24, 2012. DOI: 10.1177/0734242X11402870.

[18] M. J. Cortinhal, M. C. Mourão, and A. C. Nunes, "Local search heuristics for sectoring routing in a household waste collection context," *European Journal of Operational Research*, vol. 255, no. 1, pp. 68–79, 2016. DOI: 10.1016/j.ejor.2016.04.013.

[19] E. Angelelli and M. G. Speranza, "The application of a vehicle routing model to a waste collection problem: Two case studies," *Journal of the Operational Research Society*, vol. 53, no. 9, pp. 944–952, 2002. DOI: 10.1057/palgrave.jors.2601402.

[20] H. Tang and E. Miller-Hooks, "Interactive heuristic for practical vehicle routing problem with solution shape constraints," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 1964, no. 1, pp. 9–18, 2006. DOI: 10.1177/0361198106196400102.

[21] P. Matis, "Decision support system for solving the street routing problem," *Transport*, vol. 23, no. 3, pp. 230–235, 2008. DOI: 10.3846/1648-4142.2008.23.230-235.

[22] J. Sniezek and L. Bodin, "Using mixed integer programming for solving the capacitated arc routing problem with vehicle/site dependencies with an application to the routing of residential sanitation collection vehicles," *Annals of Operations Research*, vol. 144, pp. 33–58, 2006. DOI: 10.1007/s10479-006-0006-y.

[23] E.-G. Talbi, *Metaheuristics: From Design to Implementation*. John Wiley & Sons, 2009.

[24] J. Cirasella, D. S. Johnson, L. A. McGeoch, and W. Zhang, "The asymmetric traveling salesman problem: Algorithms, instance generators, and tests," in *Algorithm Engineering and Experimentation*, ser. Lecture Notes in Computer Science, vol. 2153, Berlin, Heidelberg: Springer, 2001, pp. 32–59. DOI: 10.1007/3-540-44808-X_3.

[25] R. Roberti and P. Toth, "Models and algorithms for the asymmetric traveling salesman problem: An experimental comparison," *EURO Journal on Transportation and Logistic*, vol. 1, no. 1, pp. 113–133, 2012. DOI: 10.1007/s13676-012-0010-0.

[26] T. Öncan, I. K. Altınel, and G. Laporte, "A comparative analysis of several asymmetric traveling salesman problem formulations," *Computers & Operations Research*, vol. 36, no. 3, pp. 637–654, 2009. DOI: 10.1016/j.cor.2007.11.008.

[27] J. Aransay, V. Galilea, and I. Hernández, "A tool to improve the efficiency of waste collection: Development and application to a case study," in *Proceedings of 8th Transport Research Arena (TRA 2020)*, Helsinki, Finland, 2020.

[28] R. Benitez-Bravo et al., "Optimization of municipal solid waste collection routes in a latin-american context," *Journal of the Air & Waste Management Association*, vol. 71, no. 11, pp. 1415–1427, 2021. DOI: 10.1080/10962247.2021.1957040.

[29] U. Boryczka and K. Szwarc, "The adaptation of the harmony search algorithm to the atsp with the evaluation of the influence of the pitch adjustment place on the quality of results," *Journal of Information and Telecommunication*, vol. 3, no. 1, pp. 2–18, 2019. DOI: 10.1080/24751839.2018.1503149.

[30] J. Brest and J. Žerovnik, "A heuristic for the asymmetric traveling salesman problem," in *Proceedings of the 6th Metaheuristics International Conference (MIC 2005)*, Vienna, Austria, 2005, pp. 145–150. [Online]. Available: https://www.researchgate.net/publication/228640882_A_heuristic_for_the_asymmetric_traveling_salesman_problem.

[31] B. Freisleben and P. Merz, "A genetic local search algorithm for solving symmetric and asymmetric traveling salesman problems," in *Proceedings of the IEEE International Conference on Evolutionary Computation*, Nagoya, Japan, 1996, pp. 616–621. DOI: 10.1109/ICEC.1996.542672.

[32] Y. Nagata and D. Soler, "A new genetic algorithm for the asymmetric traveling salesman problem," *Expert Systems with Applications*, vol. 39, no. 10, pp. 8947–8953, 2012. DOI: 10.1016/j.eswa.2012.02.029.

[33] N. V. Karadimas, K. Papatzelou, and V. G. Loumos, "Genetic algorithms for municipal solid waste collection and routing optimization," in *Artificial Intelligence and Innovations 2007: From Theory to Applications*, C. Boukis, A. Pnevmatikakis, and L. Polymenakos, Eds., ser. IFIP The International Federation for Information Processing, vol. 247, Boston, MA: Springer, 2007, pp. 223–231. DOI: 10.1007/978-0-387-74161-1_24.

[34] K. Mekamcha, M. Souier, H. N. Bessenouci, and M. Bennekrouf, "Two metaheuristics approaches for solving the traveling salesman problem: An algerian waste collection case," *Operational Research*, vol. 21, pp. 1641–1661, 2021. DOI: 10.1007/s12351-019-00529-6.

[35] E. K. Burke, P. I. Cowling, and R. Keuthen, "Effective local and guided variable neighbourhood search methods for the asymmetric travelling salesman problem," in *Applications of Evolutionary Computing*, E. Boers, Ed., ser. Lecture Notes in Computer Science, vol. 2037, Berlin, Heidelberg: Springer, 2001, pp. 203–212. DOI: 10.1007/3-540-45365-2_21.

[36] P. Viotti, A. Polettini, R. Pomi, and C. Innocenti, "Genetic algorithms as a promising tool for optimisation of the msw collection routes," *Waste Management & Research*, vol. 21, no. 4, pp. 292–298, 2003. DOI: 10.1177/0734242X0302100402.

[37] N. V. Karadimas, K. Papatzelou, and V. G. Loumos, "Optimal solid waste collection routes identified by the ant colony system algorithm," *Waste Management & Research*, vol. 25, no. 2, pp. 139–147, 2007. DOI: 10.1177/0734242X07071312.

[38] M. de Berg, O. Cheong, M. Kreveld, and M. Overmars, *Computational Geometry: Algorithms and Applications*, 3rd ed. Berlin, Heidelberg: Springer, 2008. DOI: 10.1007/978-3-540-77974-2.

A. Pseudocode of K-Means Neighbourhood Generation

Neighbourhoods used for the neighbourhood fragmentation metric are generated using the k-means clustering procedure in Algorithm 1. To avoid assigning opposite kerbsides of the same street segment to different neighbourhoods, clustering is performed on aggregated road segments. Each road segment $r \in R$ is represented by a centroid with coordinates (x_r, y_r) computed from its associated link clusters. The number of neighbourhoods is set using a target size, $k = \lceil |C| / \mathit{target\_size} \rceil$. The algorithm alternates between assigning each road centroid to its nearest centre and updating centres as the mean of assigned centroids, until convergence (i.e., maximum 100 iterations or total centre shift below $\eta = 10^{-6}$). Empty clusters are reinitialised using random road centroids. The final neighbourhood labels are then mapped back to all link clusters.

Algorithm 1 K-Means Neighbourhood Generation from Road Segment Midpoints

procedure KMeansNeighbourhoods(R, k, s)
  compute road centroids (x_r, y_r) for all $r \in R$
  initialise k centres by sampling road centroids from $\{(x_r, y_r) : r \in R\}$ using seed s
  $\delta \leftarrow \infty$
  while $\delta > \eta$ **and** iterations < 100 **do**
    for each road $r \in R$ **do**
      compute distance from (x_r, y_r) to each centre
      assign r to nearest centre
    end for
    $\delta \leftarrow 0$
    for each neighbourhood $h \in \{1, \ldots, k\}$ **do**
      if h has at least one assigned road **then**
        update centre as mean of assigned road centroids
      else
        reinitialise centre as a random road centroid
      end if
      update δ with centre shift for h
    end for
  end while
  sort centres lexicographically and relabel them $1, \ldots, k$ for consistency
  map road-level labels back to all link clusters
  return neighbourhood assignment $\nu(c)$ for all $c \in C$
end procedure

B. Pseudocode of the Optimisation Framework

The detailed procedure for the optimisation framework using SA is given in Algorithm 2.

Algorithm 2 Optimisation Framework Using a Simulated Annealing Procedure

procedure SA($T_0, \alpha, T_{final}, N_T, \hat{w}_{\mathrm{TT}}, \hat{w}_{\mathrm{NF}}, \hat{w}_{\mathrm{IC}}, \hat{w}_{\mathrm{LJ}}, \hat{w}_{\mathrm{RS}}, P(\mathrm{op}), n, \Delta_{RS}, k$)

▷ **Step 0:** Preprocessing of data

$C \leftarrow \{\mathrm{LC}\},\ Z \leftarrow \{\text{terminals}\},\ V \leftarrow C \cup Z,\ A \leftarrow \{\text{arcs}\},\ R \leftarrow \{\text{revisit pairs}\}$

$tt_{ij} \leftarrow$ travel time on each arc $(i,j) \in A$, $\quad s_i \leftarrow$ service time at each node $i \in V$

$dd_{cd} \leftarrow$ directed Euclidean distance for all $c, d \in C, c \neq d$

$d_{cd} \leftarrow$ undirected Euclidean distance for all $c, d \in C, c \neq d$

$H \leftarrow$ neighbourhoods, $\nu(c) \leftarrow$ neighbourhood labels $\forall c \in C$ (obtained from $\textsc{KMeans}(P^0_{\mathrm{LC}}, k, 42)$)

$(x_c^{\mathrm{start}}, y_c^{\mathrm{start}}), (x_c^{\mathrm{end}}, y_c^{\mathrm{end}}) \leftarrow$ endpoint coordinates for each $c \in C$

▷ **Step 1:** Obtain initial solution from travel time minimisation model

$P^0_{\mathrm{nodes}} \leftarrow$ ordered node sequence of travel time optimal route

$P^0_{\mathrm{LC}} \leftarrow$ ordered link cluster sequence obtained by removing terminals Z from P^0_{nodes}

$TT_0, NF_0, IC_0, LJ_0, RS_0 \leftarrow$ initial metrics using $P^0_{\mathrm{LC}}, Z, tt_{ij}, s_i, \nu(c), (x_c^{\mathrm{start}}, y_c^{\mathrm{start}}), (x_c^{\mathrm{end}}, y_c^{\mathrm{end}}), dd_{cd}, R, \Delta_{RS}$

▷ **Step 2:** Initialisation

$w_{\mathrm{TT}} \leftarrow \hat{w}_{\mathrm{TT}}/TT_0,\ w_{\mathrm{NF}} \leftarrow \hat{w}_{\mathrm{NF}}/NF_0,\ w_{\mathrm{IC}} \leftarrow \hat{w}_{\mathrm{IC}}/IC_0,\ w_{\mathrm{LJ}} \leftarrow \hat{w}_{\mathrm{LJ}}/LJ_0,\ w_{\mathrm{RS}} \leftarrow \hat{w}_{\mathrm{RS}}/RS_0$ ▷ Normalisation

$F(TT, NF, IC, LJ, RS) \leftarrow w_{\mathrm{TT}}TT + w_{\mathrm{NF}}NF + w_{\mathrm{IC}}IC + w_{\mathrm{LJ}}LJ + w_{\mathrm{RS}}RS$ ▷ Define objective function

$P_{\mathrm{LC}} \leftarrow P^0_{\mathrm{LC}}, \quad F \leftarrow F(TT_0, NF_0, IC_0, LJ_0, RS_0)$ ▷ Current route and objective

$P^\star_{\mathrm{LC}} \leftarrow P_{\mathrm{LC}}, \quad F^\star \leftarrow F$ ▷ Best route and objective

$TT^*, NF^*, IC^*, LJ^*, RS^* \leftarrow TT, NF, IC, LJ, RS \leftarrow TT_0, NF_0, IC_0, LJ_0, RS_0$ ▷ Best and current metrics

$T \leftarrow T_0$ ▷ Current temperature

▷ **Step 3:** SA loop with N_T iterations per temperature T

while $T > T_{\mathrm{final}}$ **do**

  $N \leftarrow 0$

  while $N < N_T$ **do**

    $N \leftarrow N + 1$

    choose operator op $\sim P(\mathrm{op})$ where op $\in \{\textsc{Relocate}, \textsc{Swap}, \textsc{OrOpt}\}$

    $P^{\mathrm{cand}}_{\mathrm{LC}}, TT^{\mathrm{cand}}, NF^{\mathrm{cand}}, IC^{\mathrm{cand}}, LJ^{\mathrm{cand}}, RS^{\mathrm{cand}}, F_{\mathrm{cand}} \leftarrow$ **None**

    for $i \in n$ **do** ▷ Sample n candidates and choose best candidate

      if *op* = Relocate **then**

        $P'_{\mathrm{LC}} \leftarrow \textsc{Relocate}(P_{\mathrm{LC}}, 42)$

      else if *op* = Swap **then**

        $P'_{\mathrm{LC}} \leftarrow \textsc{Swap}(P_{\mathrm{LC}}, 42)$

      else

        $P'_{\mathrm{LC}} \leftarrow \textsc{OrOpt}(P_{\mathrm{LC}}, 3, 42)$

      end if

      $TT', NF', IC', LJ', RS' \leftarrow$ metrics using $(Z, tt_{ij}, s_i, \nu(c), (x_c^{\mathrm{start}}, y_c^{\mathrm{start}}), (x_c^{\mathrm{end}}, y_c^{\mathrm{end}}), dd_{cd}, R, \Delta_{RS})$

      $F' \leftarrow F(TT', NF', IC', LJ', RS')$

      if $F' < F_{\mathrm{cand}}$ **then**

        $P^{\mathrm{cand}}_{\mathrm{LC}}, TT^{\mathrm{cand}}, NF^{\mathrm{cand}}, IC^{\mathrm{cand}}, LJ^{\mathrm{cand}}, RS^{\mathrm{cand}}, F_{\mathrm{cand}} \leftarrow P'_{\mathrm{LC}}, TT', NF', IC', LJ', RS', F'$

      end if

    end for

    $\Delta_F \leftarrow F_{\mathrm{cand}} - F$ ▷ Compute objective change w.r.t. current objective

    if $\Delta_F \leq 0$ **then** ▷ Acceptance test

      accept $\leftarrow$ true

    else

      $p \leftarrow \exp(-\Delta_F/T)$, draw $r \sim U[0,1]$

      accept $\leftarrow (r < p)$

    end if

    if *accept* **then** ▷ Update current solution

      $P_{\mathrm{LC}} \leftarrow P^{\mathrm{cand}}_{\mathrm{LC}}, \quad F_{\mathrm{cur}} \leftarrow F_{\mathrm{cand}}, \quad TT, NF, IC, LJ, RS \leftarrow TT^{\mathrm{cand}}, NF^{\mathrm{cand}}, IC^{\mathrm{cand}}, LJ^{\mathrm{cand}}, RS^{\mathrm{cand}}$

      if $F < F^\star$ **then** ▷ Update best-so-far solution

        $P^\star_{\mathrm{LC}} \leftarrow P_{\mathrm{LC}}, \quad F^\star \leftarrow F, \quad TT^*, NF^*, IC^*, LJ^*, RS^* \leftarrow TT, NF, IC, LJ, RS$

      end if

    end if

  end while

  $T \leftarrow \alpha T$ ▷ Cooling update after N_T iterations

end while

return $P^\star_{\mathrm{LC}}, F^*, TT^*, NF^*, IC^*, LJ^*, RS^*$ ▷ Return best-so-far solution

end procedure

B

Additional Data Description

B.1. Additional Network Representation

Figure B.1–Figure B.11 provide a visual representation of all case studies, showing both the road and abstract network. These plots illustrate how orders, terminals, and link clusters are distributed across the regions in New Zealand (set 1) and United States (set 2).

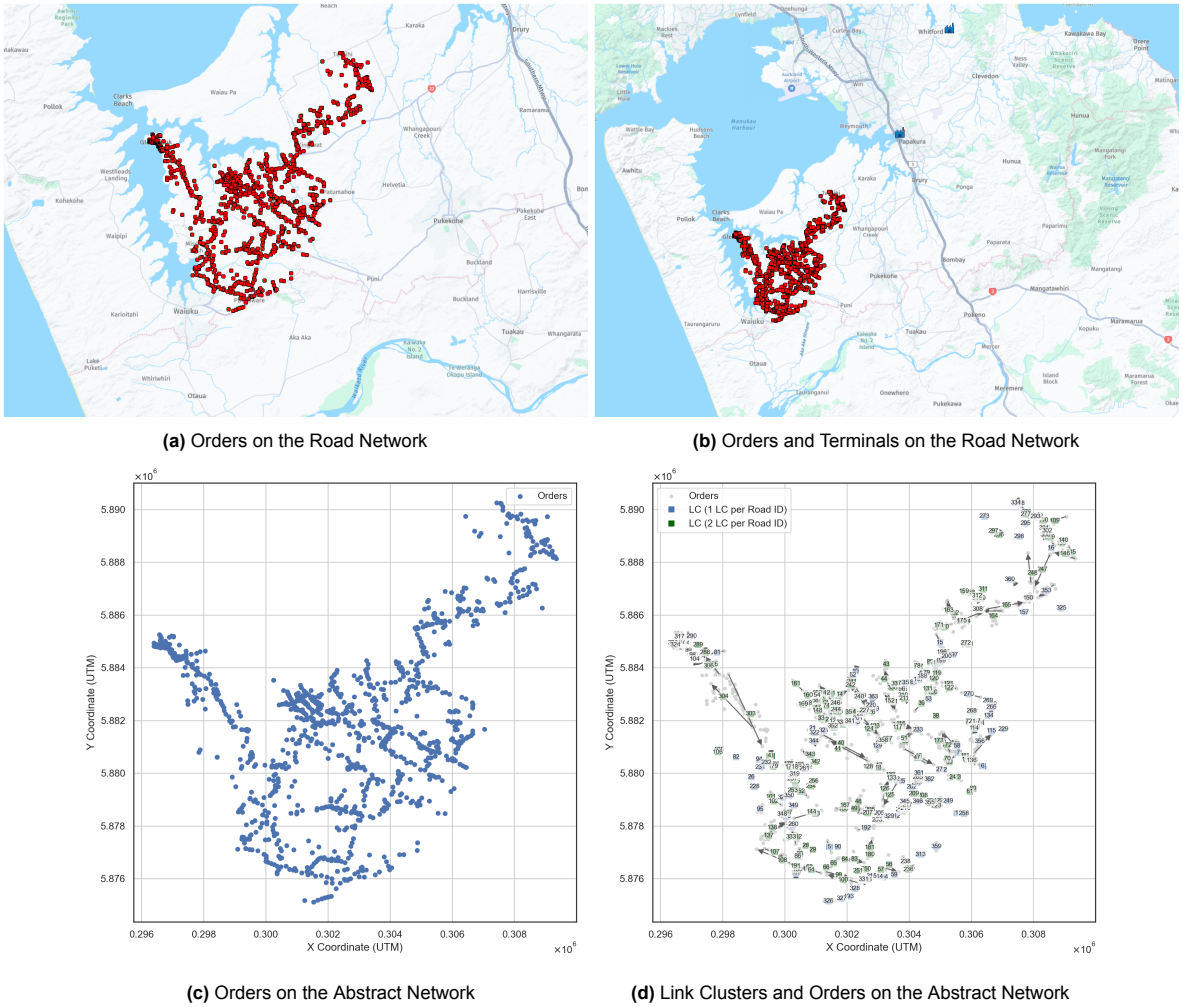

(a) Orders on the Road Network

(b) Orders and Terminals on the Road Network

(c) Orders on the Abstract Network

(d) Link Clusters and Orders on the Abstract Network

Figure B.1: Dataset Representations on the Road Network and Abstract Network for Route 2 of Set 1 (New Zealand)

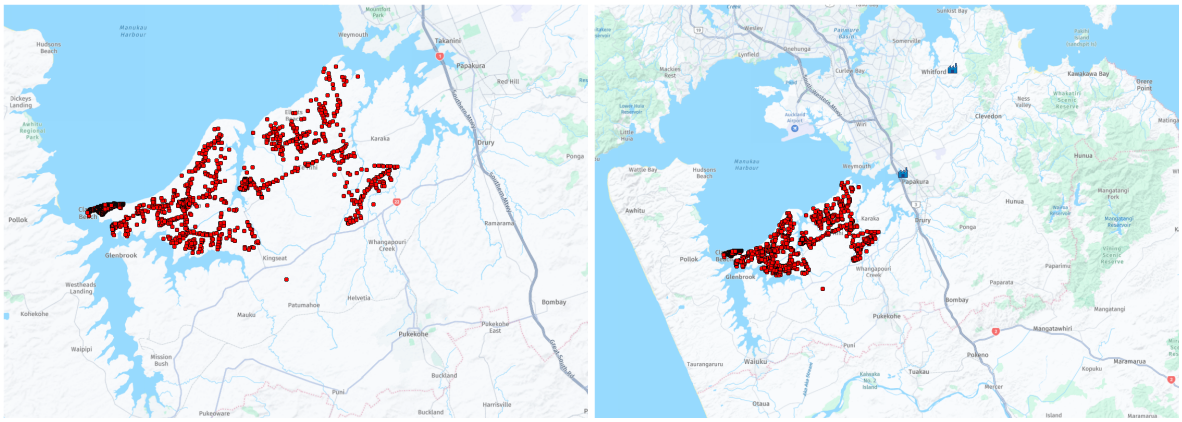

(a) Orders on the Road Network

(b) Orders and Terminals on the Road Network

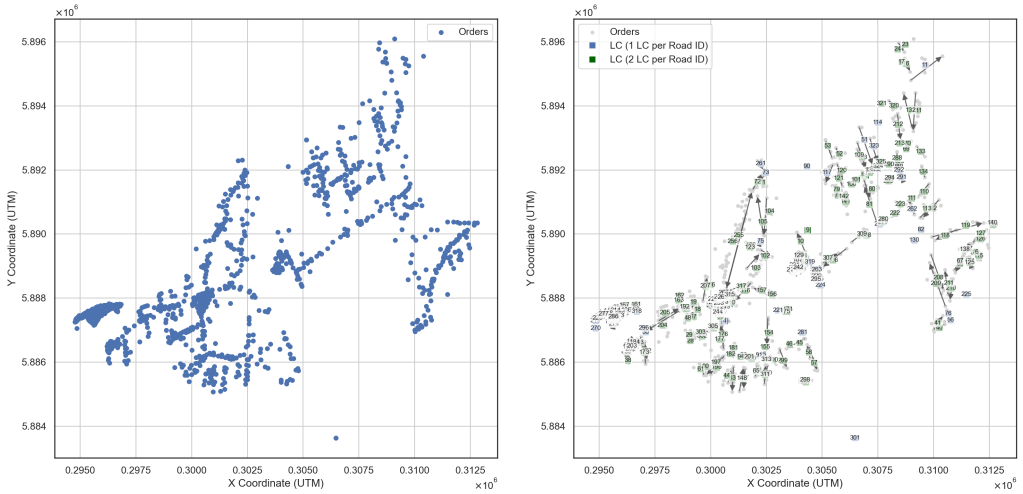

(c) Orders on the Abstract Network

(d) Link Clusters and Orders on the Abstract Network

Figure B.2: Dataset Representations on the Road Network and Abstract Network for Route 3 of Set 1 (New Zealand)

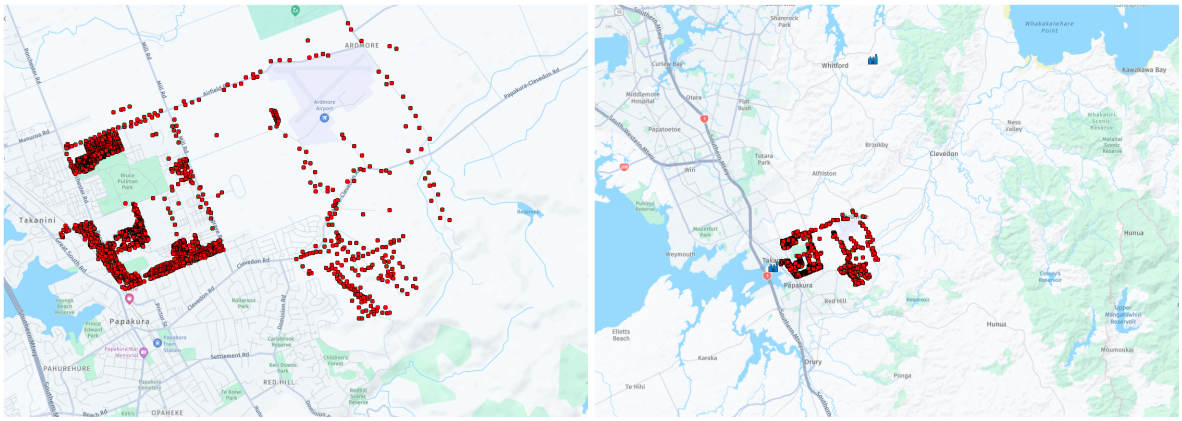

(a) Orders on the Road Network

(b) Orders and Terminals on the Road Network

Figure B.3: Dataset Representations on the Road Network and Abstract Network for Route 4 of Set 1 (New Zealand)

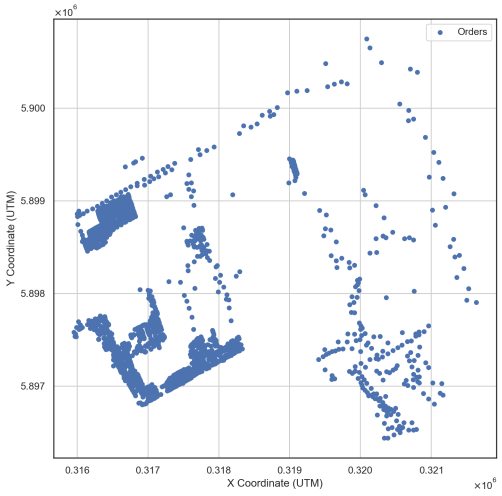

(c) Orders on the Abstract Network

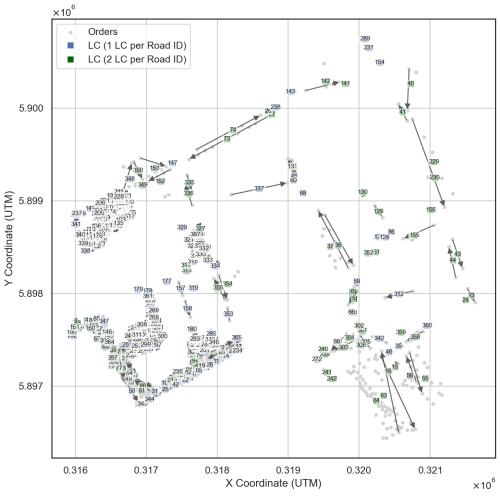

(d) Link Clusters and Orders on the Abstract Network

Figure B.3: Dataset Representations on the Road Network and Abstract Network for Route 4 of Set 1 (New Zealand) *(Continued)*

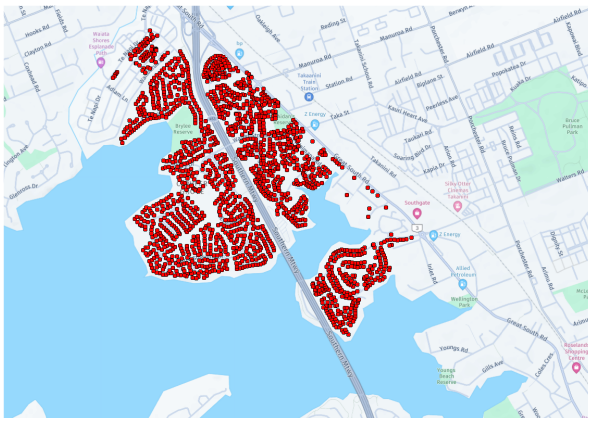

(a) Orders on the Road Network

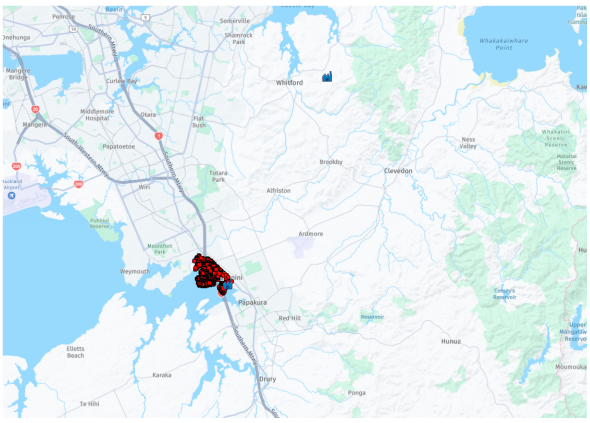

(b) Orders and Terminals on the Road Network

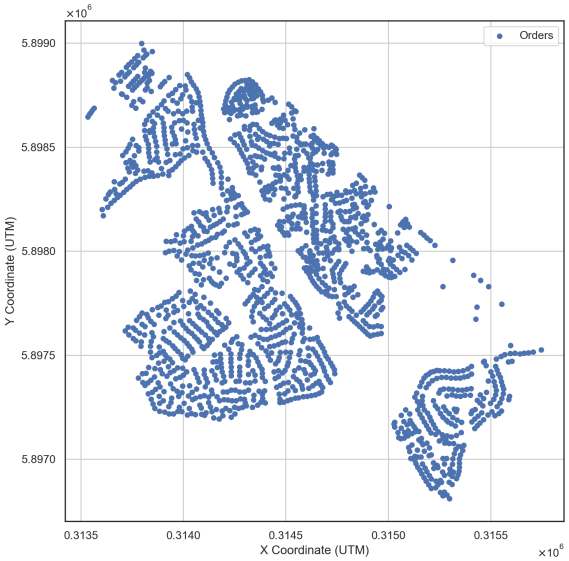

(c) Orders on the Abstract Network

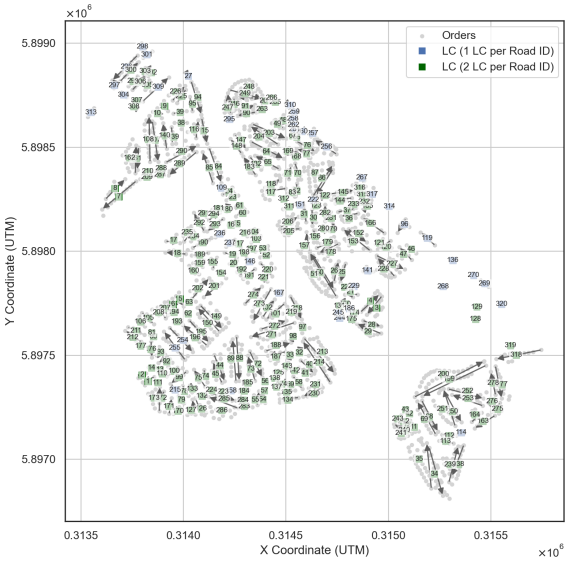

(d) Link Clusters and Orders on the Abstract Network

Figure B.4: Dataset Representations on the Road Network and Abstract Network for Route 5 of Set 1 (New Zealand)

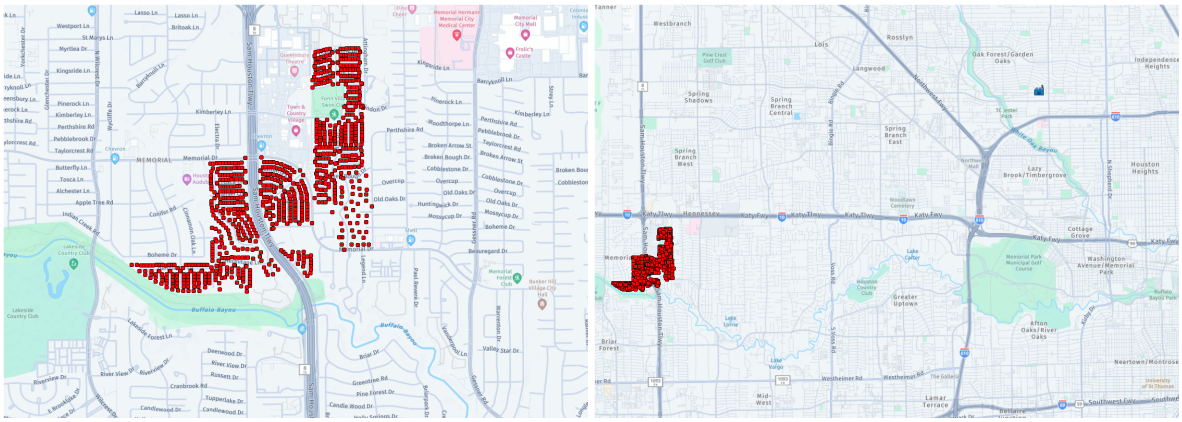

(a) Orders on the Road Network

(b) Orders and Terminals on the Road Network

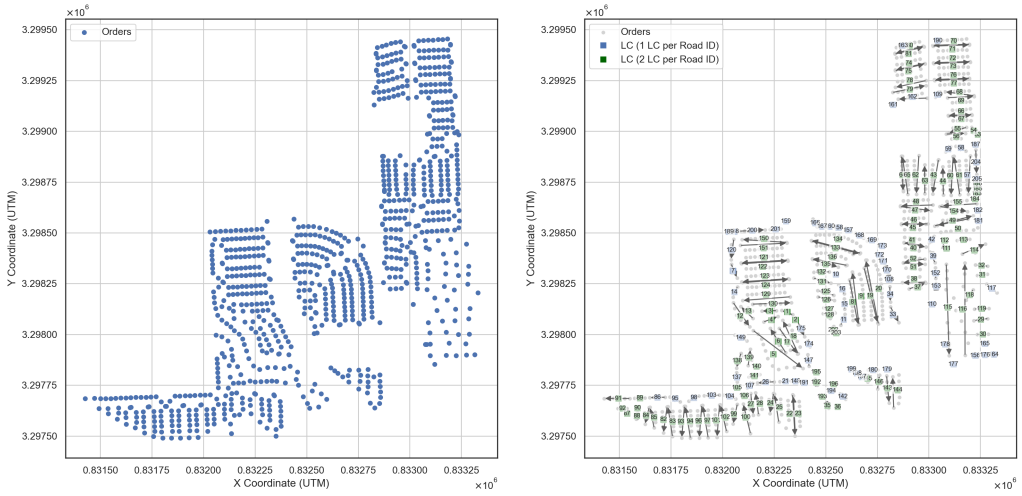

(c) Orders on the Abstract Network

(d) Link Clusters and Orders on the Abstract Network

Figure B.5: Dataset Representations on the Road Network and Abstract Network for Route 1 of Set 2 (United States)

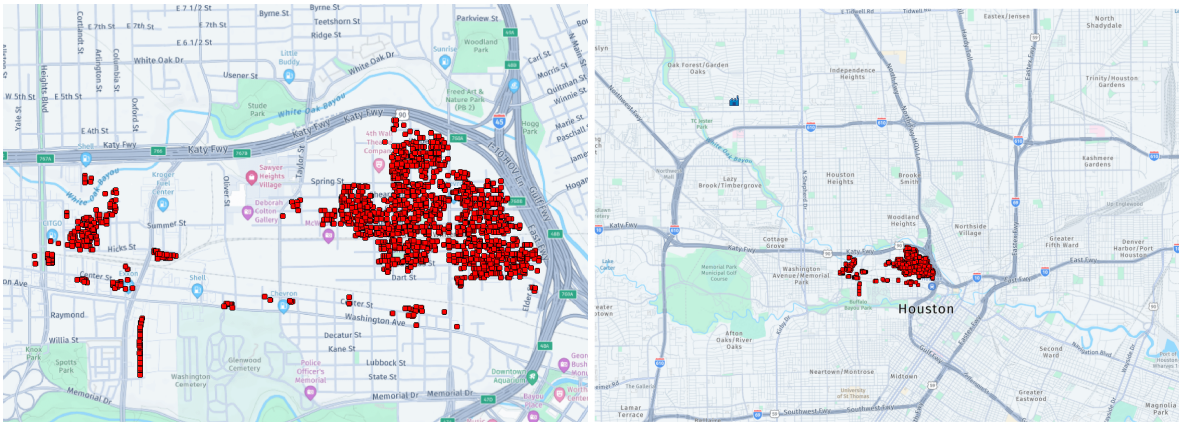

(a) Orders on the Road Network

(b) Orders and Terminals on the Road Network

Figure B.6: Dataset Representations on the Road Network and Abstract Network for Route 2 of Set 2 (United States)

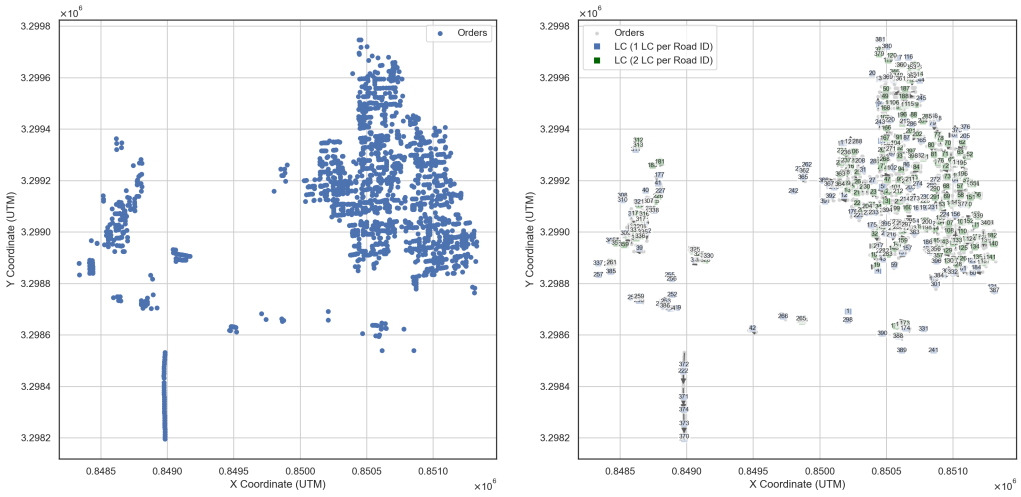

(c) Orders on the Abstract Network

(d) Link Clusters and Orders on the Abstract Network

Figure B.6: Dataset Representations on the Road Network and Abstract Network for Route 2 of Set 2 (United States) *(Continued)*

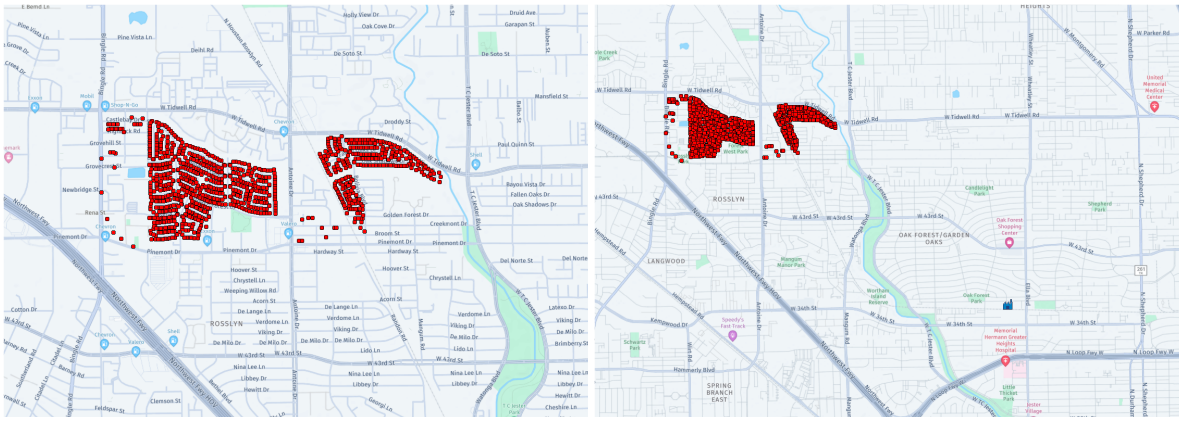

(a) Orders on the Road Network

(b) Orders and Terminals on the Road Network

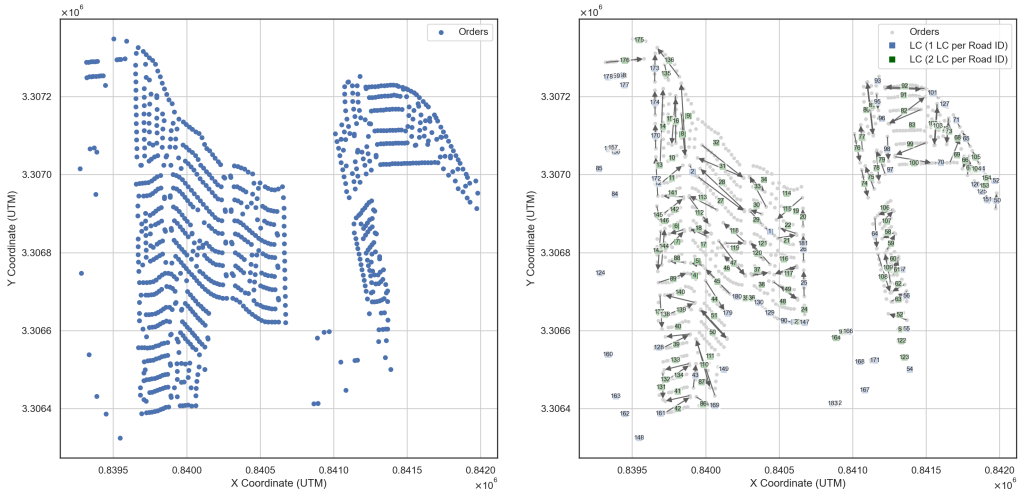

(c) Orders on the Abstract Network

(d) Link Clusters and Orders on the Abstract Network

Figure B.7: Dataset Representations on the Road Network and Abstract Network for Route 3 of Set 2 (United States)

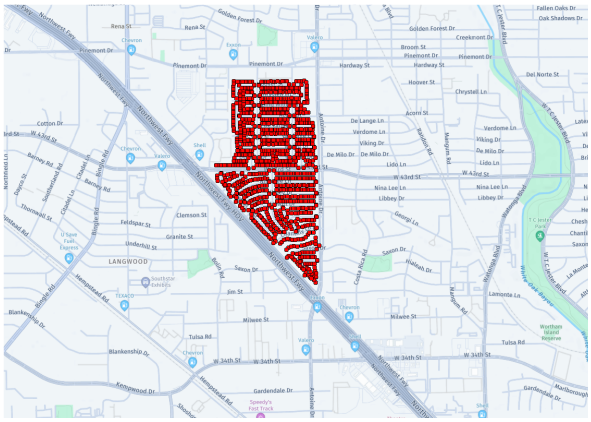

(a) Orders on the Road Network

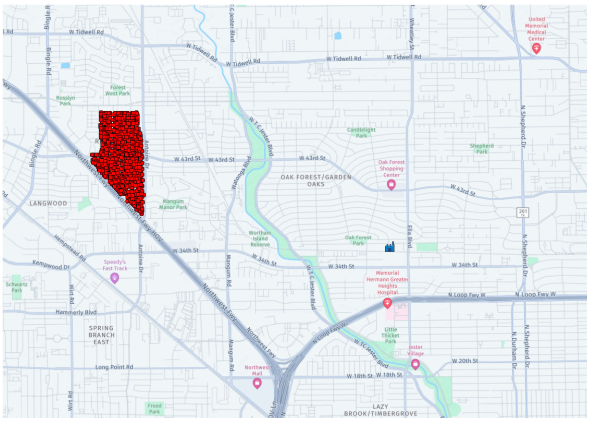

(b) Orders and Terminals on the Road Network

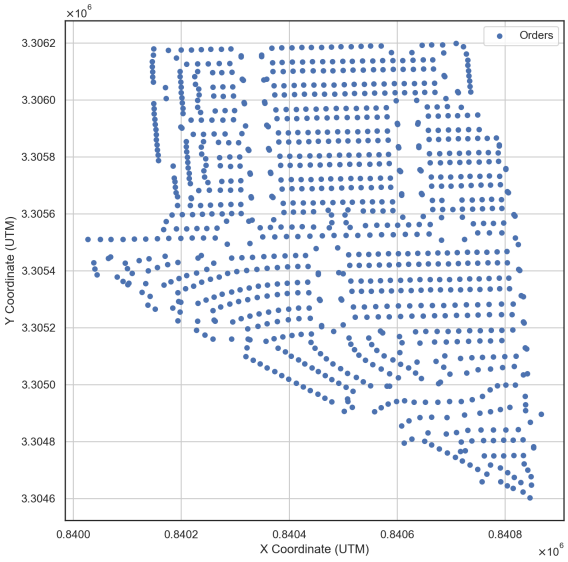

(c) Orders on the Abstract Network

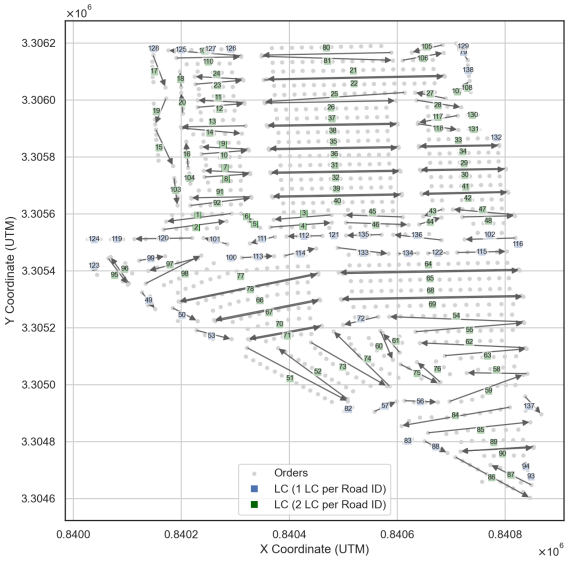

(d) Link Clusters and Orders on the Abstract Network

Figure B.8: Dataset Representations on the Road Network and Abstract Network for Route 4 of Set 2 (United States)

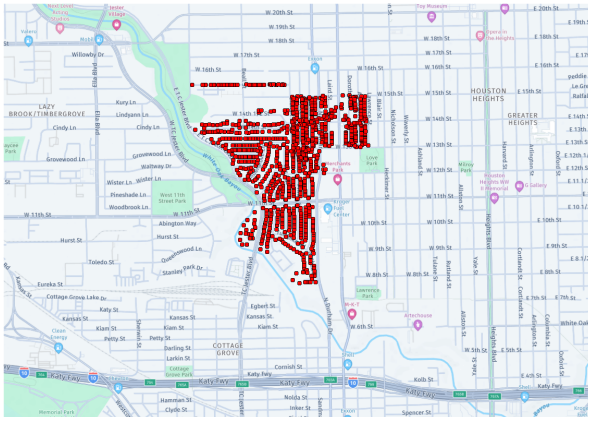

(a) Orders on the Road Network

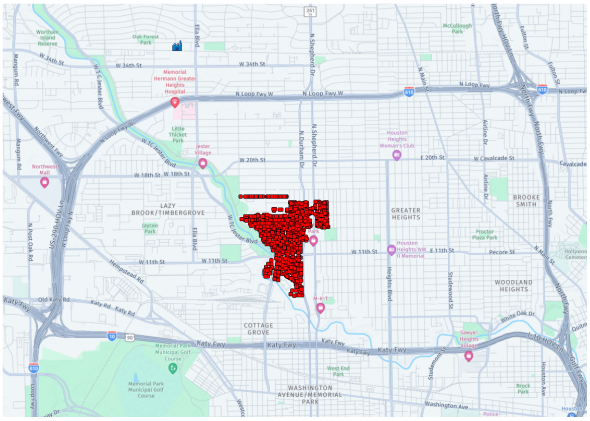

(b) Orders and Terminals on the Road Network

Figure B.9: Dataset Representations on the Road Network and Abstract Network for Route 5 of Set 2 (United States)

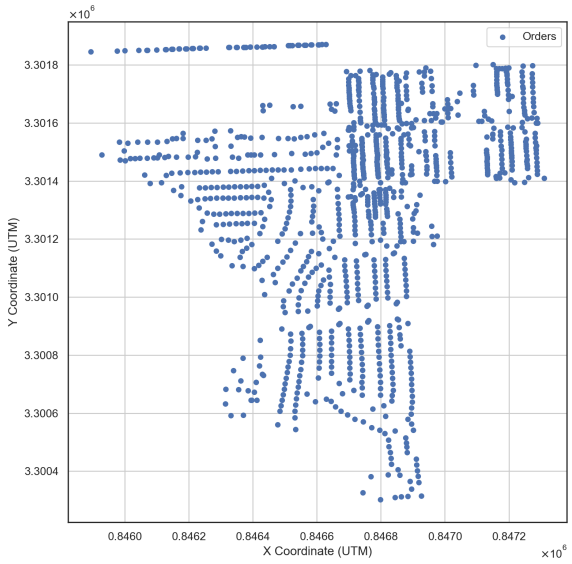

(c) Orders on the Abstract Network

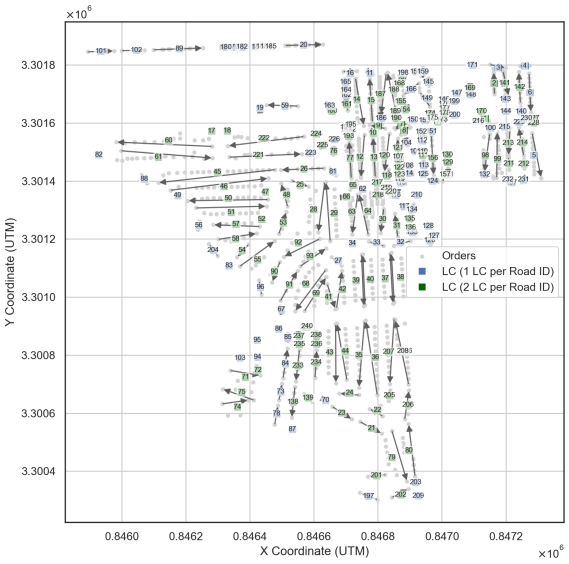

(d) Link Clusters and Orders on the Abstract Network

Figure B.9: Dataset Representations on the Road Network and Abstract Network for Route 5 of Set 2 (United States) *(Continued)*

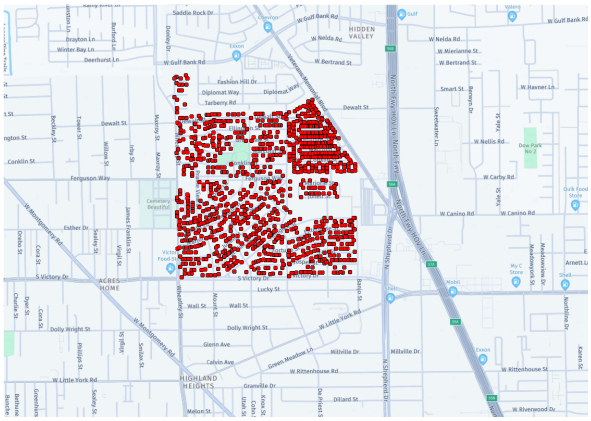

(a) Orders on the Road Network

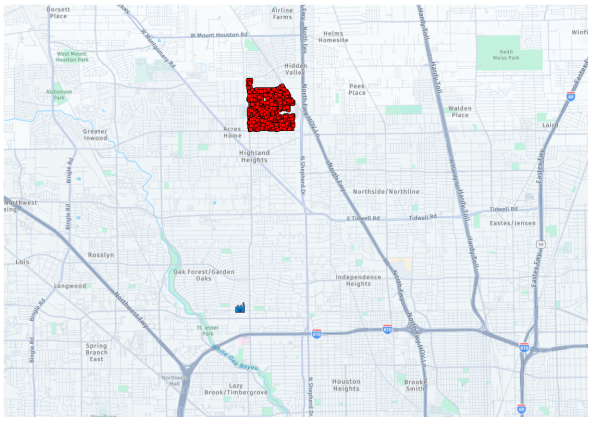

(b) Orders and Terminals on the Road Network

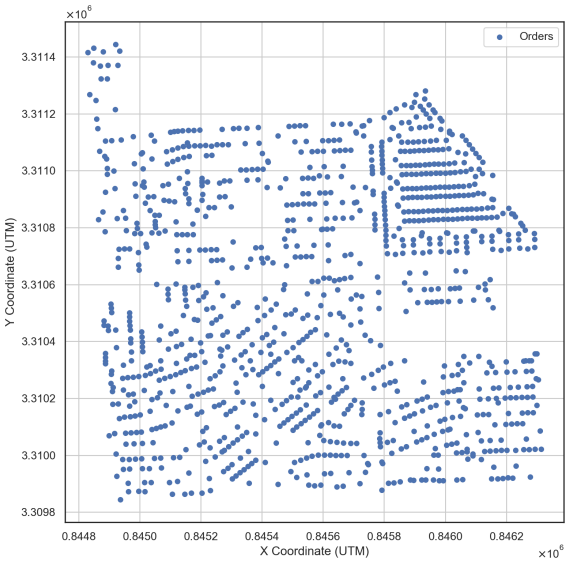

(c) Orders on the Abstract Network

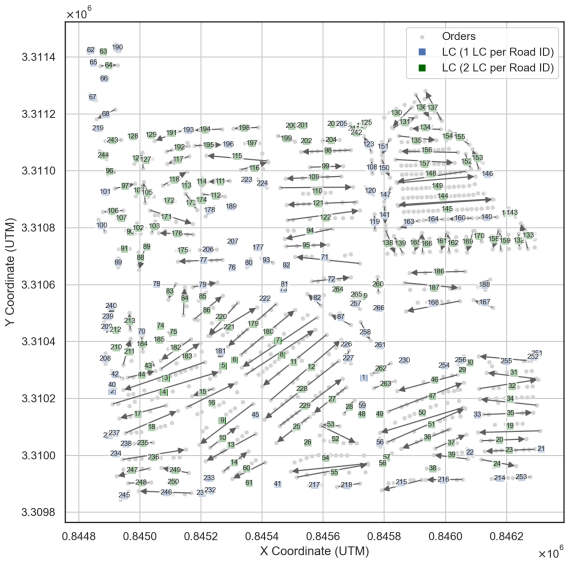

(d) Link Clusters and Orders on the Abstract Network

Figure B.10: Dataset Representations on the Road Network and Abstract Network for Route 6 of Set 2 (United States)

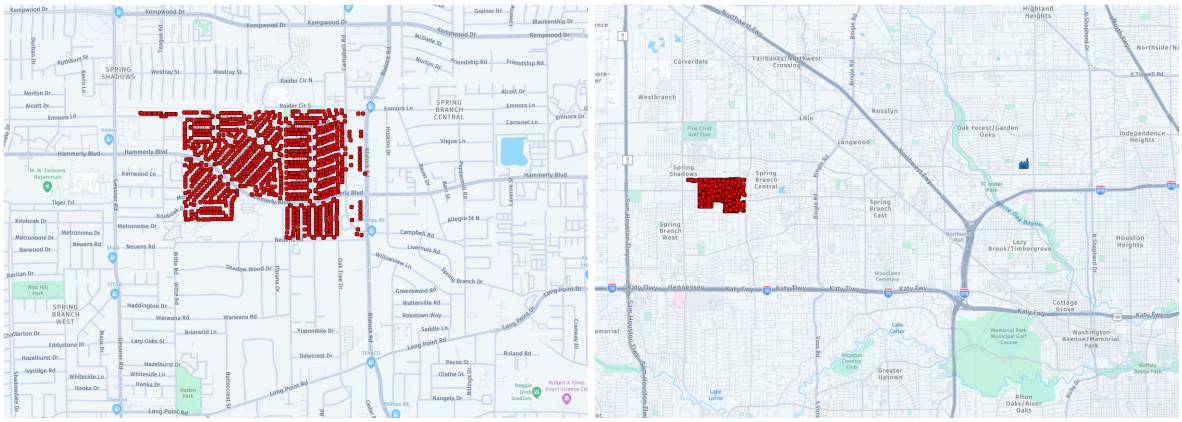

(a) Orders on the Road Network

(b) Orders and Terminals on the Road Network

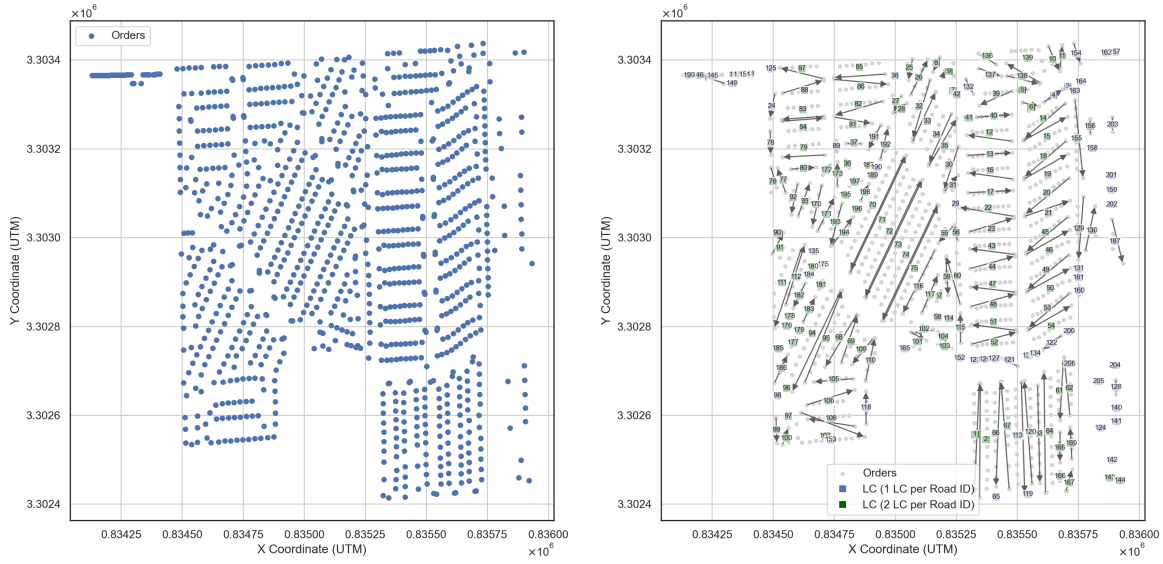

(c) Orders on the Abstract Network

(d) Link Clusters and Orders on the Abstract Network

Figure B.11: Dataset Representations on the Road Network and Abstract Network for Route 7 of Set 2 (United States)

B.2. Additional Descriptive Statistics

This section presents the statistics for all case studies in addition to the summaries in Section 3.2.

B.2.1. Network Size and Composition

The descriptive statistics on network composition and size are shown for each individual case study route in Table B.1, including the number of orders and link clusters, convex hull area, and kerbside composition. These values complement the summary table in Subsection 3.2.1.

Table B.1: Network Size and Composition Statistics of the Case Studies

Route ID	Area (m^2)	Total Orders	Total LCs	Average Orders per LC	Maximum Orders per LC	Ratio Single-to-Dual
Set 1 (New Zealand)						
1.1	653 762	903	194	4.65	22	0.94
1.2	113 204 551	1172	363	3.23	28	1.00
1.3	117 982 848	1526	325	4.70	37	0.66
1.4	15 815 393	1762	367	4.80	41	1.99

Continued on next page

Table B.1: Network Size and Composition Statistics of the Case Studies *(Continued)*

Route ID	Area (m^2)	Total Orders	Total LCs	Average Orders per LC	Maximum Orders per LC	Ratio Single-to-Dual
1.5	2 655 361	1815	320	5.67	28	0.32
Set 2 (United States)						
2.1	2 006 389	933	205	4.55	17	1.11
2.2	3 114 554	1618	398	4.07	24	1.32
2.3	2 107 399	967	183	5.28	20	1.05
2.4	872 531	969	138	7.02	19	0.82
2.5	1 478 783	1074	240	4.48	21	1.53
2.6	1 934 322	1127	266	4.24	21	1.09
2.7	1 452 361	1070	206	5.19	18	0.78

B.2.2. Spatial Distribution

The spatial descriptive statistics of all case studies are shown in Table B.2, supporting the summaries reported Subsection 3.2.2. These metrics fall into three groups: (i) centrality and compactness (standard distance, geometric offset), which describe the location and spread of the service area; (ii) density-based measures (global density), which capture the concentration of link clusters; and (iii) distance-based metrics (NND-ratio), which assess how link clusters are spaced relative to randomness.

Standard distances vary from roughly 370 metres in the most compact networks to over 5 kilometres in the most extensive ones. Most routes have small geometric offsets (0.05–0.12), suggesting balanced and symmetrical spatial distributions, while a few cases (e.g., routes 1.3, 1.4, 2.2) show larger values above 0.20, indicating mild skewness caused by elongated or uneven layouts. Global densities ($\lambda = n/A$) also vary substantially. Moreover, the clustering is particularly strong for routes 1.2–1.4, 2.2, and 2.3 ($R < 1$, $z < -10$), while routes 2.1, 2.5, and 2.7 show almost random spacing ($R{\approx}1$, $|z| < 1$). A few routes, such as 2.4, 2.6, display mild spatial dispersion ($R > 1$, $z > 3$).

The nearest-neighbour statistic of Clark and Evans (1954) describes how link cluster midpoints are spaced relative to a random point pattern. It compares the observed mean nearest-neighbour distance $\bar{D}$ with the theoretical mean under complete spatial randomness (CSR), $\bar{D}_E = 1/(2\sqrt{\lambda})$. The ratio $R = \bar{D}/\bar{D}_E$ acts as a clustering index, where $R < 1$ indicates clustering, $R \approx 1$ randomness, $R > 1$ dispersion. It significance is assessed using the analytical z-score $z = (\bar{D} - \bar{D}_E)/SE$, where $SE = 0.26136/\sqrt{n^2/A}$ reflects the expected variability of $\bar{D}_E$ under CSR. The constant 0.26136 derives from the theoretical variance of a Poisson process. Following standard thresholds, $|z| > 1.96$ indicates a significant deviation from CSR. Although this analytical approximation may slightly underestimate variability in irregular areas, it is computationally efficient and sufficiently accurate for this study, as the service areas are well defined.

Table B.2: Spatial Descriptive Statistics of the Case Studies

Route ID	Standard Distance (m)	Geometric Offset	Global Density λ (m^{-2})	NND-Ratio R	Z-Score of R
Set 1 (New Zealand)					
1.1	376.47	0.116	0.000297	1.15	4.07
1.2	4664.19	0.052	0.000003	0.67	−12.10
1.3	5487.70	0.236	0.000003	0.59	−14.13
1.4	1555.27	0.243	0.000023	0.64	−13.28
1.5	690.26	0.057	0.000121	0.97	−0.95
Set 2 (United States)					
2.1	693.53	0.051	0.000102	1.00	0.03
2.2	859.12	0.304	0.000128	0.70	−11.60

Continued on next page

Table B.2: Spatial Descriptive Statistics the Case Studies *(Continued)*

Route ID	Standard Distance (m)	Geometric Offset	Global Density λ (m^{-2})	NND-Ratio R	Z-Score of R
2.3	809.52	0.136	0.000087	0.90	−2.72
2.4	463.84	0.056	0.000158	1.16	3.62
2.5	468.39	0.189	0.000162	1.03	0.85
2.6	595.85	0.035	0.000138	1.12	3.87
2.7	543.07	0.011	0.000142	1.04	1.01

Centrality and Compactness

Each dataset was analysed using four measures, namely the mean centre, geometric median centre, standard distance, and standard deviational ellipse. The mean centre gives the average location of all link clusters, but is sensitive to outliers. The geometric median provides a more robust alternative by minimising the total Euclidean distance to all clusters, making it less sensitive to extreme points.

Spatial spread around the centre is captured by the standard distance, the spatial equivalent of standard deviation. Small values indicate compact routes, whereas large values indicate dispersed service areas. Together with the mean and geometric median, it is used to compute the geometric offset between the mean and geometric median relative to overall dispersion. The standard deviational ellipse extends this by showing both spread and orientation. Circular ellipses represent compact networks, whereas elongated ones reveal directional bias, often caused by linear street layouts.

Figure B.12 and Figure B.13 show the link cluster midpoints, mean centre (black), geometric median (red), and one-standard-deviation ellipse for all case study routes. The mean and geometric median almost always coincide, indicating mostly symmetrical distributions, with only a few routes (e.g., 1.3, 1.4, 2.2) showing imbalance. The ellipses reveal compact, circular routes like 1.1 and 2.5, while routes such as 1.3 and 1.4 are more elongated. Overall, the figures demonstrate a mix of spatial distributions while preserving comparable network structures across regions.

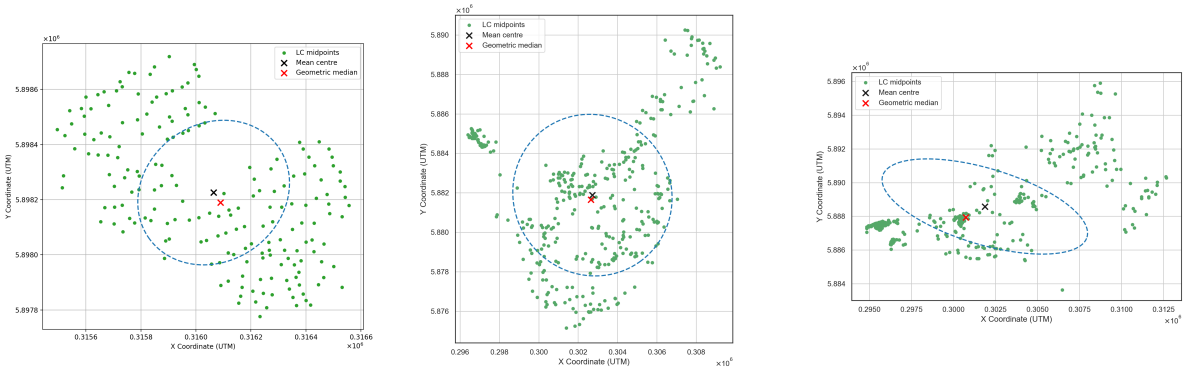

(a) Route 1.1 **(b)** Route 1.2 **(c)** Route 1.3

Figure B.12: Centrality and Compactness Measures of Link Cluster Midpoints for Routes of Set 1 (New Zealand)

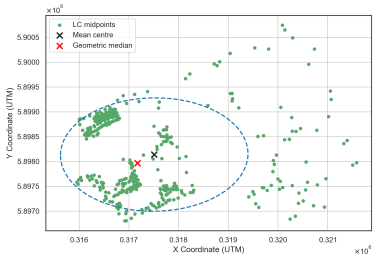

(d) Route 1.4

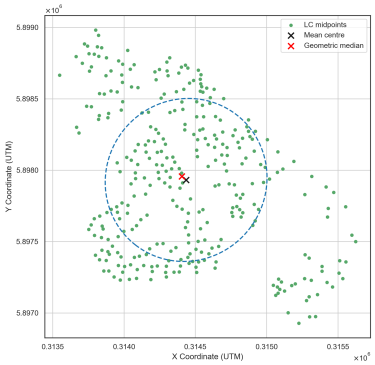

(e) Route 1.5

Figure B.12: Centrality and Compactness Measures of Link Cluster Midpoints for Routes of Set 1 (New Zealand) *(Continued)*

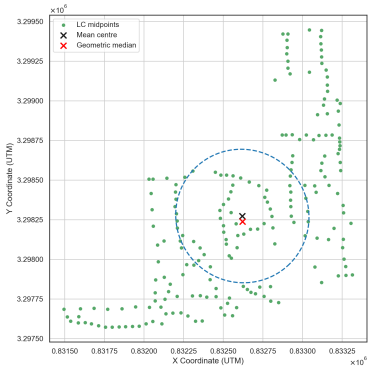

(a) Route 2.1

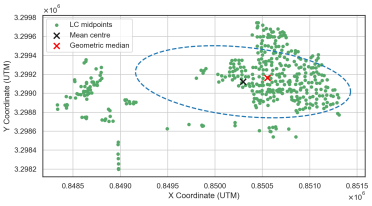

(b) Route 2.2

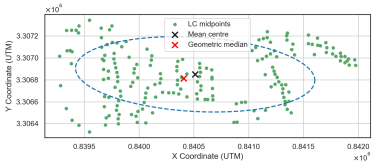

(c) Route 2.3

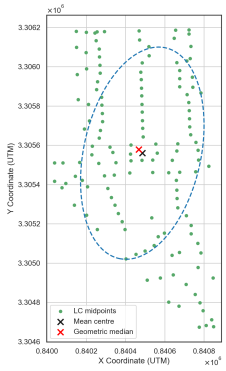

(d) Route 2.4

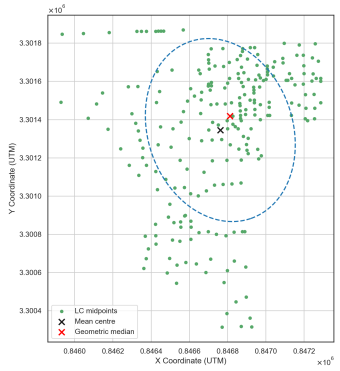

(e) Route 2.5

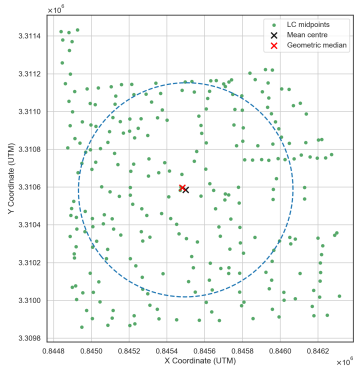

(f) Route 2.6

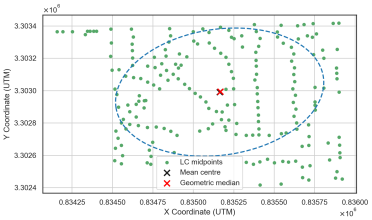

(g) Route 2.7

Figure B.13: Centrality and Compactness Measures of Link Cluster Midpoints for Routes of Set 2 (United States)

Local Density Distribution

The spatial density was analysed using the global density and local density surfaces obtained from kernel density estimation (KDE). The global density (Table B.2) reflects the intensity of link clusters within each route's convex hull area, while the KDE reveals detailed patterns such as local peaks.

A two-dimensional Gaussian KDE (Scott, 1992) was computed on a 120×120 grid covering the bounding box of all link clusters with 5% padding. The bandwidth was selected automatically using Scott's rule of thumb, $h = n^{-1/6}\bar{\sigma}$, where $\bar{\sigma}$ is the mean standard deviation. This ensures consistent smoothing across routes of different spatial scales. The resulting surfaces were scaled to express point intensity (link clusters per m^2), allowing direct comparison with global densities.

The estimated local density surfaces for all routes are shown in Figure B.14 and Figure B.15. Darker areas indicate higher concentrations of link clusters. Across the datasets, spatial heterogeneity is visible. Several routes (e.g., 1.1, 1.4, 2.2) contain strong local maxima, indicating concentrated service zones. Other routes (e.g., 1.3, 2.3, 2.4) display elongated high-density regions. In contrast, route 2.1 is more uniformly distributed, reflecting a homogeneous street layout. This is consistent with the compactness measures observed in Table B.2.2. Overall, these plots show that the routes capture a range of spatial distributions within each region, from compact to dispersed patterns. At the same time, both regions maintain broadly similar route structures and density patterns across regions.

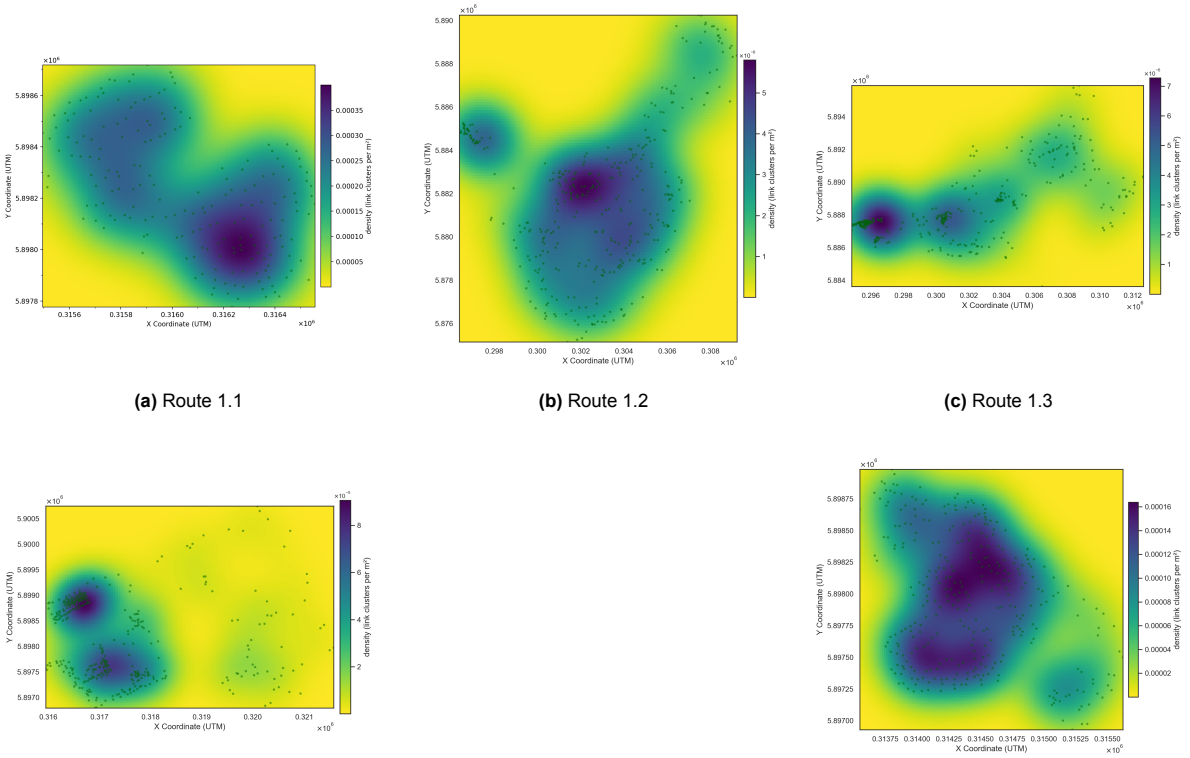

(a) Route 1.1 **(b)** Route 1.2 **(c)** Route 1.3

(d) Route 1.4 **(e)** Route 1.5

Figure B.14: Local Density of Link Cluster Midpoints for Routes of Set 1 (New Zealand)

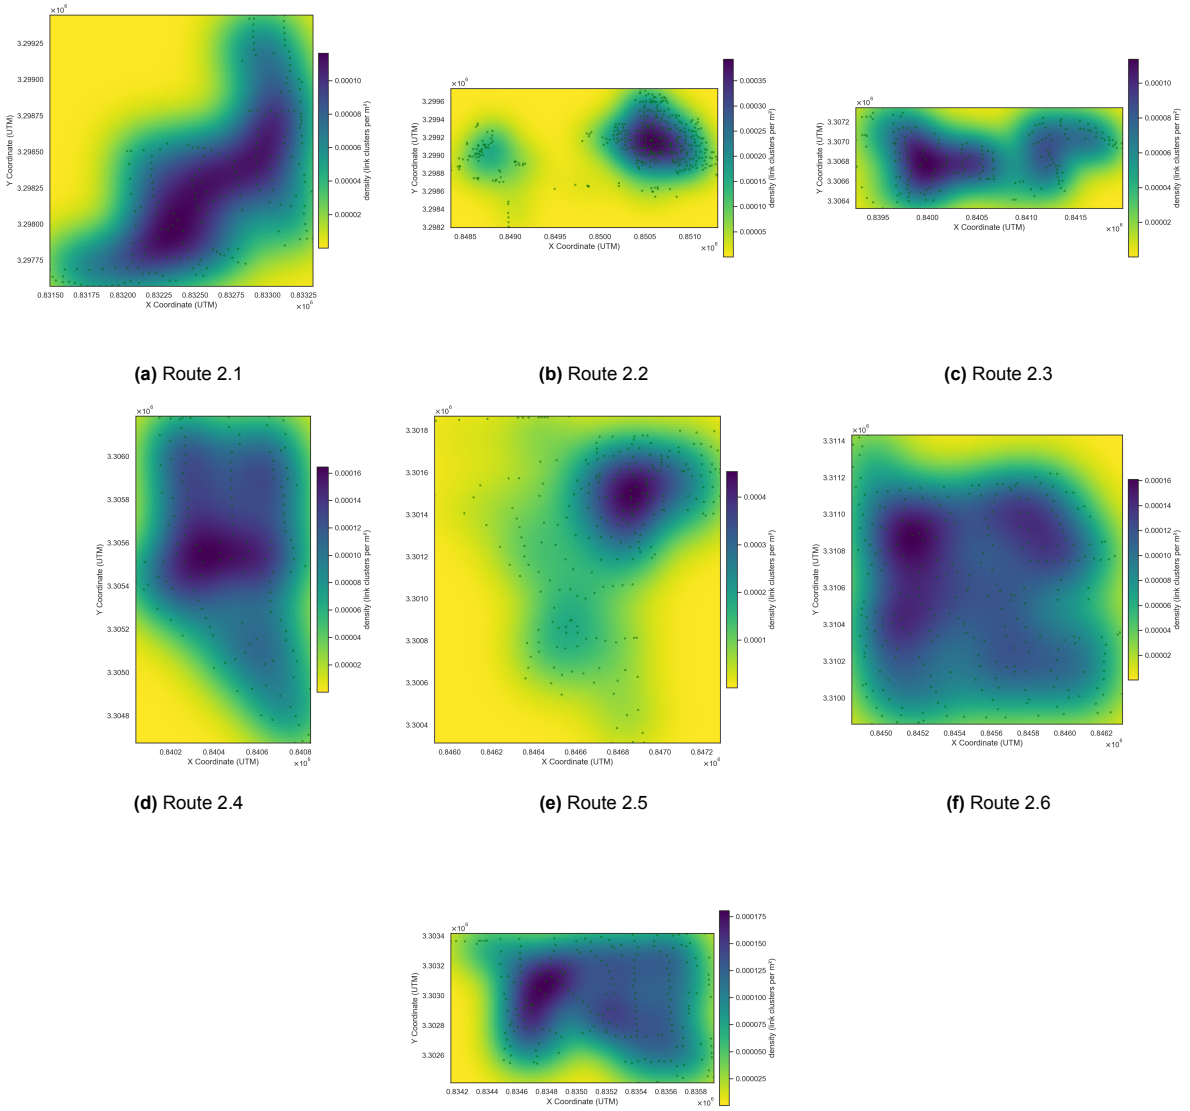

(a) Route 2.1

(b) Route 2.2

(c) Route 2.3

(d) Route 2.4

(e) Route 2.5

(f) Route 2.6

(g) Route 2.7

Figure B.15: Local Density of Link Cluster Midpoints for Routes of Set 2 (United States)

C

Pseudocode

C.1. K-Means Neighbourhood Generation

The neighbourhoods for the neighbourhood fragmentation metric are created using k-means clustering at the road-segment level. Since link clusters can represent opposite sides of the same street segment, clustering directly on link clusters may assign two clusters from the same road segment to different neighbourhoods. To prevent this, k-means is applied to aggregated road segments, each identified by a road link ID and represented by a centroid (x_r, y_r) computed as the mean midpoint of its link clusters C_r. The number of neighbourhoods is set by the target size, $k = \lceil |C| / \textit{target\_size} \rceil$. The k-means procedure is initialised by selecting k road centroids at random (seed $s = 42$). The algorithm alternates between assigning each road segment to the nearest centre and updating centres as the mean of their assigned segments. The process iterates until convergence, stopping when 100 iterations are reached or the total centre shift falls below convergence threshold $\eta = 10^{-6}$. Empty neighbourhoods are reinitialised with random road centroids. After convergence, centres are sorted lexicographically by (x, y) and relabelled $1, \ldots, k$ for consistency. The resulting neighbourhood labels are mapped back to all link clusters. The pseudocode is given below.

Algorithm 2 K-Means Neighbourhood Generation from Road Segment Midpoints

procedure KMeansNeighbourhoods(R, k, s)
    compute road centroids (x_r, y_r) for all $r \in R$
    initialise k centres by sampling road centroids from $\{(x_r, y_r) : r \in R\}$ using seed s
    $\delta \leftarrow \infty$
    while $\delta > \eta$ **and** iterations < 100 **do**
        for each road $r \in R$ **do** ▷ Step 1: assignment step
            compute distance from (x_r, y_r) to each centre
            assign r to nearest centre
        end for
        $\delta \leftarrow 0$
        for each neighbourhood $h \in \{1, \ldots, k\}$ **do** ▷ Step 2: update step
            if h has at least one assigned road **then**
                update centre as mean of assigned road centroids
            else
                reinitialise centre as a random road centroid
            end if
            update δ with centre shift for h
        end for
    end while
    sort centres lexicographically and relabel them $1, \ldots, k$ for consistency
    map road-level labels back to all link clusters
    return neighbourhood assignment $\nu(c)$ for all $c \in C$
end procedure

C.2. Operators

This section presents the neighbourhood operators used to generate new candidate solutions within the optimisation framework explained in Chapter 5.

C.2.1. Random Relocate Operator

This operator randomly selects one link cluster and inserts it at a different position in the sequence, as described in Algorithm 3.

Algorithm 3 Relocate Operator (Random Relocate Move on Link Cluster Sequence)

procedure Relocate(P_{LC}, s)
  $n \leftarrow |P_{\text{LC}}|$
  select index i uniformly from $\{1, \ldots, n\}$ using seed s
  $c \leftarrow P_{\text{LC}}[i]$
  $P^{rem} \leftarrow P_{\text{LC}}$ with c removed from index i ▷ $|P^{rem}| = n - 1$
  select insertion index j uniformly from $\{1, \ldots, n\} \setminus \{i\}$ ▷ Positions in P^{rem}
  $P' \leftarrow P^{rem}$ with c inserted at index j
  return P'
end procedure

C.2.2. Random Swap Operator

This operator randomly selects two link clusters and swaps their positions in the sequence, as described in Algorithm 4.

Algorithm 4 Swap Operator (Random Swap of Two Link Clusters)

procedure Swap(P_{LC}, s)
  $n \leftarrow |P_{\text{LC}}|$
  sample two distinct indices $i \neq j$ randomly from $\{1, \ldots, n\}$ using seed s
  $P' \leftarrow P_{\text{LC}}$
  swap $P'[i]$ and $P'[j]$
  return P'
end procedure

C.2.3. Or-Opt Operator

This operator relocates a randomly selected consecutive block of link clusters to a new position in the sequence, as described in Algorithm 5.

Algorithm 5 Or-Opt operator (Random Relocation of a Consecutive Block of Length ℓ)

procedure OrOpt(P_{LC}, ℓ, s)
  $n \leftarrow |P_{\text{LC}}|$
  select start index i uniformly from $\{1, \ldots, n - \ell + 1\}$ using seed s
  $B \leftarrow (P_{\text{LC}}[i], \ldots, P_{\text{LC}}[i + \ell - 1])$ ▷ Block
  $P^{rem} \leftarrow P_{\text{LC}}$ with block B removed from index i ▷ $|P^{rem}| = n - \ell$
  select insertion index j uniformly from $\{1, \ldots, |P^{rem}| + 1\} \setminus \{i\}$ ▷ Positions in P^{rem}
  $P' \leftarrow P^{rem}$ with block B inserted at index j
  return P'
end procedure

D

Additional Results of Single-Objective Optimisation

This appendix contains the route visualisations for the exact models on the test set. Only the travel-time-optimal route is shown in the main text in Subsection 4.3.2, whereas the remaining seven routes generated by the visual attractiveness models are included here for reference. Additionally, Figure D.1 shows the manually defined neighbourhoods used to compute the neighbourhood fragmentation objective and the corresponding post-hoc metric values. Because the test set is small and extracted from a larger case study, neighbourhoods were defined manually based on spatial proximity and road network connectivity, as inspected using AMCS software. This is done to avoid unrealistic groupings that may arise from automatic clustering.

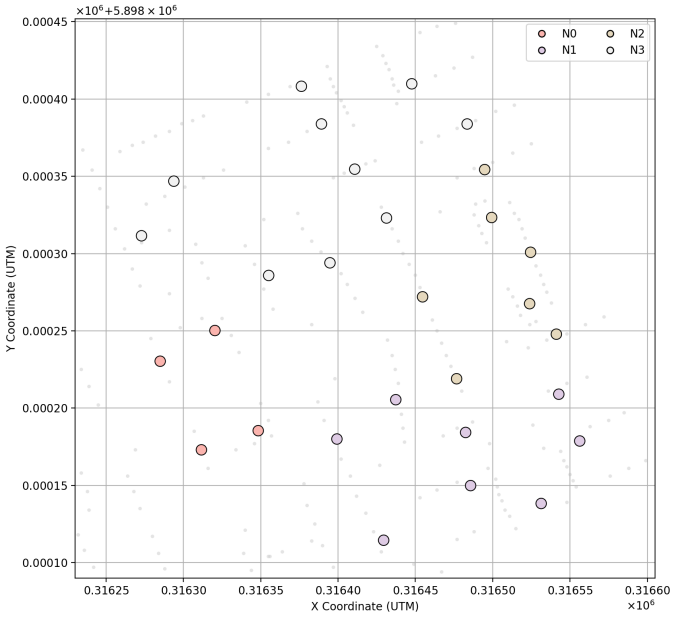

Figure D.1: Manually Defined Neighbourhoods for the Test Set

All route plots in Figure D.2, Figure D.3, and Figure D.4 use a consistent visual encoding. Link clusters are shown at their midpoints with visitation sequence labels and each link cluster's start and end order nodes are visualised explicitly, with directed arcs drawn between these endpoints. The four route segments are colour-coded, revisit pairs are connected with green lines, and the longest 20% of directed moves between successive link clusters are highlighted in purple. The figures are organised accord-

ing to the three objective categories identified in the quantitative analysis, namely grouping-based, sequence-based, and single-purpose objectives.

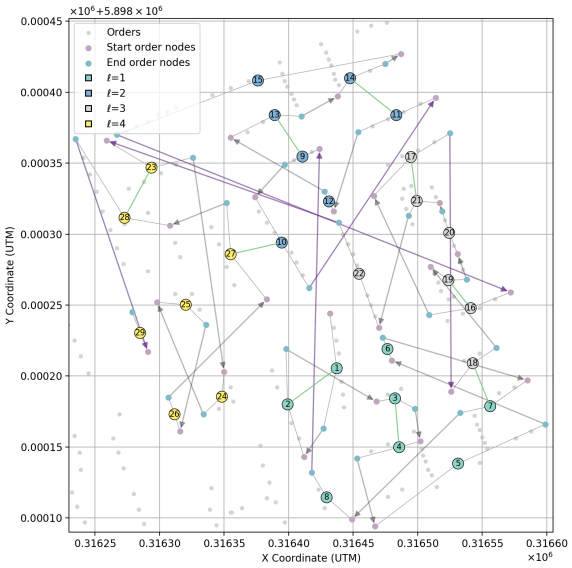

(a) Route of Model Optimised for Medoid Spread (MS)

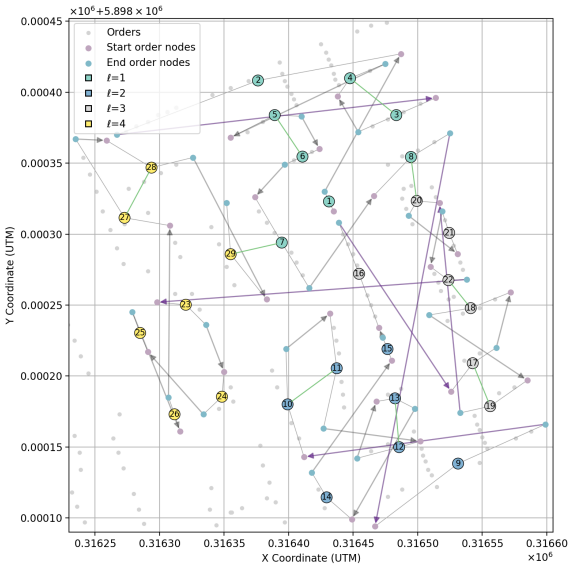

(b) Route of Model Optimised for Pairwise Spread (PS)

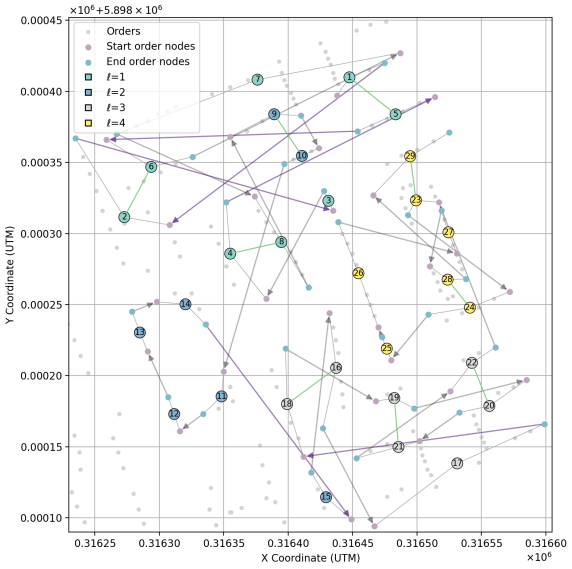

(c) Route of Model Optimised for Neighbourhood Fragmentation (NF)

Figure D.2: Routes of Models with Grouping-Based Objectives (MS, PS, NF) on the Abstract Network for the Test Set

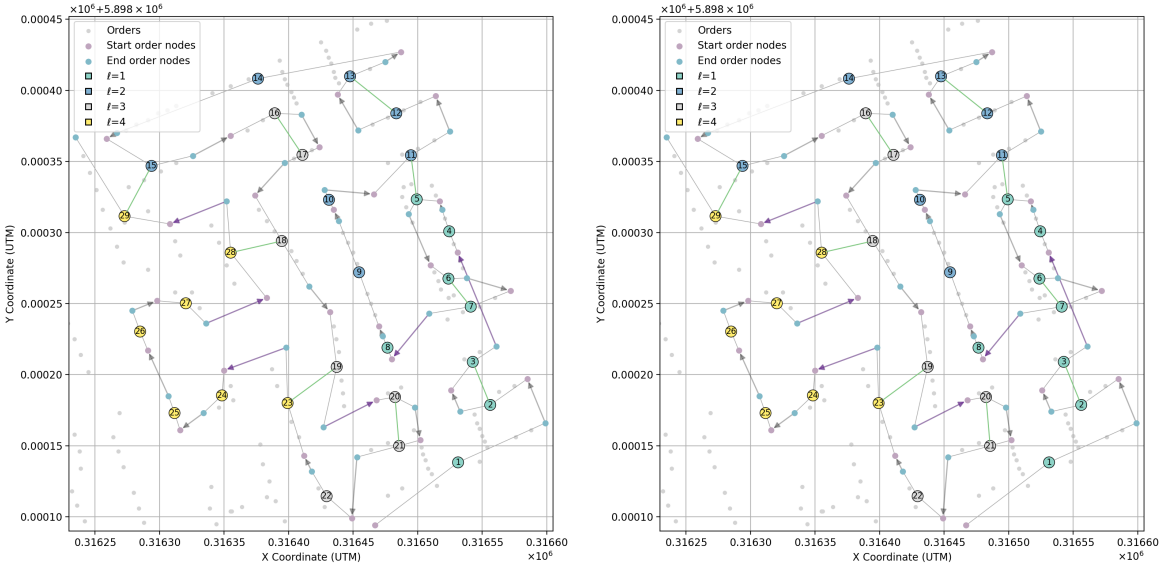

(a) Route of Model Optimised for Top Longest Jumps (TJ)

(b) Route of Model Optimised for Local Jumps (LJ)

Figure D.3: Routes of Models with Sequence-Based Objectives (TJ and LJ) on the Abstract Network for the Test Set

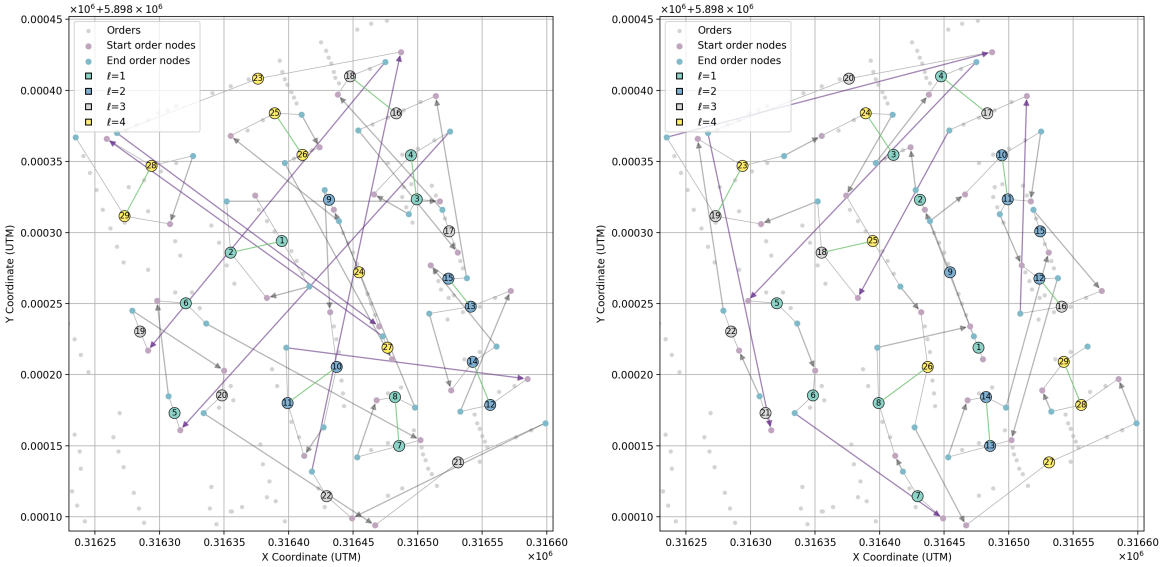

(a) Route of Model Optimised for Revisit Span (RS)

(b) Route of Model Optimised for Intra-Route Crossings (IC)

Figure D.4: Routes of Models with Single-Purpose Objectives (RS and IC) on the Abstract Network for the Test Set

E

Additional Results of Case Studies

This appendix contains supplementary results for the numerical and visual evaluation of the optimisation framework relative to existing AMCS routing solutions in Chapter 7.

E.1. Additional Numerical Results

This section presents the raw route metrics behind the relative results in Table 7.1. First, Table E.1 shows the travel-time-optimal routes from the current AMCS method. Second, Table E.2 reports routes from the exact model minimising travel time with a time limit of 300 seconds. These are similar to the AMCS routes in total travel time. Finally, Table E.3 presents the best solutions found by the optimisation framework. All metrics are calculated as explained in Subsection 4.2.2. For the optimisation framework, the reported objective values correspond to the best-found solutions. The initial solution from the exact model minimising travel time has an objective value of 5.5.

Table E.1: Raw Metrics for the AMCS Baseline Routes

Case Study	TT (s)	NF	RS (s)	IC	LJ (m)	Runtime (s)
Set 1 – Route 1	6123	47	47 930	86	11 527	93
Set 1 – Route 2	17 533	162	295 813	101	141 831	193
Set 1 – Route 3	14 647	94	162 353	81	103 299	133
Set 1 – Route 4	12 004	83	48 296	91	45 692	361
Set 1 – Route 5	10 238	92	77 743	92	17 927	1156
Set 2 – Route 1	12 317	90	92 891	61	16 510	5
Set 2 – Route 2	16 104	131	187 575	124	25 459	32
Set 2 – Route 3	8955	34	41 878	40	13 149	4
Set 2 – Route 4	5449	26	157 693	26	8683	4
Set 2 – Route 5	10 868	171	167 271	66	16 517	13
Set 2 – Route 6	12 230	136	160 076	68	21 042	34
Set 2 – Route 7	10 369	111	227 124	50	14 449	10

Table E.2: Raw Metrics for the Initial Routes (Exact Model Minimising Travel Time)

Case Study	TT (s)	NF	RS (s)	IC	LJ (m)	Runtime (s)	Gap (%)
Set 1 – Route 1	6112	48	55 205	70	11 065	300	0.0
Set 1 – Route 2	17 513	179	336 666	92	137 834	300	0.0
Set 1 – Route 3	15 729	87	231 126	99	107 594	300	7.3
Set 1 – Route 4	11 845	101	59 231	77	45 212	13	0.0
Set 1 – Route 5	10 312	97	104 829	97	18 302	300	4.4
Set 2 – Route 1	12 329	75	103 238	52	15 889	300	2.6

Continued on next page

Table E.2: Raw Metrics for the Initial Routes (Exact Model Minimising Travel Time) *(Continued)*

Case Study	TT (s)	NF	RS (s)	IC	LJ (m)	Runtime (s)	Gap (%)
Set 2 – Route 2	17 145	235	196 957	124	28 250	305	8.6
Set 2 – Route 3	8782	39	68 205	38	14 738	119	0.0
Set 2 – Route 4	5408	26	163 391	33	8824	300	0.6
Set 2 – Route 5	10 595	140	131 178	63	15 303	6	0.0
Set 2 – Route 6	12 157	150	251 688	75	21 203	6	0.0
Set 2 – Route 7	10 307	114	192 149	52	14 151	9	0.0

Table E.3: Raw Metrics for the Optimisation Framework Routes

Case Study	TT (s)	NF	RS (s)	IC	LJ (m)	F	Runtime (s)
Set 1 – Route 1	8892	38	876	2	12 321	3.82	7681
Set 1 – Route 2	30 097	81	47 804	27	190 718	4.94	49 278
Set 1 – Route 3	24 286	53	12 743	8	151 316	4.25	39 301
Set 1 – Route 4	11 845	77	59 231	101	45 212	5.5	49 812
Set 1 – Route 5	16 155	70	5540	11	25 980	4.43	37 591
Set 2 – Route 1	17 316	17	7503	6	20 089	3.80	11 757
Set 2 – Route 2	22 620	75	11 337	19	29 443	3.64	59 602
Set 2 – Route 3	12 299	15	7239	2	16 348	3.74	6625
Set 2 – Route 4	7413	7	15 803	2	9571	3.60	4001
Set 2 – Route 5	15 184	20	9340	19	16 472	3.84	15 776
Set 2 – Route 6	17 289	42	18 213	17	25 138	3.96	19 581
Set 2 – Route 7	14 148	15	13 577	2	17 184	3.59	11 627

E.2. Additional Visual Results

In this section Figure E.1–Figure E.11 compare abstract network visualisations of the AMCS baseline routes and the optimisation framework routes for all case studies, except route 1 in set 1 (see Section 7.3). Link clusters are shown at their midpoints and labelled by visitation order, with grey arrows indicating travel direction. Revisit pairs are marked with green lines, crossing arcs in blue, and terminal nodes and arcs omitted. Link clusters are coloured by neighbourhood or a colour gradient for visitation order.

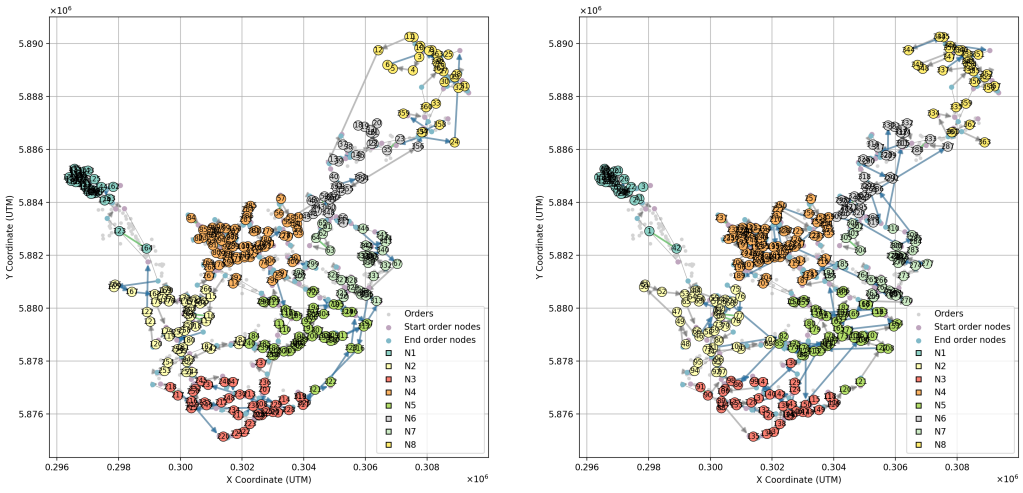

(a) AMCS Baseline Route with Neighbourhood Colouring

(b) Optimisation Framework Route with Neighbourhood Colouring

Figure E.1: Comparison of Routing Solutions on the Abstract Network for Route 2 of Set 1 (New Zealand)

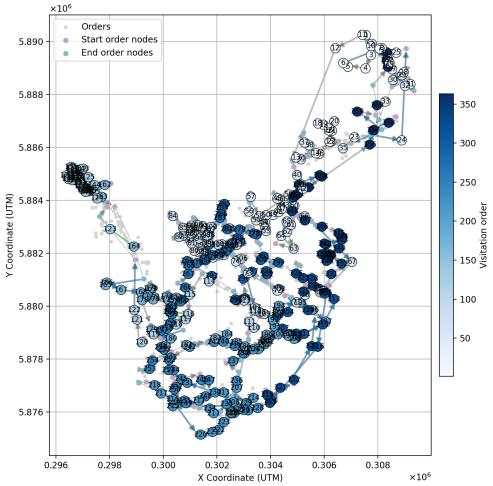

(c) AMCS Baseline Route with Visitation Order Colouring

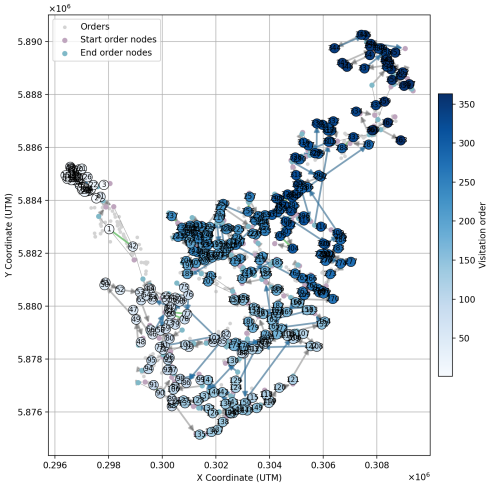

(d) Optimisation Framework Route with Visitation Order Colouring

Figure E.1: Comparison of Routing Solutions on the Abstract Network for Route 2 of Set 1 (New Zealand) *(Continued)*

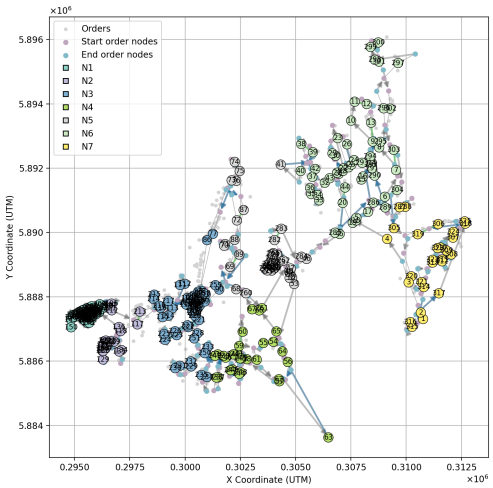

(a) AMCS Baseline Route with Neighbourhood Colouring

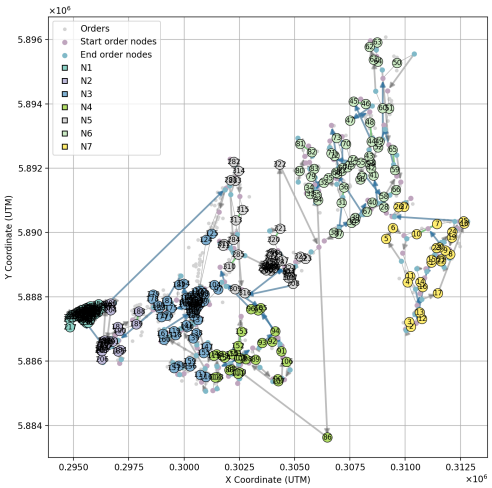

(b) Optimisation Framework Route with Neighbourhood Colouring

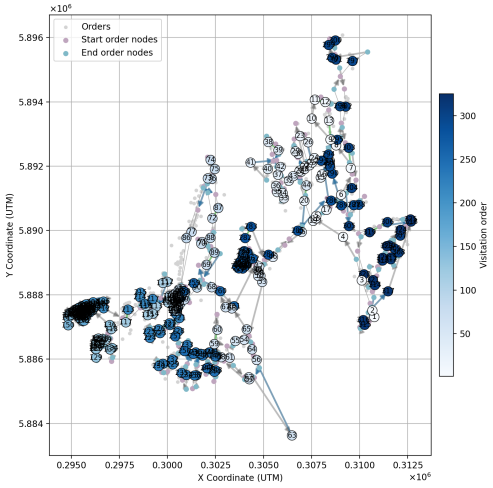

(c) AMCS Baseline Route with Visitation Order Colouring

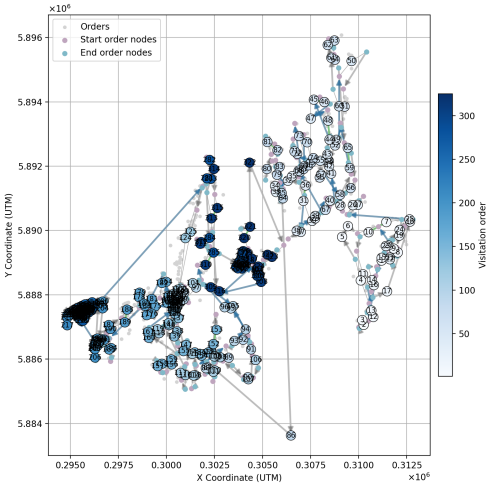

(d) Optimisation Framework Route with Visitation Order Colouring

Figure E.2: Comparison of Routing Solutions on the Abstract Network for Route 3 of Set 1 (New Zealand)

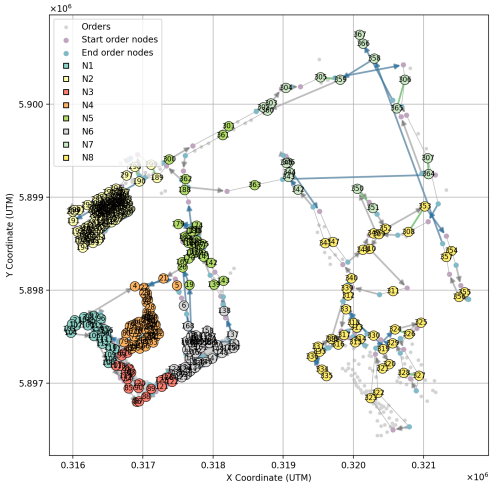

(a) AMCS Baseline Route with Neighbourhood Colouring

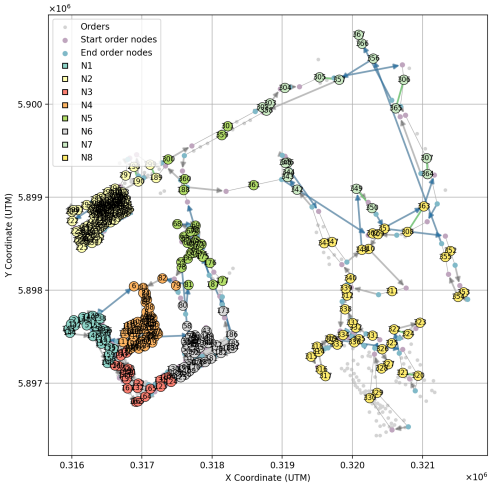

(b) Optimisation Framework Route with Neighbourhood Colouring

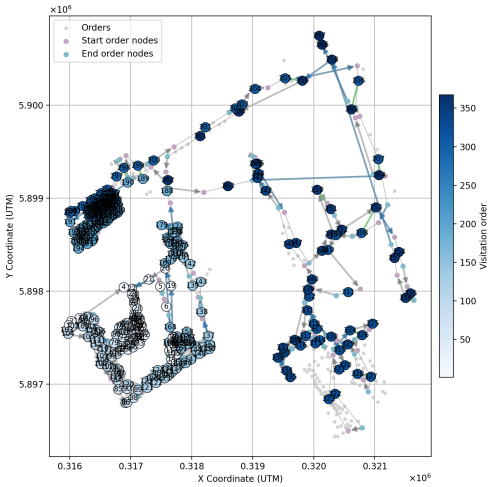

(c) AMCS Baseline Route with Visitation Order Colouring

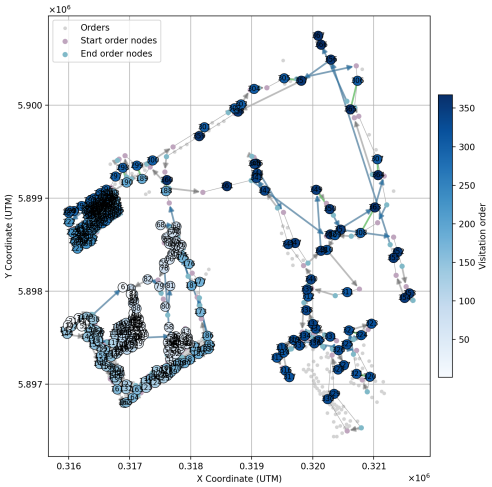

(d) Optimisation Framework Route with Visitation Order Colouring

Figure E.3: Comparison of Routing Solutions on the Abstract Network for Route 4 of Set 1 (New Zealand)

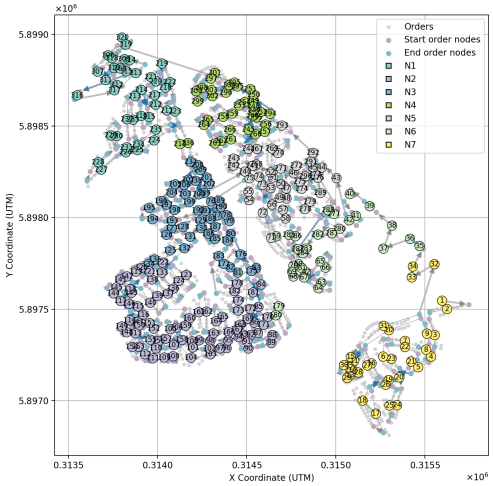

(a) AMCS Baseline Route with Neighbourhood Colouring

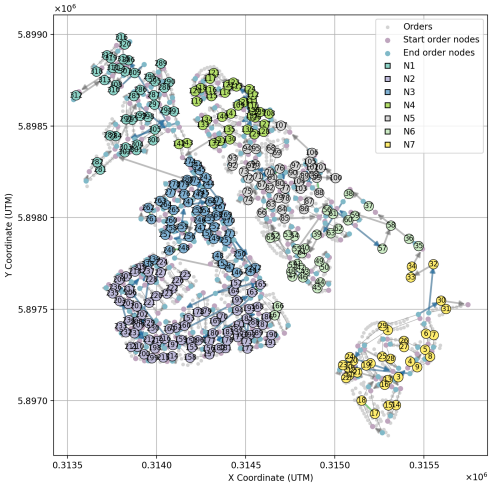

(b) Optimisation Framework Route with Neighbourhood Colouring

Figure E.4: Comparison of Routing Solutions on the Abstract Network for Route 5 of Set 1 (New Zealand)

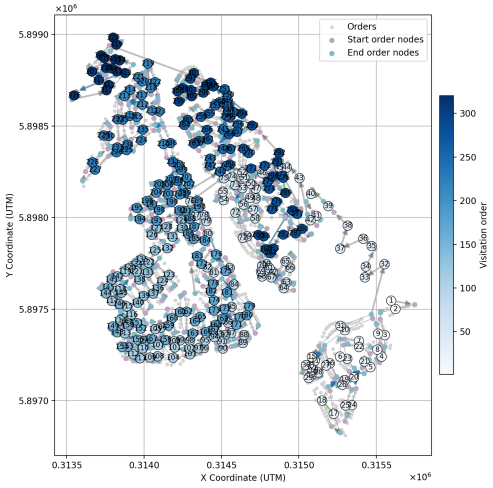

(c) AMCS Baseline Route with Visitation Order Colouring

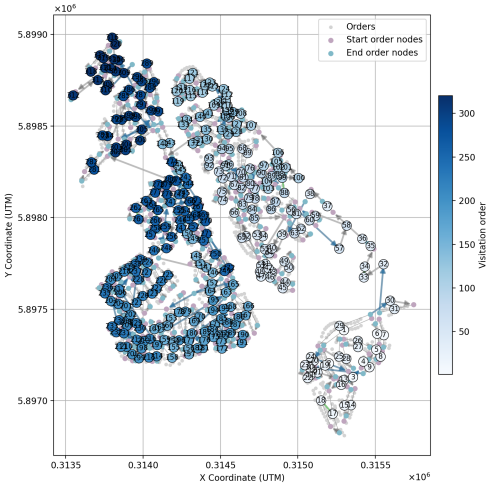

(d) Optimisation Framework Route with Visitation Order Colouring

Figure E.4: Comparison of Routing Solutions on the Abstract Network for Route 5 of Set 1 (New Zealand) *(Continued)*

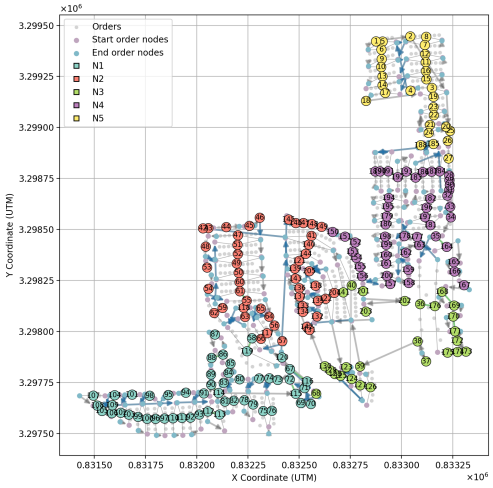

(a) AMCS Baseline Route with Neighbourhood Colouring

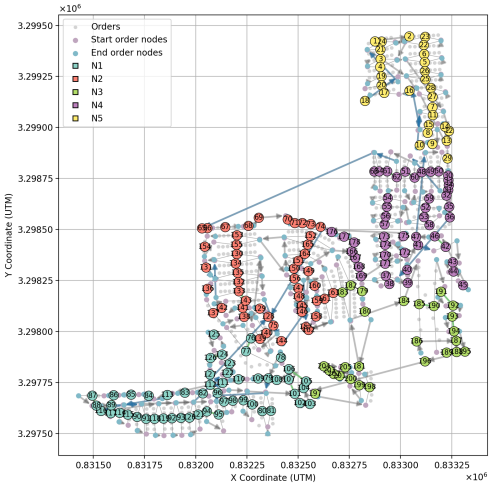

(b) Optimisation Framework Route with Neighbourhood Colouring

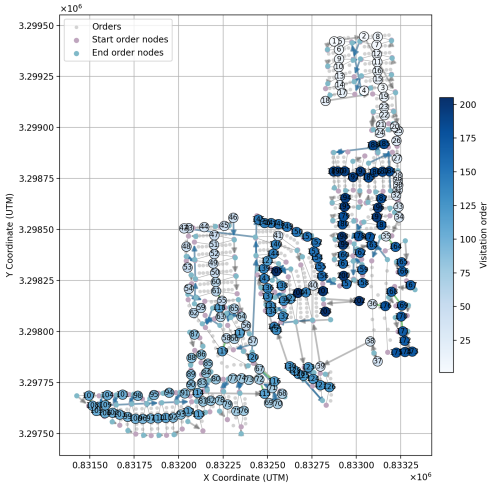

(c) AMCS Baseline Route with Visitation Order Colouring

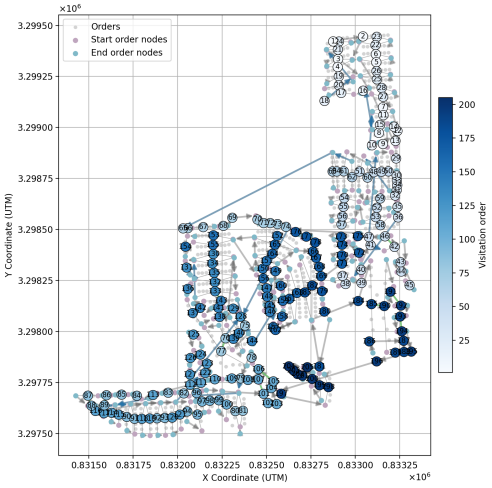

(d) Optimisation Framework Route with Visitation Order Colouring

Figure E.5: Comparison of Routing Solutions on the Abstract Network for Route 1 of Set 2 (United States)

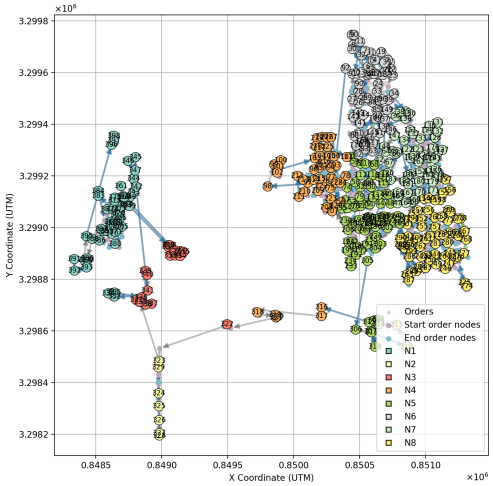

(a) AMCS Baseline Route with Neighbourhood Colouring

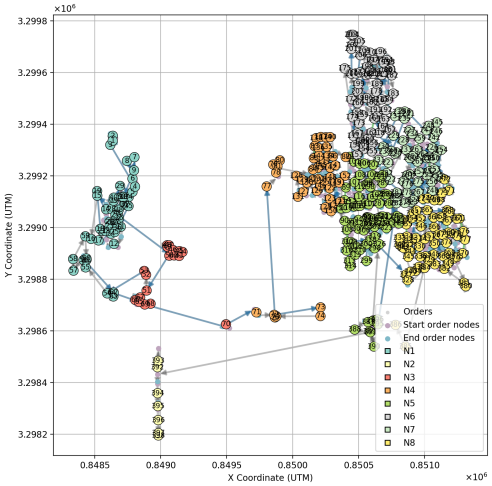

(b) Optimisation Framework Route with Neighbourhood Colouring

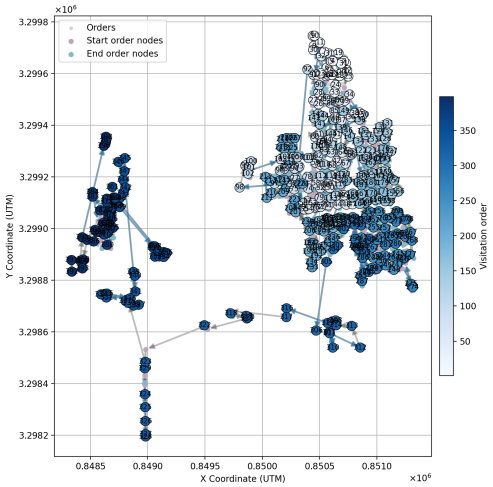

(c) AMCS Baseline Route with Visitation Order Colouring

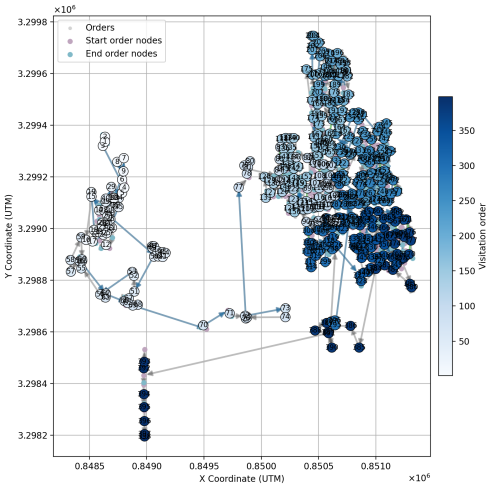

(d) Optimisation Framework Route with Visitation Order Colouring

Figure E.6: Comparison of Routing Solutions on the Abstract Network for Route 2 of Set 2 (United States)

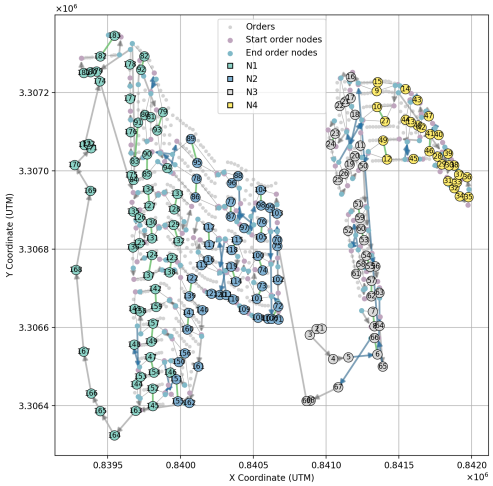

(a) AMCS Baseline Route with Neighbourhood Colouring

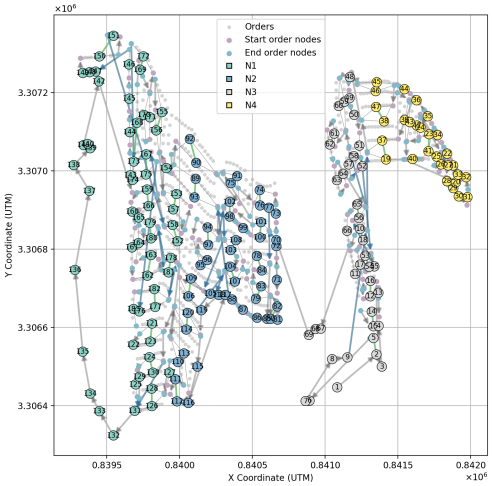

(b) Optimisation Framework Route with Neighbourhood Colouring

Figure E.7: Comparison of Routing Solutions on the Abstract Network for Route 3 of Set 2 (United States)

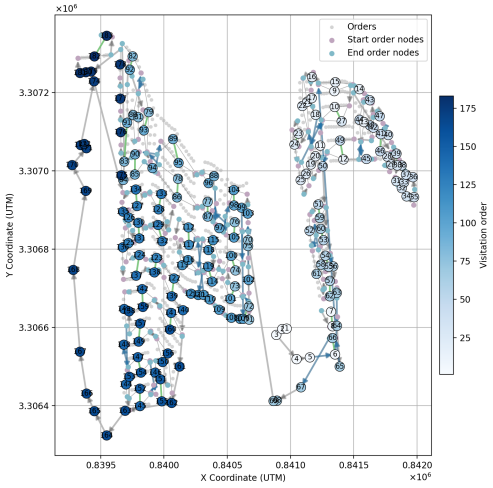

(c) AMCS Baseline Route with Visitation Order Colouring

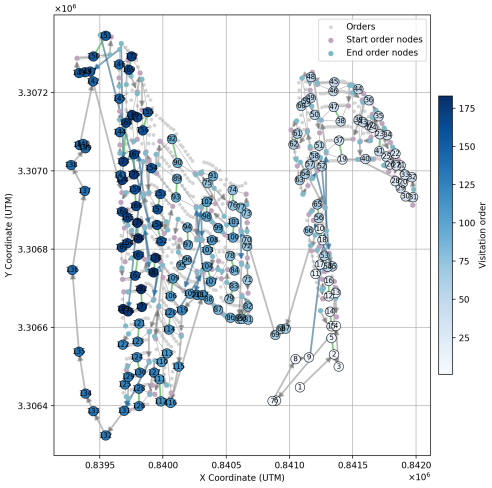

(d) Optimisation Framework Route with Visitation Order Colouring

Figure E.7: Comparison of Routing Solutions on the Abstract Network for Route 3 of Set 2 (United States) *(Continued)*

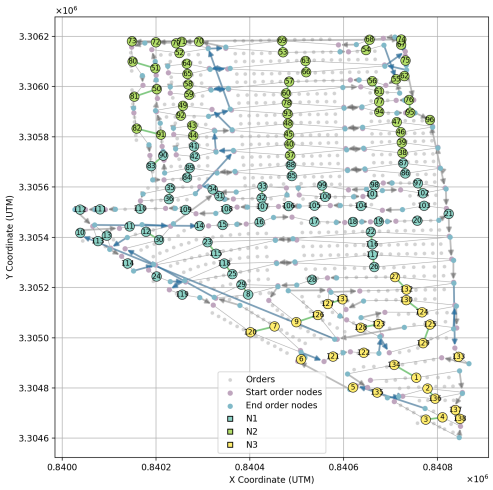

(a) AMCS Baseline Route with Neighbourhood Colouring

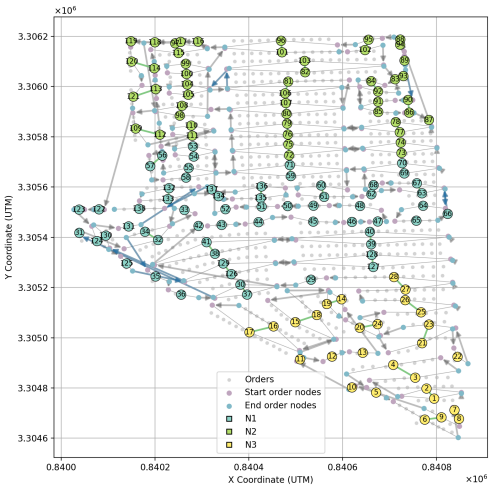

(b) Optimisation Framework Route with Neighbourhood Colouring

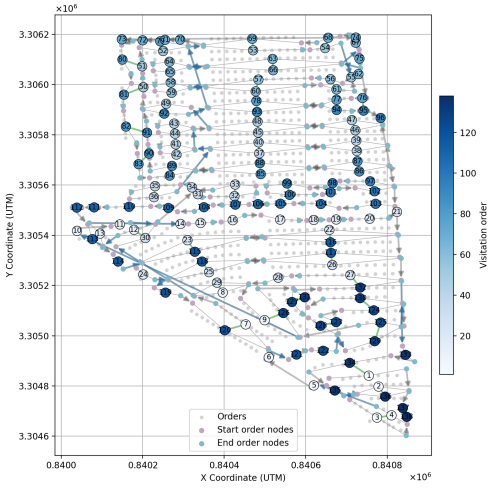

(c) AMCS Baseline Route with Visitation Order Colouring

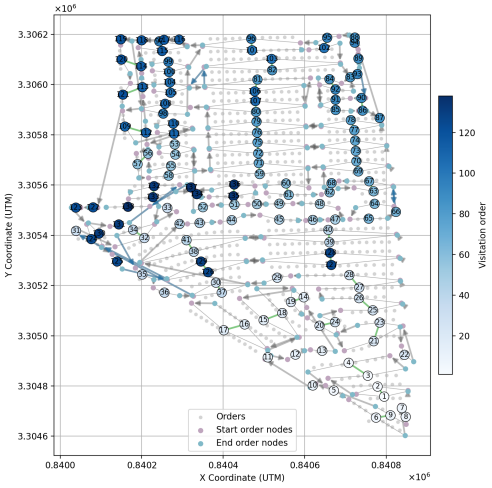

(d) Optimisation Framework Route with Visitation Order Colouring

Figure E.8: Comparison of Routing Solutions on the Abstract Network for Route 4 of Set 2 (United States)

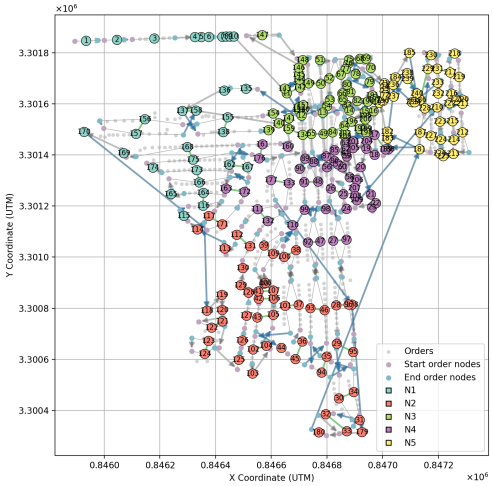

(a) AMCS Baseline Route with Neighbourhood Colouring

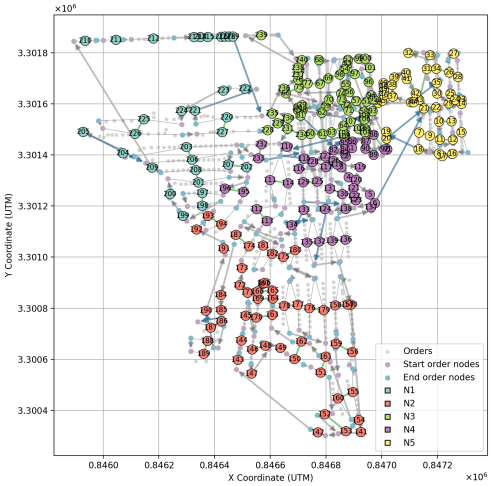

(b) Optimisation Framework Route with Neighbourhood Colouring

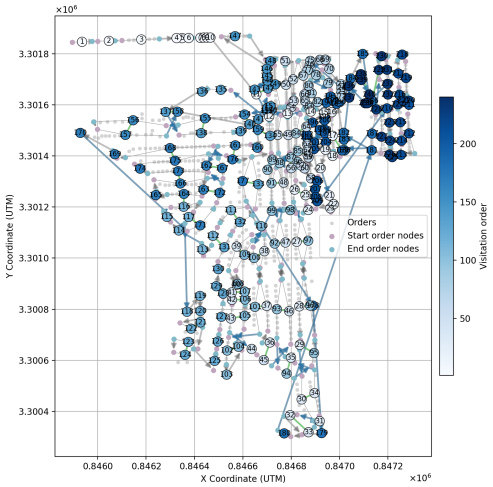

(c) AMCS Baseline Route with Visitation Order Colouring

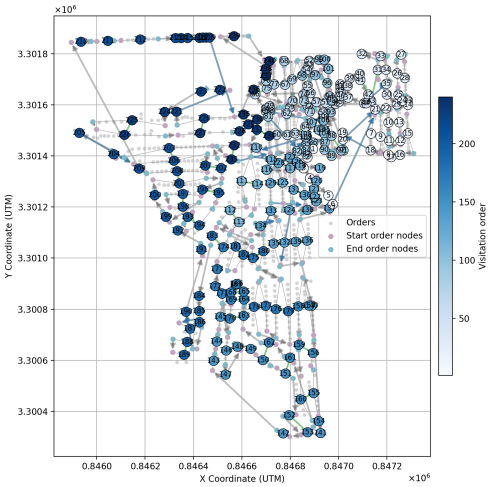

(d) Optimisation Framework Route with Visitation Order Colouring

Figure E.9: Comparison of Routing Solutions on the Abstract Network for Route 5 of Set 2 (United States)

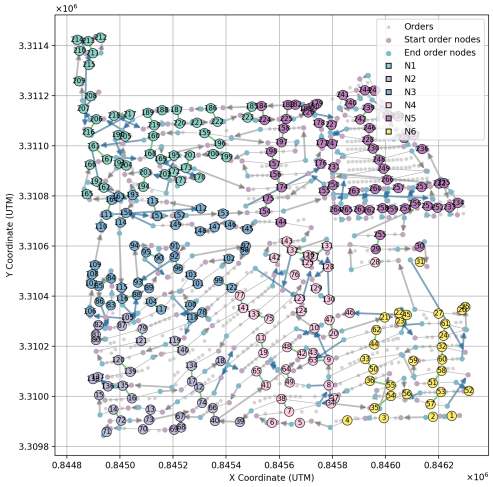

(a) AMCS Baseline Route with Neighbourhood Colouring

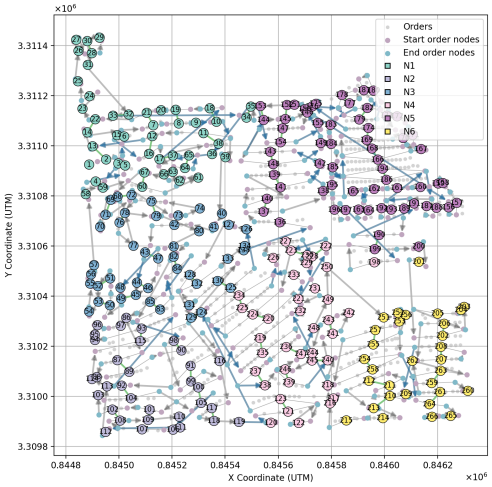

(b) Optimisation Framework Route with Neighbourhood Colouring

Figure E.10: Comparison of Routing Solutions on the Abstract Network for Route 6 of Set 2 (United States)

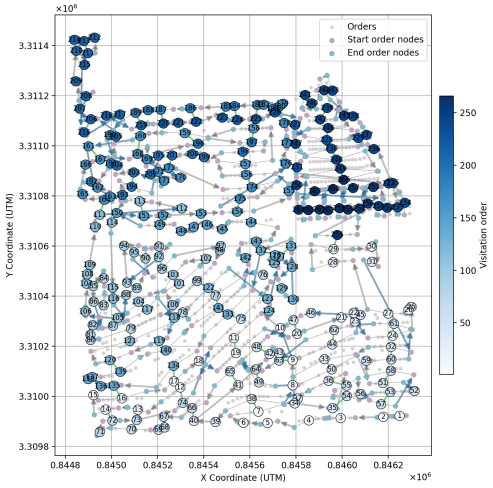

(c) AMCS Baseline Route with Visitation Order Colouring

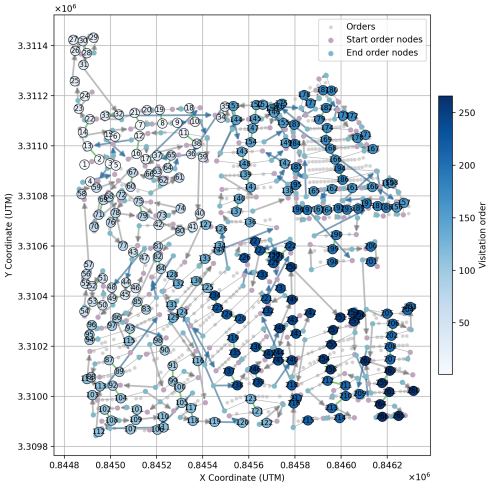

(d) Optimisation Framework Route with Visitation Order Colouring

Figure E.10: Comparison of Routing Solutions on the Abstract Network for Route 6 of Set 2 (United States) *(Continued)*

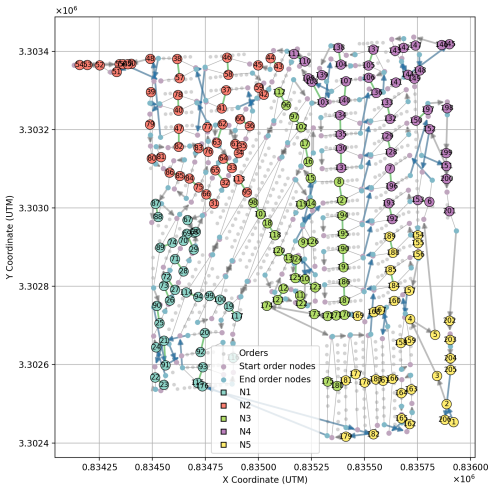

(a) AMCS Baseline Route with Neighbourhood Colouring

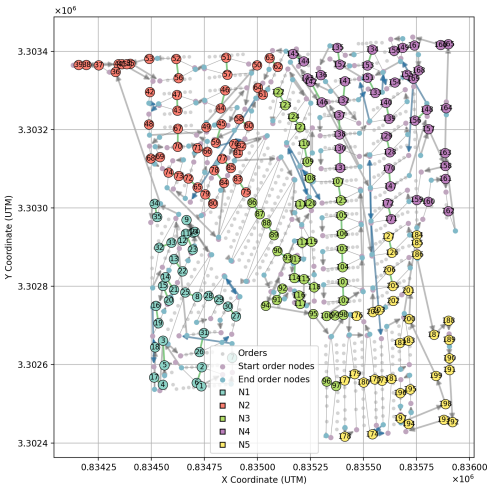

(b) Optimisation Framework Route with Neighbourhood Colouring

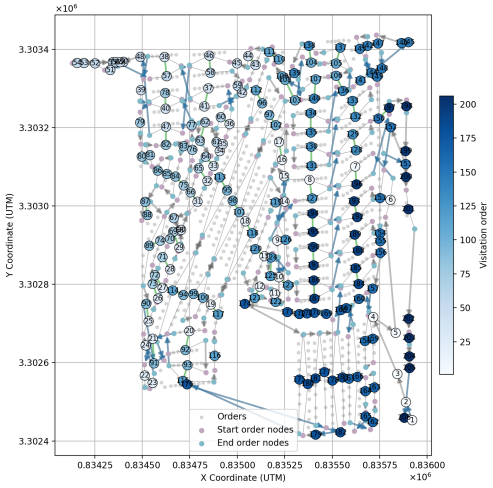

(c) AMCS Baseline Route with Visitation Order Colouring

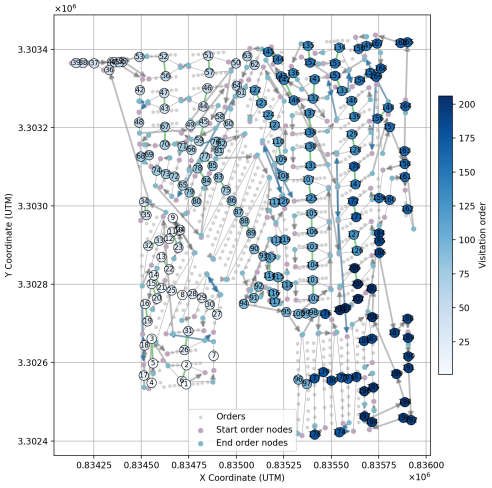

(d) Optimisation Framework Route with Visitation Order Colouring

Figure E.11: Comparison of Routing Solutions on the Abstract Network for Route 7 of Set 2 (United States)