ROTTERDAM SCHOOL OF MANAGEMENT,

ERASMUS UNIVERSITY ROTTERDAM

MASTER THESIS

MScBA BUSINESS ANALYTICS & MANAGEMENT

Optimizing Routes in Attended Home Services: Balancing Efficiency, Complexity, and Customer Service

14.06.2024

Author: Félix Raphael (682511)

Thesis Coach: dr. L.P. Luuk Veelenturf

Thesis Co-reader: prof. dr. ir. N.A.H. Niels Agatz

Acknowledgments

I would like to express my gratitude and appreciation towards everyone who has supported me throughout my studies at Rotterdam School of Management, and in particular throughout the process of completing this thesis.

With that, I first want to thank my thesis supervisor and co-reader for their constant support and strong engagement. In our close collaboration, I gained valuable knowledge and I highly appreciate having worked with you. I very much enjoyed our discussions and exchanges.

Also, I want to express my gratitude towards the developers of the PyVRP solver, Wouda et al. (2024). The Python package has greatly facilitated my work through its vast capabilities and its ease of use, which I am very thankful for.

Finally, I want to thank my family, friends and partner for the mental support they provided me with along my studies.

Preface

Félix Raphael

Rotterdam, 14.06.2024

Executive Summary

Companies operating in the sector of Attended Home Deliveries and -Services face the challenge of planning efficient routes to deliver to their customers while maintaining high customer service levels. Numerous companies such as the Dutch e-grocer Picnic already tackle this challenge by leveraging extensively studied, sophisticated methods such as a priori optimization.

This thesis investigates how companies can balance planning complexity, route efficiency and customer service levels by leveraging optimized appointment-day offerings and a priori routes to minimize their total travel distance in delivery operations. For this, four different strategies are explored in a comparative analysis on synthetic data from 2,500 delivery instances across three Dutch provinces.

In a first stage, the strategies partition customers into groups, for each of which routes are pre-planned and different appointment-day choices are assigned. After customer preferences are incorporated, the strategies, in a second stage, finalize daily routes to visit each customer. This is either done by adhering to the initial visiting order from the pre-planned routes or by reoptimizing each daily route.

The study provides support for several findings from academic literature. Further, the results highlight the importance of companies selecting the best grouping method in the first stage of their strategy. In contrast, reoptimizing routes is found to only add a marginal improvement to the efficiency of routes. Overall, the study demonstrates that optimizing appointment-day offerings can yield substantial benefits to the efficiency of routes, thereby contributing to the overall profitability of a company.

Contents

List of Figures

List of Tables

List of Abbreviations

AHD Attended Home Delivery

AHS Attended Home Service

ALNS Adaptive Large Neighbourhood Search

CCP Capacitated Clustering Problem

CVRP Capacitated Vehicle Routing Problem

CVRS Computerized Vehicle Routing and Scheduling

DTWAVRP Discrete Time Window Assignment Vehicle Routing Problem

DSS Decision Support System

GA Genetic Algorithm

HCCSP Home Care Crew Scheduling Problem

HGS Hybrid-Genetic Search

K-GWO K-means Grey Wolf Optimizer

MDVRPTW Multi-Depot Vehicle Routing Problem with Time Windows

MILP Mixed-Integer Linear Program

ORION On-Road Integrated Optimization and Navigation

OSRM Open Source Routing Machine

PSO Particle Swarm Optimization

RAS Routing and Appointment Scheduling

SMTWAVRP Stochastic Multi-Period Time Window Assignment Vehicle Routing Problem

STSM Strategic Time Slot Management

TS Tabu Search

TSMP Time Slot Management Problem

TWAVRP Time Window Assignment Vehicle Routing Problem

UPS United Parcel Service

VRP Vehicle Routing Problem

VRPHTW Vehicle Routing Problem with Hard Time Windows

VRPMPTW Multi-Objective Vehicle Routing Problem with Multiple Prioritized Time Windows

VRPSTW Vehicle Routing Problem with Soft Time Windows

VRPTW Vehicle Routing Problem with Time Windows

1 Introduction

Companies in the Attended Home Delivery (AHD) and Attended Home Service (AHS) sector face a daunting challenge: designing efficient routes to deliver their services while adhering to various constraints and customer requirements. Addressing such a problem typically involves a multi-stage approach, in which decisions taken in the previous stage can greatly affect the feasibility and economic profitability of the following decision stages (Cordeau et al., 2023, p. 550). In particular, the optimization of AHD/AHS problems typically requires solving a two-stage problem, combining appointment scheduling and solving a Vehicle Routing Problem (VRP) with the goal of reducing operating cost, minimizing travel time and travel distance, while meeting customers' requirements (Bruck et al., 2020, p. 137). In particular, Rowell et al. (2012) find that customer requirements, total cost, travel time and travel distance are weighed as the most important routing factors by surveyed companies.

In increasingly complex supply chains and customer networks, companies are turning to computerized planning methods to devise delivery routes and schedules, thereby outperforming previously adopted manual approaches, as showcased by Côté et al. (2024). While larger corporations may benefit from Computerized Vehicle Routing and Scheduling (CVRS) systems, the software tools are often less suitable and financially overwhelming for smaller companies. In fact, AHS companies often manually plan deliveries and allocate significant resources to assess schedules (Bruck et al., 2020, p. 137). Despite low adoption rates, studies indicate substantial benefits of CVRS implementation, including transport cost reduction, reduction in fuel and environmental impact, improved customer service and effective strategic planning (Nicolas Rincon-Garcia & Cherrett, 2018, p. 119). Notably, UPS' On-Road Integrated Optimisation and Navigation (ORION) project highlights the transformative potential of algorithmic routing solutions. ORION is estimated to save UPS $300 to $400 million annually, while contributing to UPS' sustainability efforts by reducing driven miles and fuel consumption (Holland et al., 2017, p. 19).

Utility providers, typically operating within the AHS sector, often face challenges in accommodating customer scheduling preferences for installation or maintenance services (Cordeau et al., 2023, p. 550). To overcome these challenges, companies optimize for the time windows that they offer to their customers to choose from, a field of research that is referred to as demand management and time slot management. Services in the AHS sector can generally be distinguished into ordinary (planned) and extraordinary (emergency) categories (Cordeau et al., 2023, pp. 549–550). While some AHD/AHS problems are characterized by a dynamic customer landscape and therefore display demand and planning uncertainty, more ordinary AHD/AHS operations, such as oil deliveries to customers, exhibit deterministic and more plannable demand patterns, allowing for more predictable scheduling and routing approaches.

Addressing AHD/AHS issues often requires quick solutions for each new appointment within a service day, thereby requiring vast resource availability such as for computation. A priori optimization offers an efficient and resource-friendly alternative through the adoption of pre-planned routes for known customer locations. These pre-planned routes specify the order of customer visits for a delivery vehicle and can serve as a starting point for reoptimization strategies, if time permits on the day of service (Bertsimas et al., 1990; Campbell & Thomas, 2008).

This thesis investigates *how ordinary AHS providers can leverage optimized appointment-day offerings and a priori routes to minimize their total travel distance and effectively manage the trade-off between planning complexity, routing efficiency, and customer service*. This is done by comparing four different strategies on synthetic customer locations of a real-world company that operates in the ordinary AHS sector, thereby closely examining the underlying decision stages of appointment-day assignment and final route determination. This problem statement is addressed while serving the following sub-questions:

1. *What are key characteristics and advantages of scientifically validated methods for optimizing appointment-day offerings and a priori routes in order to minimize the total travel distance?*

2. *How can the performance of different strategies, which optimize appointment-day offerings and a priori routes, be accurately assessed ?*

3. *How do different strategies perform in terms of minimizing total travel distance and which factors influence their effectiveness?*

4. *Which insights can be derived from the comparative performance analysis of different strategies, and how can these insights inform actionable strategies for balancing complexity, efficiency, and customer service?*

As highlighted by the example of the ORION project, solving this problem is of high managerial relevance for companies in the AHD/AHS sector, especially those providing ordinary, plannable services. This high practical relevance is supported by the observation that researchers in this field often collaborate with industry partners (Waßmuth et al., 2023, p. 812). By implementing optimized scheduling and routing operations, companies may be able to reduce resource utilization and improve their overall resource allocation. Resulting cost savings could further improve a company's profitability. Additionally, the impact of demand- and time slot management on customer service levels magnifies the importance of effectively balancing the trade-off between operational complexity and routing efficiency. In other words, deciding on appointment-day offerings and developing a priori routes ultimately has implications for the level of customer service, which is why the findings of this study provide actionable insights for guiding ordinary AHD/AHS companies towards more efficient and customer-centric planning methods in developing delivery routes.

This study further contributes to academic literature in the related fields of demand management, time slot management, a priori optimization and vehicle routing in several ways. First, this thesis provides novel insights by comparing different multi-stage strategies while leveraging different academically validated methods in an ordinary AHS context, which, to the best of my knowledge, has not been addressed in academic literature to this date. Moreover, this work addresses previous assertions made in academic literature. In particular, it investigates the impact of clustering on routing performance and its implications for the trade-off between routing efficiency and planning complexity. The analysis further provides insights into the effect of offering appointment-days on the customer service level, addressed in Agatz et al. (2011), Bühler et al. (2016), Côté et al. (2024), and Zhan et al. (2021), as well as on the trade-off between total cost and service level (Agatz et al., 2011; Beheshti et al., 2015).

The remainder of this thesis is structured as follows. Section 2 entails a brief review of relevant academic literature, addressing sub-question 1. Section 3 links the studied problem to a real-world case and details underlying assumptions and considerations. Section 4 presents the general design of the decision making process involved, details the methodology of the 4 studied strategies, provides a description of upper and lower bound solutions and finally presents the solving method used in this study for solving routing problems. Section 5 lays out the experimental design by explaining the performance evaluation of strategies, thereby addressing sub-question 2. Moreover, this section formulates expectations different strategies regarding efficiency, complexity and service level, and finally details the data generation process. Section 6 explores the results of this study and with that investigates sub-question 3. Lastly, section 7 addresses sub-question 4 by providing conclusions and presents limitations of the study as well as suggestions for future research.

2 Literature Review

Within this section, an introduction on Attended Home Delivery and -Services characteristics and processes is presented in chapter 2.1. Further, academic literature on demand management, VRPs and their solution approaches are detailed in chapters 2.2, 2.3 and 2.4.

2.1 AHD and AHS

Many of the products and services that consumers order online, such as furniture and groceries, can only be delivered to customers in their presence, which is also referred to as attended delivery. This condition makes AHD especially challenging, as delivery failures are generally costly (Visser & Savelsbergh, 2019; Waßmuth et al., 2023). Furthermore, to reduce the risk of delivery mishaps, service providers usually provide customers with options to select delivery time slots (Visser & Savelsbergh, 2019, p. 1). Service providers may thereby offer different time windows to different customers, e.g. based on their geographical delivery zone (Agatz et al., 2011; Bruck et al., 2020; Côté et al., 2024), or based on the customers' risk and value profiles (Beheshti et al., 2015, p. 404). Apart from reducing the risk of delivery failure, time-window offerings improve customer service, as suggested by Agatz et al. (2011), Bühler et al. (2016), Côté et al. (2024), and Zhan et al. (2021). Agatz et al. (2011) further identify the trade-off between service level and delivery cost as a key challenge for e-grocers operating in the AHD sector. They find that, while being more convenient for customers, narrow delivery time slots reduce routing efficiency (Agatz et al., 2011, p. 449).

Typically, as suggested by Campbell and Savelsbergh (2005) and Waßmuth et al. (2023), the fulfillment process of AHD providers can be summarized in three main stages. Figure 1 provides an overview of the process as suggested by Waßmuth et al. (2023, p. 804). In the first of three stages, denoted as the *order capture* stage, the order is placed by the customer (Waßmuth et al., 2023, p. 802). Customers can usually choose their service appointment to their own convenience, given the availabilities displayed to them by the service provider. Some providers tailor the set of time-window options to the customer's shopping history or delivery location for instance. Even more so, some providers nudge customers' choices by limiting time slot availability or implementing fees tied to specific time slots, all aimed at minimizing delivery costs (Bühler et al., 2016, p. 78) or increasing revenues (Waßmuth et al., 2023, p. 802). Once the company has confirmed the order, stage 2 consists of the *order assembly*. During this second stage, the provider takes action in preparing the delivery of the order, which typically includes order picking, sorting, and packaging (Waßmuth et al., 2023, p. 802). The final stage of *order delivery* entails transporting products or services to the customer's home within a specified timeframe. During this phase, the service provider assigns customers to vehicles and maps out delivery routes, a process that can be represented as a Vehicle Routing Problem (VRP) (Waßmuth et al., 2023, p. 802).

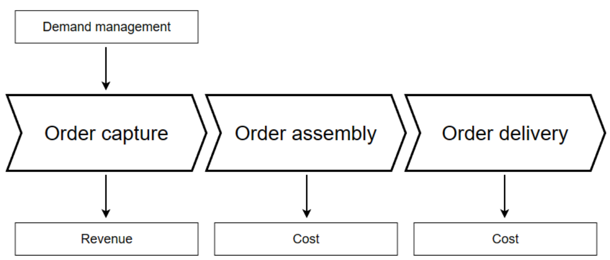

Figure 1: AHD process according to Waßmuth et al. (2023, p. 804)

Bruck et al. (2020) define AHS as "service delivery systems in which a supplying company and a customer agree on a time window during which the customer will be home and the service will be performed" (Bruck et al., 2020, p. 137). They further describe AHS as the combination of two stages. In the initial phase, called the booking process, customers reserve a service within the available time slots displayed by the company. In the service execution stage, the distribution company dispatches technicians to the customers' locations to carry out the services (Bruck et al., 2020, p. 140).

As opposed to the two stages described in Bruck et al. (2020, p. 140), it should be acknowledged that AHS providers can incorporate an intermediate *order assembly* stage as well, as it was described for AHD providers by Waßmuth et al. (2023, p. 802). This stage becomes crucial, particularly for service providers that also need to provide parts or goods as part of their services. For instance, companies that deliver furniture to their customers combine delivery and service, which is why there is a need for these companies to handle tasks such as supplying goods, managing warehouses, or loading trucks.

2.2 Demand Management in AHD

Demand management can be beneficial to the profits of companies in the AHD sector in two main ways (Waßmuth et al., 2023, p. 801). Waßmuth et al. (2023, p. 801) assert that demand management can amplify a company's revenue by means of prioritizing high-value customers or by expanding the company's customer base, enabled by better resource and capacity utilization. On the other hand, demand management can also boost company profits through cost savings that result from a more efficient order-to-delivery process (Waßmuth et al., 2023, p. 801). Bühler et al. (2016, p. 79) contend that the demand management of AHS providers is mainly motivated by cost reduction in the service delivery. The authors assert that, while revenue is generated within the first stage of order capture in AHS settings, service providers must anticipate delivery routing cost while handling demand management decisions in order to maximize profits (Bühler et al., 2016, p. 89).

Demand management is a field of research that is especially relevant to the first stage of order capture, as it aims to increase revenue. Generally speaking, ”demand management for AHD aims to generate customer demand, at the same time, shape it in a way that benefits the fulfillment process” (Waßmuth et al., 2023, p. 802). Waßmuth et al. (2023) provide an extensive review of the current literature revolving around demand management in the AHD sector. In that, the authors outline a planning framework and further distinguish three planning levels, encompassing the strategic, tactical and operational levels, as well as two levers, namely offering and pricing, through which the profitability of a company could be enhanced. Further, Bühler et al. (2016, p. 79) argue that incentives such as pricing that are strategically placed during the booking process, can occur in a static or dynamic way.

While strategic and tactical decisions occur before the first stage of order capture, the operational decision-making arises within that stage, in which customer demand is being communicated to and captured by the provider (Waßmuth et al., 2023, p. 804). For instance, Visser and Savelsbergh (2019) propose a two-stage stochastic programming formulation, which the authors refer to as Strategic Time Slot Management (STSM), while Waßmuth et al. (2023, pp. 808–809) classify the studied problem as tactical offering by the definition layed out in their literature review paper. Similar to the benefits of demand management addressed in Waßmuth et al. (2023), Visser and Savelsbergh (2019) contend that while STSM reduces the quantity of time slots displayed to customers to choose from, it generally enables more cost effective operations. Furthermore, Agatz et al. (2011) investigate a Time Slot Management Problem (TSMP) of tactical nature, in which a set of time-slot offerings needs to be selected for the different service regions, while ensuring an acceptable level of customer service. In their suggested method, time slots are geographically allocated before the demand of customers is known. Based on Waßmuth et al. (2023)'s classification, this thesis incorporates operational offering decisions, in which service options presented to customers during the order capture process must be determined (Waßmuth et al., 2023, p. 806).

2.3 Vehicle Routing Problems

In this subsection, an overview of VRPs and their derivatives that can be considered most relevant to this thesis, is provided. After discussing the relevance of the TSMP and demand management, which are related to the first stage of order capture in the AHD/AHS process, this thesis also focuses on the stage of order delivery, which usually entails solving a routing problem (Bühler et al., 2016; Waßmuth et al., 2023).

2.3.1 VRPs and their Derivatives

The Vehicle Routing Problem (VRP) made its first appearance in Dantzig and Ramser (1959) under the name of the Truck Dispatching Problem, described as a "generalization of the Traveling Salesman Problem (TSP)" (Dantzig & Ramser, 1959, p. 80).

In a simple VRP, a depot provides goods to a certain number of customers. A fleet is responsible for transporting these goods to geographically dispersed customer locations, given a set of constraints. These constraints may for instance include capacity constraints as well as time-window constraints. Each route of a VRP must start and end at a depot location, as depicted in the exemplary VRP diagram in Figure 2. In this diagram, every node symbolizes a customer location, interconnected by lines and arrows indicating the sequence of the corresponding route. Therewhile, the main objective of a VRP model is to minimize the total transportation cost, where the transportation cost can be assumed to be proportional to the vehicle's travel path. In other words, a shorter travel path will yield lower transportation cost through a reduction in fuel consumption and driver's working time (Zhang et al., 2022, pp. 195–197).

Since its first appearance in 1959, the VRP has been studied extensively in research, in many different variations of the initial problem. Reviewing relevant academic literature, Zhang et al. (2022) delve into the model description and solution methods of derivative types of the vehicle routing problem, including the Capacitated Vehicle Routing Problem (CVRP), Vehicle Routing Problem with Time Windows (VRPTW) and Dynamic Vehicle Routing Problem (DVRP).

Bruck et al. (2020, p. 140) assert that for most AHS problems, creating routing plans for technicians can be modeled as a VRPTW. Similarly, this thesis is specifically tailored to AHD/AHS providers which offer planned services to their customers and therefore deliver these services within given time windows. In particular, this work explores the impact of offering appointment-days to customers and investigates different strategies on how these offerings can best be assigned to customers. Considering this objective, the VRPTW and the Time Window Assignment Vehicle Routing Problem (TWAVRP) are of special interest for this thesis and are introduced in the following.

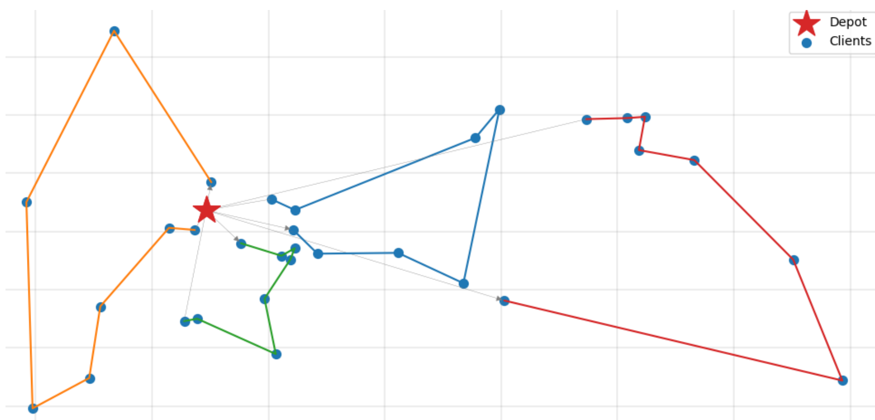

Figure 2: VRP diagram example

2.3.2 Vehicle Routing Problem with Time Windows

The Vehicle Routing Problem with Time Windows (VRPTW) extends the VRP by imposing a constraint that mandates the start of each service at customer locations within specified time windows. In academic literature, this problem is frequently differentiated into two categories: VRPTW with hard time windows (VRPHTW), where time windows must be strictly adhered to, and VRPTW with soft time windows (VRPSTW), allowing for violations of soft time windows with the payment of suitable penalties (Beheshti et al., 2015, p. 402). Further derivatives of the VRPTW include problems with multiple time windows and problems with fuzzy time windows (Beheshti et al., 2015, p. 403). Useful applications of the VRPTW encompass bank deliveries, postal deliveries and school bus routing, among others, as pointed out by Bräysy and Gendreau (2005a, p. 105).

In VRPTW, distributors typically want to maximize customer satisfaction while minimizing the total transportation cost at the same time. However, there is a trade-off between customer satisfaction and distribution cost, which causes the cost to increase when maximizing customer satisfaction, as pointed out by Beheshti et al. (2015, pp. 402–404). Thus, the authors design and solve a multi-objective VRPTW in order to arrive at an appropriate trade-off between distribution cost and customer satisfaction. In their case study, the distributor proposes a set of non-overlapping time windows and the customers prioritize these delivery time windows. Routing is then performed while incorporating customer preferences. In fact, this set of constraints makes the studied problem a multi-objective Vehicle Routing Problem with Multiple Prioritized Time Windows (VRPMPTW). Rasmussen et al. (2012) solve a Home Care Crew Scheduling Problem (HCCSP), which generalises the VRPTW in the setting of home care services. Similar to Beheshti et al. (2015), the authors implement a multi-criteria objective function, minimizing the overall operational costs while maximizing the service level. This reflects the importance of maintaining a certain service level while optimizing for cost in delivery or transportation in a vehicle routing problem. Highlighting the relevance of VRPTW and derivatives of this VRP type in the home care and home health care sector, Liu et al. (2019) address a home-caregiver scheduling and routing problem with stochastic travel and service times, which can, according to the authors, be viewed as a special Multi-Depot Vehicle Routing Problem with Time Windows (MDVRPTW) with stochastic travel and service times. In this case, multi depot refers to the different home locations of the caregivers, given that each of their individual routes starts and ends at their personal home location.

2.3.3 Time Window Assignment Vehicle Routing Problem

In practical scenarios, retailers heavily rely on pre-determined time windows for operational tasks like inventory management and personnel scheduling. The Time Window Assignment Vehicle Routing Problem (TWAVRP) addresses this need through a two-stage approach. Initially, time windows are assigned to each customer before their demand is known. Subsequently, after customer demand is revealed for each day of the time period, vehicle routes are constructed while adhering to the assigned time windows from the first stage (Spliet & Gabor, 2015, p. 379).

Spliet and Gabor (2015) solve a variant of TWAVRP referred to as the discrete TWAVRP (DT-

WAVRP). Unlike TWAVRP, DTWAVRP involves selecting a single time window from a finite set of candidate time windows for each customer, instead of using continuous time windows. Additionally, the authors propose approximating the probability distributions of customer demands with a finite set of possible demand scenarios. They highlight that, particularly when considering multiple scenarios, DTWAVRP essentially entails solving several instances of the Vehicle Routing Problem with Time Windows (VRPTW), each corresponding to a distinct scenario and interconnected by the choice of time windows (Spliet & Gabor, 2015, pp. 379–380). Building on Spliet and Gabor (2015), a recent study carried out by Côté et al. (2024) tackles another variant of the TWAVRP, which the authors define as a Stochastic Multi-Period Time Window Assignment Vehicle Routing Problem (SMTWAVRP). The study is based on a real case provided by a Canadian retailer that sells and delivers furniture and appliances. Côté et al. (2024) thereby introduce a novel problem by introducing uncertainty for the number of customers, their locations, demands and service times. In the first stage of their model, the number of visits and time windows offered per geographical delivery zone is optimized. After customers have expressed their individual time-window preferences, the second stage takes the first stage solution as an input and constructs optimized routes to visit the customers while considering customer preferences. Similarly, Zhan et al. (2021) solve an integrated routing and appointment scheduling (RAS) problem with stochastic service times, motivated by the practices of home services. Unlike Spliet and Gabor (2015) and Côté et al. (2024), the authors restrict their model to time points instead of time windows. Moreover, customers' preferences are not represented in Zhan et al. (2021)'s model, as opposed to Côté et al. (2024).

2.4 Solution Approaches to VRPs

There are numerous different approaches to solve vehicle routing problems. In the following, some of the methods that are relevant to this study are presented, namely clustering (2.3.1), a priori optimization (2.3.2) and different solving algorithms (2.3.3) .

2.4.1 Clustering in Vehicle Routing

In the context of solving Vehicle Routing Problems, this part of the literature review revolves around the so-called ”cluster-first route-second” method, in which customers are first grouped into feasible clusters (cluster-first), before establishing efficient routes for each of the clusters (Korayem et al., 2015, p. 1). In a two-stage solution approach to a VRP, clustering is often used in the first stage.

Clustering has been found to be effective in managing the burden of mathematical complexity in solving a VRP (Nallusamy et al., 2009, p. 131). Furthermore, Wang et al. (2015, p. 1427) claim both the clustering of customer and the optimization of vehicle routing to be critical for the successful implementation of VRP into large-scale logistics networks. Customer clustering could thereby not only improve the logistics system efficiency, but also reduce the operational cost (Wang et al., 2015, p. 1428). Further, Rasmussen et al. (2012, p. 599) argue that while clustering may compromise optimality, it enables the resolution of larger instances of problems. The authors further posit that run time decreases significantly by adopting clustering before solving the VRPTW.

The business case of a Canadian retailer presented in Côté et al. (2024) provides real-world evidence for adopting a location-based clustering approach in practical settings. Specifically, the retailer promotes the strategy of consolidating deliveries to customers residing in the same area on the same delivery day, thereby minimizing unnecessary travel distances. This partitioning of customers further allows to assign different time-window options to different sets of customers and control for number of visits for different delivery zones (Côté et al., 2024). Similarly, Nallusamy et al. (2009) use a K-means clustering algorithm to group 180 cities to 6 different vehicles in an optimal manner. Bührmann and Bruwer (2021, p. 39) acknowledge the usefulness of clustering in vehicle routing in assisting managers to group and assign customers to vehicles. The authors compare a K-medoids and a K-means clustering approach to solving a the CVRP without clustering and find that the K-medoids algorithm outperforms the more frequently used K-means approach in solving a CVRP. However, Bührmann and Bruwer (2021, p. 39) find both clustering methods to provide a more logical grouping of customer locations compared to solving the CVRP without clustering inputs. While the authors acknowledge several practical advantages of clustering in VRP such as facilitated planning and assignment of resources or the allocation of newly added customers in the distribution network, the study shows that the algorithm used to optimize the routes performs better in terms of minimizing cost and distance travelled when running the CVRP without prior clustering of customers. In fact, the authors find that the travel cost increases with the number of clusters. The reason for that is that clustering, while being beneficial from an operational point of view, adds an extra constraint to the method and reduces the solution space, thereby leading to a decrease in performance in terms of cost and distance traveled (Bührmann & Bruwer, 2021, p. 39). Korayem et al. (2015) describe the clustering phase as a Capacitated Clustering Problem (CCP), since clusters have to maintain the vehicle capacity constraint, with each vehicle serving one cluster, which is why the authors consider the demand of each customer when clustering. Korayem et al. (2015) employ a K-means clustering algorithm and acknowledge that it suffers from the possibility of falling into local minima (Korayem et al., 2015, p. 2). To avoid forming infeasible clusters, the authors develop two new heuristics. Bruck et al. (2020) develop a Decision Support System (DSS) to support an Italian multi-utility company facing an AHS problem. The first module of the DSS implements a Mixed-Integer Linear Program (MILP) with distance inputs received from an Open Source Routing Machine (OSRM) to generate geographical clusters. Based on the formed clusters, appointment scheduling and routing are optimized in stages 2 and 3 of the DSS.

When grouping customers into clusters, customers should share common features such as geospatial location, as discussed previously, or demand (Wang et al., 2015, p. 1428). In similar terms, there are other parameters than the intuitive location by which customers are grouped in academic literature revolving around vehicle routing. In particular, in the case presented by Beheshti et al. (2015), the underlying company partitions its customer landscape based on the customers' value and risk profile. In Rasmussen et al. (2012), a range of visit clustering schemes are devised for addressing the HCCSP, each grounded in soft preference constraints.

2.4.2 A Priori Optimization

In an a priori framework, the set of potential clients remains fixed and known, yet each day requires serving only a subset of these clients. In the a priori optimization problem, the "master solution" specifies the solution to any potential instance such that for any arising instance, the customers are visited in the same order as determined by the master solution, which encompasses the entire potential client set. This way, a minimal additional travel distance, compared to if a delivery person had followed the optimal tour specifically tailored to that subset of clients, is ensured (Schalekamp, 2007, p. 1).

Similarly, Campbell and Thomas (2008) define an a priori, or pre-planned route as "a route which specifies an ordering of all possible customers that a particular driver may need to visit" (Campbell & Thomas, 2008, p. 1). The authors assert that many delivery companies have long used a priori routes to overcome the difficulties caused by the fact that only a subset of their customers need to be visited each day. Visser and Savelsbergh (2019, p. 2) support that strategy by claiming that maintaining partial routes and schedules facilitates adjustments required by the occurrence of demand variations, and that it further enables the control of the number of required delivery vehicles. Further, the authors observe that, while online grocery retailers face a challenge in handling variability in daily demand, their customers also exhibit recurring ordering patterns that express in favorite time slots and delivery days. As an example, Visser and Savelsbergh (2019) focus on time slot management for online grocery retailers and specifically motivate their study by the example of Picnic, an online grocery retailer that is based in The Netherlands and operates in all major city centers of the country (picnic.app/nl/). Picnic strategically plans delivery routes for each day of the week in advance, covering all customer locations, and allocates time slots to these routes before orders are placed. During the ordering process, available time slots are managed based on these pre-established routes. The customers' selected time slots are incorporated into the corresponding delivery route associated with the pre-planned route, on which the customers were assigned to (Visser & Savelsbergh, 2019, pp. 2–3). The authors contend time slot management to be fairly easy when adopting such an a priori route approach, following pre-planned routes and skipping locations according to the planned schedule (Visser & Savelsbergh, 2019, p. 3).

Further, Visser and Savelsbergh (2019) identify several more advantages of designing a priori routes. First, a priori routes integrate customer order patterns, particularly beneficial in urban areas, enhancing operational efficiency. Second, the elimination (re)optimization after cut-off times enables extended cut-off times, offering increased customer convenience and earlier fulfillment center operations. Moreover, a priori routes foster delivery route consistency, enabling drivers to become familiar with routes, thus improving customer service, especially in densely populated city centers (Visser & Savelsbergh, 2019, p. 3).

There are two general approaches that can be adopted when using a priori routes: the *reoptimization approach* and the *skip approach*. In the *reoptimization approach*, delivery routes are reoptimized after customer demand and a subsequent schedule is known. However, in many cases, companies may not have the resources to reoptimize on each instance or service day, especially at larger scale. Nevertheless, these a priori routes provide a good starting point for reoptimization, in case time and resource availabilities allow for it (Bertsimas et al., 1990; Campbell & Thomas, 2008). Often, companies cannot follow this reoptimization approach, which would intuitively appear more ideal, as pointed out by Bertsimas et al. (1990, p. 1020), due to the approach being too expensive and time consuming, a lack of advance information, or because the company prioritizes regularity of service (Jaillet, Patrick, 1988, p. 929). This drawback, in turn, gives ground to the second approach used in a priori optimization, referred to as the *skip approach*. The skip approach overcomes the mentioned limitations of reoptimizing delivery routes. It involves pre-planning a tour through all potential customer locations, regardless of daily demand variations. During execution, the vehicle follows this predetermined route and its subsequent sequence of customer locations, skipping customers who do not require service on a particular day (Bertsimas et al., 1990; Campbell & Thomas, 2008; Jaillet, Patrick, 1988; Visser & Savelsbergh, 2019) .

2.4.3 Solving Algorithms

Solving VRPs is referred to as NP-hard, a classification originating in computational complexity theory, in brief, meaning that no algorithm is able to solve these problems in polynomial time (Molina et al., 2020, p. 2). Accordingly, this classification is especially true for extensions of the basic VRP, such as the CVRP or VRPTW, which add complexity to the problem. Hence, finding optimal solutions to large-scale problems is not practicable, which justifies the usage of classical heuristics and sophisticated metaheuristics in order to find solutions of good quality when dealing with VRPs (Iswari & Asih, 2018; Molina et al., 2020). Bräysy and Gendreau (2005a, p. 105) highlight the importance of heuristics, asserting their capability to produce high-quality solutions in limited time. They further constate a trade-off between the run time of a heuristic and its solution quality. Hence, also considering the high complexity level of problems such as the VRPTW and its wide applicability, the authors express the importance of finding a compromise in this trade-off, such that good quality solutions could be produced within a reasonable time scope (Bräysy & Gendreau, 2005a, p. 105).

Classical Heuristics are characterized by their focused exploration of the solution space, often yielding high-quality results in reasonable time frames. Additionally, these methods are adaptable to accommodate various constraints present in real-world scenarios, making them practical choices for commercial applications (Laporte et al., 2000, p. 286). Classical heuristics can typically be divided into *Constructive Heuristics* and *Improvement Heuristics* (Laporte, 2007, p. 814). The first, also referred to as route construction heuristics, sequentially select nodes or arcs until a viable solution is achieved. Nodes are thereby typically chosen according to a cost minimization criterion, while ensuring compliance with constraints such as vehicle capacity and time windows. Zhan et al. (2021, p. 198) distinguish between the two-stage method, scanning algorithm, C-W saving algorithm, nearest neighbor method and recently inserted method, as different classes of constructive heuristics. Improvement heuristics iteratively improve a solution by exploring neighboring solutions to the problem (Bräysy & Gendreau, 2005a, p. 109). In Zhan et al. (2021, p. 198)'s algorithm classification, improved heuristics include the k-opt algorithm and λ-interchange algorithm.

Metaheuristics, unlike classical heuristics, "allow the exploration of the solution space beyond the first local minimum encountered" (Laporte, 2007, p. 815). In other words, the emphasis of such methods lies in the comprehensive exploration of promising regions within the solution space (Laporte et al., 2000, p. 286). In that, the algorithms even tolerate non-improving and infeasible intermediate solutions during the search process. At the same time, they incorporate techniques derived from classical construction and improvement heuristics. As a result, the solution quality obtained by metaheuristics is usually superior to the solution quality obtained with classical heuristics (Bräysy & Gendreau, 2005b; Laporte, 2007). However, in the context of the trade-off between run time and solution quality (Bräysy & Gendreau, 2005a, p. 105), this higher solution quality of metaheuristics also comes at the price of an increased computing time (Laporte et al., 2000, p. 286). Laporte (2007) classifies metaheuristics into three categories. The first, referred to as *local search* includes Tabu Search (TS) and Adaptive Large Neighbourhood Search (ALNS). *Population search* or *population-based methods* encompass metaheuristic methods such as the Genetic Algorithm (GA) and Particle Swarm Optimization (PSO) and are characterized by solutions moving between populations with specific evaluation processes. In particular, a GA uses genetic operations such as mutation and crossover, while PSO adopts elements of swarm behavior, like updating position and velocity in the evaluation (Iswari & Asih, 2018; Laporte, 2007). The third class Laporte (2007) addresses is referred to as *Learning Mechanisms* and includes neural networks and ant colony optimization, which attempt to replicate ants' trail-marking behavior, thereby favoring the emergence of efficient paths.

Many advancements have been made towards optimizing and extending some of these metaheuristics. For instance, Thangiah et al. (1991) have developed a two-moduled GA named "GIDEON", in which the first module performs genetic sectoring, while the second module consists of a local route optimization. Furthermore, Korayem et al. (2015) study a Grey Wolf Optimizer (GWO), a new metaheuristic which draws inspiration from the leadership hierarchy and hunting mechanism observed in grey wolves. The GWO, extended by K-means (K-GWO), thereby showed better accuracy in solving a CVRP, compared to GAs and PSO. Similarly, Vidal et al. (2013) propose a Hybrid Genetic Search (HGS) algorithm with adaptive diversity management which addresses a wide range of VRPTW, thereby outperforming other existing approaches. It combines exploratory strengths of genetic algorithms with effective local search-based improvement methods and diversity management (Vidal et al., 2013, p. 477). This HGS algorithm is of particular interest for this thesis, as it serves as the foundation for the route-solving algorithm used.

3 A Case study: A Dutch Car Services Provider

3.1 Problem Description

This thesis introduces a problem that can be projected onto many companies operating in the AHD/AHS sector. In particular, the study specifically addresses companies in the sector that provide planned services (AHS). In other words, subjects include companies that provide services that are characterized by a certain regularity in their occurrence and planning beforehand.

For the ease of applicability and interpretability, the presented solution approaches are constructed around a case of a Dutch car services provider. While the company exists, the data used in this thesis is synthetically generated, of which the generation process is presented in detail in section 5.3. The company's main value proposition is the changing of car tires at customers' home locations. With that, the provider has a rather static customer landscape, where customers usually need to be visited twice a year, for putting on winter or summer tires accordingly. Nevertheless, some new customer dynamics might arise. This thesis, however, focuses on the plannable side of the company's operations, i.e. existing, recurring customers, and therefore neglects the arrival of new customers.

This thesis investigates the impact of employing a priori routes in combination with providing appointment-day offerings on the planning complexity, route efficiency and customer service of the company's operations. In particular, different strategies, including different appointment-day assignment and subsequent routing strategies are evaluated and assessed by the primary metric of total expected distance traveled, further addressed in subsection 5.1.

3.2 Assumptions

First, it is assumed that the company operates a fleet of m vehicles, each with a homogeneous capacity c, which represents the number of customers a single vehicle can visit, taking into account the number of tires it can carry. Hence, the company's final routes encompass c customer locations in their sequence. The a priori route for a group of customers, however, encompasses a multiple of these c customers. To be precise, one a priori route includes $c \times T$ customers, where T is the number of appointment-day options given to each customer. The company is assumed to operate out of one depot location, from which the vehicles are sent out on their routes to customers and come back to after completing their route. Further, customer locations are known by the company and static, meaning that no new customers join and no customers drop out of the customer network.

Moreover, a planning period of *n* days is observed, during which each customer has to be served once and only once. Further, it is assumed that each customer in the considered customer network requires a visit within the planning period. Additionally, customers are assumed to be indifferent between appointment-day offerings. In other words, it is assumed that, given the appointment-day options displayed to a customer, they are equally likely to choose either of the days for their appointment. Based on these customer decisions, this thesis considers every possible combination of customers choosing from the options displayed to them, which is referred to as *customer decision scenarios*, but evaluates the route performance on a subset of S randomly drawn scenarios, to allow for computation in a reasonable amount of time.

Further, it is assumed that one vehicle has to visit c customers within one appointment-day. Hence, the limiting factor in the planning of vehicle routes is the truck capacity c, not the time dimension, which is why a CVRP is solved without considering day-specific time windows or any other temporal constraints. Consequently, only the sequence in which customers are visited is considered in the final route solution, rather than the specific arrival or departure times at each customer location. Once the routing decisions are made within the company's process, customers are informed about a more narrow, day-specific time frame, in which they can expect to be visited. This last step, however, goes beyond the final routing decision and is therefore not subject to this thesis. Similarly, this thesis neglects the stage of *order assembly*, introduced by (Waßmuth et al., 2023, p. 802). The reason for that is that this study solely aims at optimizing for the decisions of appointment-day assignments and the subsequent routing decision in a two-stage framework.

Given the provided problem description and assumptions, the problem in this study can be described as a DTWAVRP in the context of a AHS for ordinary, i.e. recurring, plannable, services.

4 Methodology

As a first element of this section, the general design of the decision making process studied in this thesis is explained. Next, a detailed overview of the different strategies is provided, followed by a subsection, which introduces the upper and lower bound methods. Finally, the CVRP- & TSP-solving method employed in this study is presented.

4.1 General Design of the Decision Making Process

The solution approaches in this thesis revolve around a decision making process that is motivated by Waßmuth et al. (2023, p. 802), Campbell and Savelsbergh (2005, p. 2) and Bruck et al. (2020, p. 140). The high-level decision-making sequence that can be assumed for this study provides the frame to the proposed strategies. It is visualized in Figure 3 and is detailed in the following.

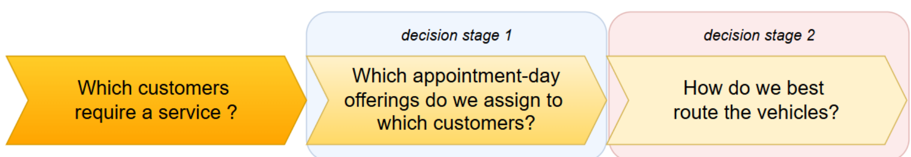

Figure 3: AHS decision making process

4.1.1 Decision Stage 1: Assigning Appointment-Day Offerings to Customers

The provider has to first decide which appointment-days should be offered to which group of customers. Prior to this, customer needs are known by the provider because it usually follows a strict pattern. This stage involves partitioning customers into groups, motivated by studies such as Bührmann and Bruwer (2021), Côté et al. (2024), and Nallusamy et al. (2009) and many more, mentioned in section 2.3.1. The partitioning of customers, further explained in ”Stage 1 Approaches”, divides customers into groups, in each of which a customer is assigned to the same set of T appointment-day offerings.

4.1.2 Decision Stage 2: Establishing Final Routes

During the second stage, after customers chose their appointment day, the provider's decision consists of designing the final route of a specific vehicle on a given appointment-day. In other words, the route that will be followed by the technician to visit customers on a given day has to be finalized, based on the partitioning of decision stage 1 and the customer decisions. In particular, these final routes can either be built using the a priori routes' customer sequence determined in stage 1, or by reoptimizing and thereby solving a new routing problem with a TSP, two methods which are further detailed in ”Stage 2 Approaches”.

4.2 Strategy Specifications

Stage 1	Stage 2	Abbreviation
CVRP	Skip	*CVRP-Skip*
KmeansTSP	Skip	*KmeansTSP-Skip*
CVRP	TSP	*CVRP-TSP*
Kmeans	TSP	*Kmeans-TSP*

Table 1: Two-stage strategies overview

Based on the decision making process described in the previous section, 4 different strategies are explored in this study, considering two different methods for stage 1 and two methods for stage 2 of the decision making process. An overview of the different strategy characteristics is provided by Figure 4 and Table 1. Appendix B provides pseudo-algorithms for the high-level methodology (Algorithm 1) as well as for each decision stage (Algorithms 2 & 3). In the following, each strategy is explained in depth.

4.2.1 Stage 1 Approaches

The first stage 1 approach, which is denoted as *stage1-CVRP* and depicted in Figure 5 (a), partitions customers by solving a CVRP, comprising all the customers that need to be visited over the planning period of n days, similar to the route-based clustering approach adopted in Becker (2023). In solving this CVRP, the capacity parameter, i.e., how many customers are assigned to one a priori route, is defined by $c \times T$, since customers with the same appointment-day offerings should be grouped together. This approach is employed in strategies *CVRP-Skip* and *CVRP-TSP*, which differ

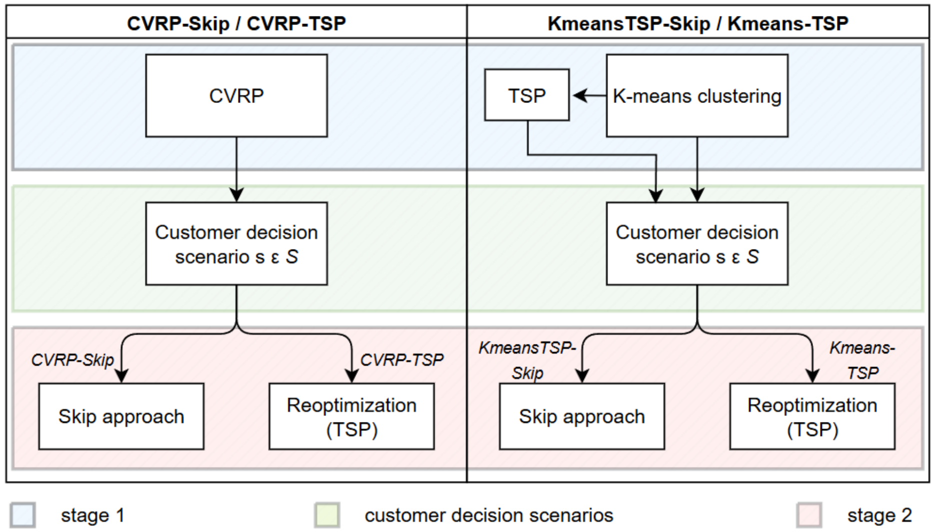

Figure 4: Two-stage strategies visualized

from each other during stage 2. The resulting routes are then fixed to a priori routes, on which customers have the same appointment-day offerings assigned to. These a priori routes consequently represent the partitioning of customers and give ground to the customer decision scenarios. In the fictive scenario of Figure 5, there are 16 customers, represented by the nodes, and one depot. As described, on each of the a priori routes, shown by the arrows that connect customers and the depot, customers are offered the same appointment-day, as depicted by the different colors in the nodes.

The second approach in this decision stage is denoted as *stage1-Kmeans*, visualized in Figure 5 (b), and involves solving a CCP, as suggested by Korayem et al. (2015), by applying a clustering algorithm that partitions the geospatial customer landscape into groups of size $c \times T$. The advantages of partitioning customers were detailed in section 2.3.1 about clustering literature in the field of vehicle routing, where sufficient support is provided for using this method in this first stage of assigning appointment-day options to customers (Bührmann & Bruwer, 2021; Côté et al., 2024; Nallusamy et al., 2009; Wang et al., 2015). This thesis does not experiment with different clustering algorithms but applies the K-means algorithm due to its widespread use, ease of implementation, and proven effectiveness in producing good results (Bührmann & Bruwer, 2021; Korayem et al., 2015; Nallusamy et al., 2009). More specifically, the K-means algorithm employed in this thesis uses the geospatial information of customer locations and uses the Haversine distance calculation to create clusters. The Haversine distance determines "the shortest distance between two points on a sphere using their latitudes and longitudes measured along the surface" (GeeksforGeeks, 2022). The formula for the Haversine distance calculation, which is used for any distance calculation throughout this study, is provided in Appendix A. Please refer to equation (1) for determining the number of clusters K in the application of K-means clustering in this study, where c stands for the vehicle capacity and T denotes the number of appointment-days offered to each customer. This *stage1-Kmeans* approach is adopted by the strategies *KmeansTSP-Skip* and *Kmeans-TSP*. While stage 1 ends here for the *Kmeans-TSP* strategy, a TSP needs to be solved for each resulting cluster for strategy *KmeansTSP-Skip*, making it *stage1-KmeansTSP*. The resulting routes are then, like in the *stage1-CVRP* method, fixed to a priori routes, on which each customer gets the same appointment-day options displayed for each route respectively. Fixing a priori routes is not required for strategy *Kmeans-TSP* since the first stage does not solve a routing problem.

$$\text{number of clusters K} = \frac{\text{Number of customers}}{c \times T} \tag{1}$$

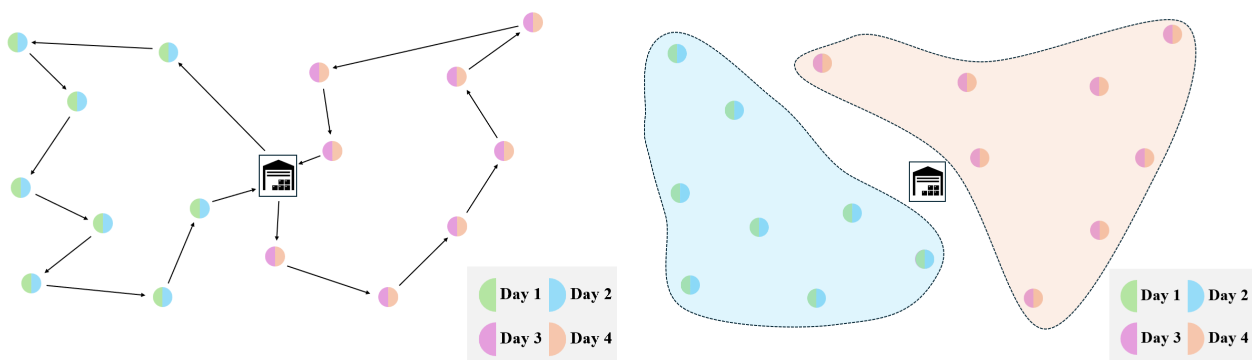

(a) Stage1-CVRP (b) Stage1-Kmeans

Figure 5: Comparison of stage 1 approaches

4.2.2 Customer Decision Scenarios

As mentioned before, in this study, each customer is given the choice to decide between T appointment days. The customer decision scenarios, previously introduced in section 3.2, represent various combinations of customers' appointment-day choices. In each of the scenarios, a customer chooses exactly one of the offered appointment days. As a consequence, for each group of customer resulting from stage 1, each customer decision scenario $s \in S$ is considered as an input for stage 2. Hence, for each strategy, stage 2 solves one routing problem for each customer decision scenario respectively, an approach that is motivated by Spliet and Gabor (2015, pp. 379–380).

4.2.3 Stage 2 Approaches

In decision stage 2, this study design differentiates between two approaches, introduced in section 2.3.2 about a priori optimization.

In the first approach, denoted as *stage2-Skip* approach, based on customers' choices, customers are visited on the chosen day of the appointment and skipped otherwise. Hence, the final route that results from this approach follows the visiting sequence fixed by its respective a priori route, which is the outcome of *stage1-CVRP* and *stage1-Kmeans-TSP* respectively, and skip customers that have to be visited on the other day accordingly. This *stage2-Skip* method is employed by *CVRP-Skip* and *KmeansTSP-Skip*. The approach is visualized in Figure 6 (a) for the *CVRP-Skip* strategy. Figure 6 is based on the *stage1-CVRP* approach depicted in Figure 5 (a). The grey arrows in Figure 6 (a) represent the remaining, unused paths of the a priori routes fixed in Figure 5 (a), while the colored arrows show the actual final routes after skipping customers respectively.

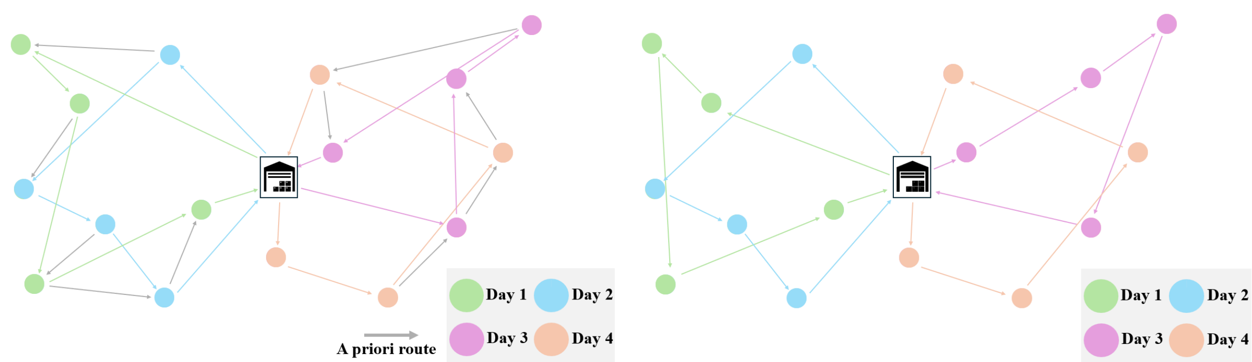

(a) Stage2-Skip approach

(b) Stage2-TSP approach

Figure 6: Comparison of stage 2 approaches

The second stage2-approach, denoted as *stage2-TSP* takes on a reoptimization strategy that is visualized in Figure 6 (b). For the *CVRP-TSP* strategy, the a priori routes from *stage1-CVRP* are reoptimized after customers have decided on their appointment-day by solving a new TSP for the customers assigned to each day and truck. In the case of strategy *Kmeans-TSP*, the final route for each appointment day and truck is determined by solving a TSP based on the clusters from *stage1-Kmeans* and the customer decisions. In other words, since choosing an appointment-day allows to further assign customers to day/vehicle-specific groups, a new TSP for this specific group of customers can be solved, which then gives the final route for a given truck on a given day. This strategy experiments with the observation made by Bertsimas et al. (1990) and Campbell and Thomas (2008), contending that a priori routes offer a good starting point for reoptimization, given the necessary resource availability, i.e. mainly computation and time.

4.3 Upper- and Lower Bound

To better understand the results of different strategies, they are presented alongside one upper- and one lower bound. While the upper bound is expected to show the worst performance at a high customer service level, the lower bound is expected to display the best performance in minimizing the total distance traveled at a low customer service level, accordingly.

4.3.1 Upper Bound

As an upper bound with the highest expected distance, a *free-choice* strategy is adopted, in which each customer in the full customer landscape for the full planning period of n days is given free choice over their appointment-day. In other words, for this method, $T = n$. After having grouped customers to days, a CVRP is solved for each day with the customers who chose this appointment day. In each of these CVRPs, the number of vehicles per day m_t is given by equation (2):

$$\text{number of vehicles } m_t \text{ in Upper Bound CVRP} = \frac{\text{Number of customers}_t}{n \times c} \quad \text{for } t \in \{1, 2, \ldots, n\} \tag{2}$$

4.3.2 Lower Bound

As a lower bound with the lowest expected distance, a *no-choice* strategy is adopted, in which customers have no decision power over the appointment-day they are assigned to. For this, a CVRP is solved for the full customer landscape and planning period of n days. The capacity, i.e., the number of vehicles needed, for solving each CVRP is given by equation (3). The total number of vehicles available is, however, used over the full span of n days, which is why the number of vehicles needed per appointment-day is again given by equation (2).

$$\text{number of vehicles } m \text{ in Lower Bound CVRP} = \frac{\text{Number of customers}}{c} \tag{3}$$

4.4 CVRP & TSP Solving Method

In order to solve the routing problems for each strategy in stages 1 and 2 respectively, this study will make use of a recently developed open-source Python library named *PyVRP* (Wouda et al., 2024). With this package, Wouda et al. (2024) aim to provide flexible implementation and a state of the art VRP solver for researchers and practitioners (Wouda et al., 2024, p. 1). The solver currently supports two VRP variants, the CVRP and VRPTW. Its implementation is based on the open-source Hybrid Genetic Search (HGS) - CVRP implementation, published by Vidal (2022) and extends it by supporting VRPTW and redesigning it as a highly customizable Python package. The package uses a variant of Vidal et al. (2013)'s HGS algorithm and consists of a genetic algorithm, a population, and a local search improvement method (Wouda et al., 2024).

5 Experimental Design

This section is divided into three parts. The first part details the method that is used to evaluate the performance of strategies in this study. Next, expectations about strategy-specific performances are addressed in the light of the efficiency, complexity, customer service level trade-off. Finally, a description of the data and data generation process is provided.

5.1 Strategy Performance Evaluation

This study evaluates and finally compares different strategies in order to gain insights about each method's performance, while drawing implications about efficiency, complexity and customer service level. In that, it should be specified how the performance of a strategy is measured, in order to evaluate it next to other methods in a comparative setting. This thesis takes the expected value of the traveled distance, which is the objective that is to be minimized, for each of the strategies. In particular, this entails taking the expected total distance of routes over the different customer decision scenarios, from which the expected value over all different customer location datasets is taken. For each of the strategies, this further yields the expected value of total distance traveled for each of the five studied datacases, which are further specified in section 5.3. The mathematical representation for the expected value for a datacase is provided by equation (4), in which n_I represents the number of datasets $i \in I$ included in the datacase. Further, n_S represents the number of customer decision scenarios $s \in S$ for each of the datasets, while d stands for the distance of a route $r \in R$. Taking the expected value over the five datacases gives the overall solution over all provinces and dataset sizes. This approach of taking the expected value follows the reasoning suggested by Schalekamp (2007) , contending that in an a priori framework, the goal was to obtain a good average solution and with that to optimize for the expected value of the objective.

$$E[D] = \frac{1}{n_I} \sum_{i \in I} \left(\frac{1}{n_S} \sum_{s \in S} \left(\sum_{r \in R} d_r \right) \right) \tag{4}$$

5.2 Results Expectations

The comparative analysis of this thesis builds on real-world assumptions and implications about the individual strategy's planning *complexity*, route *efficiency* and level of *customer service*. In other words, the different approaches chosen can vary in different dimensions, of which three will be explored closely in this study, allowing to draw implications about different advantages and disadvantages in the implementation of such strategies for specifically AHS companies.

In terms of *complexity* in determining final routes, the lower bound can be considered the least complex because it only consists of one stage in which one larger CVRP is solved, without any customer decisions involved. The upper bound is slightly more complex but still relatively simple, as it only requires solving one smaller CVRP for each day of the planning period, after customers have made their appointment-day decisions. At a similar level of complexity, the *CVRP-Skip* is considered relatively simple. While it solves the same CVRP as in the lower bound solution, a small layer of complexity is added by adopting a skip approach as a second stage after incorporating customer's decisions. On the third level of complexity, the *KmeansTSP-Skip* strategy can be allocated, since it comprises two components in stage 1, solving a K-means algorithm with rather low complexity and TSPs for each resulting cluster. Hence, it is reasonable to deem it as slightly more complex than the *CVRP-Skip* strategy. While both reoptimization strategies *CVRP-TSP* and *Kmeans-TSP* are characterized by a higher complexity than the previously mentioned methods, the *CVRP-TSP* is considered the most complex, because solving a CVRP is assumed to be more computationally expensive than solving a K-means clustering algorithm.

In terms of *efficiency*, optimality is expected to be compromised in using strategies *CVRP-Skip, CVRP-TSP, KmeansTSP-Skip* and *Kmeans-TSP*, due to the clustering and overall partitioning done in stage 1. However, these two-stage solution approaches may enable solving problems at larger scale in a reasonable run time (Rasmussen et al., 2012, p. 599). Hence, building on Bührmann and Bruwer (2021, p. 39), it can be expected that the lower bound outperforms the other methods in terms of total expected distance traveled. At the same time, the upper bound is expected to yield the highest total distance, because customers are given more choices, thereby reducing the solution space of potential routes. Furthermore, the strategies *CVRP-TSP* and *Kmeans-TSP* can be expected to perform better than their respective Skip-approach variant, *CVRP-Skip* and *KmeansTSP-Skip*, respectively, since reoptimizing is considered to be the more ideal solution approach (Bertsimas et al., 1990, p. 1020).

The lower bound is considered providing the highest level of *customer service*, given the choice between the full planning period given to customers. This reasoning builds on the findings of Agatz et al. (2011), Bühler et al. (2016), Côté et al. (2024), and Zhan et al. (2021), contending that time-window offerings improve customer service. The strategies *CVRP-Skip, CVRP-TSP, KmeansTSP-Skip* and *Kmeans-TSP* therefore fulfill a lower, but comparable threshold of customer service. The lowest level of customer service is consequently provided by the upper bound, offering customers no decision power over their appointment scheduling.

This rationale of how strategies balance the three dimensions is visualized in Figure 7, where the dimensions are evaluated on a scale from 1 to 5. The higher the score, the more the characteristic is magnified for the respective strategy or bound.

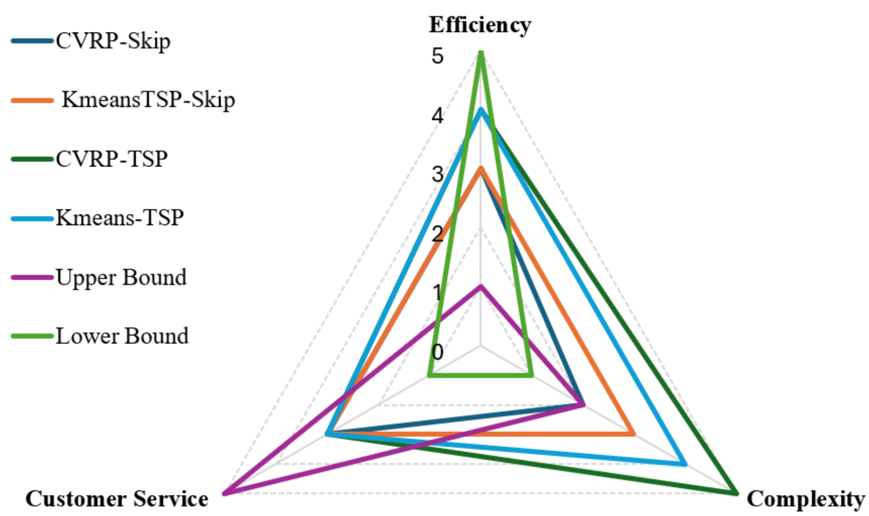

Figure 7: Balance of complexity, customer service & efficiency by strategy

5.3 Data

The goal of this section is to explain how synthetic data of customer locations is generated for this thesis. The overall objective of this task is to generate different scenarios of customer locations that are as realistic as possible, which will be done by assigning population density related weights to potential customer locations when randomly generating customer locations. Inspired by Bührmann and Bruwer (2021)'s approach, using different sets of customer locations allows to reduce bias and achieve overall more representative performance values when comparing the different solution approaches. In that, customer locations within The Netherlands are generated in different provinces, where each province has one designated depot location that is allocated by searching the internet for appropriate depot locations. An appropriate location, in this context, is a location which is situated in an industrial park within the respective province. Other requirements to the depot location are not set, which is why the designation of a depot simply follows an educated guess. It is therefore assumed that each province has its own depot from which technicians are sent out to customers, who are located in the same province as the depot. Figure 14 in Appendix A illustrates a sample customer landscape for the province of Utrecht, featuring 32 customer locations, represented by the blue dots, spread across a 4-day planning period, considering a single delivery vehicle with a capacity of $c = 8$ customers, along with a depot location, represented by the red dot. The process adopted in this thesis to generate synthetic datasets of customer locations is described in the following.

As a first step, postal codes are scraped from nld.postcodebase.com, resulting in 461,947 Dutch postal codes, in which the first 4 characters are numbers, followed by two letters. These 4-digit postal codes that are used in The Netherlands, are from here on referred to as PC4 codes. A PC4 code thereby encompasses several of these 6-character zip codes. As a second step, as the objective of the data generation process is to generate customer locations in a realistic way, the overarching PC4 postal codes and their population are retrieved from a dataset publicly provided by Statistics Netherlands (CBS, n.d.). Next, for each PC4, the population density is calculated as follows:

$$PC4Density_p = \frac{PC4Population_p}{\sum\limits_{p \in P} PC4Population_p} \quad \text{for } p \in P \tag{5}$$

In equation (5), one PC4 code is denoted as p, element in the set P of all PC4 codes. Once the population density is calculated, PC4 codes in the amount of desired customer locations are randomly drawn with the population density measure as an assigned weight, i.e. probability, from the subset of PC4 codes for the respective province. In a next step, for each selected PC4 code, one corresponding 6-character zip code is randomly chosen with equal weight distribution. This zip code serves as the representative location for one customer. Finally, latitude and longitude coordinates from a publicly available data file are appended for each of the zip codes in the generated dataset (Yurchak & Casares, 2021).

As a result of the data generation process, 5 different *datacases* arise. In particular, three different datacases of different number of customers, are created for Utrecht, which allows to evaluate the results on an increasing scale of customer locations, while controlling for the geographical area. While the Utrecht datacases therefore include 32, 64 and 128 customers, the Overijssel and South-Holland data comprises 64 customers in each of the datasets. This in turn allows to assess the different methods for geographical areas with different characteristics, while controlling for the sample size. Further, for each of the 5 datacases, the random sampling of customer locations is carried out 10 times. As a result, the datacases *Utrecht32, Utrecht64, South-Holland64, Overijssel64 and Utrecht128* are obtained with 10 datasets each, amounting to a total of 50 different *datasets*.

In the Utrecht datasets with 32 customers, it is assumed that 1 vehicle with capacity $c = 8$ is used to visit the customer locations within the 4 days. For the datasets with 64 and 128 customers, 2 and 4 vehicles are required on each day respectively.

As each province is characterized by a different size and spread of customer locations, it is important to interpret results in the light of each datacase. Figure 8 displays density plots representing the varying data characteristics for datacases *Utrecht64, South-Holland64, Overijssel64*. Specifically, it shows the sum of all values in the distance matrix for each dataset of 64 customers, grouped by province. *Overijssel64* (Figure 8 (c)) has the longest distances between entities of the distance matrices with a mean of 33.72 kilometers, followed by *South-Holland64* with a mean distance of 22.54 kilometers (Figure 8 (b)) and *Utrecht64* with a mean of 17.19 kilometers (Figure 8 (a)).

As for the mentioned customer decision scenarios, each customer is given $T = 2$ appointment-days to choose from, except for the upper bound, for which $T = n = 4$. In the case of $T = 2$, for

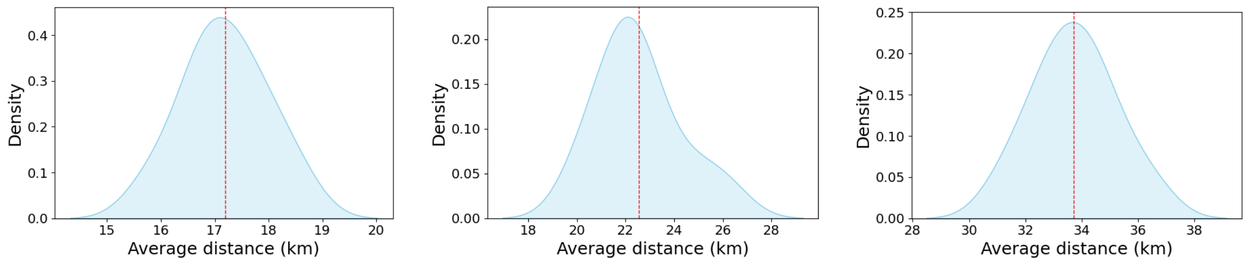

(a) Utrecht64 (b) South-Holland64 (c) Overijssel 64

Figure 8: Average distance distributions for datasets with 64 customer by province

each decision scenario, half of the customers chooses the first option, while the other half chooses the second option. For the upper bound, $\frac{1}{4}$ of the customers choose each of the 4 appointment-day options respectively.

For strategies *CVRP-Skip, CVRP-TSP, KmeansTSP-Skip and Kmeans-TSP*, there are $\binom{16}{8} = 12,870$ different combinations of customers deciding for one of the 2 days offered. To allow for computation in a reasonable time, 50 randomly drawn scenarios, out of these 12,870 combinations, are used to extend the datasets. Hence, for each cluster or customer group that results from the first stage of a strategy, 50 possible scenarios are considered for each dataset. To be able to accurately compare the results between the methods, the same randomly drawn customer decision scenarios are used across all 4 strategies. As a result, the 4 different strategies are tested on 2,500 different instances, across 50 datasets and 50 customer decision scenarios.

For the upper bound strategy, for the scenarios with 32 customers and one vehicle, there are $\binom{32}{8} = 10,518,300$ combinations of the 32 customers directly choosing one appointment-day out of the 4 days available. Accordingly, 50 scenarios are drawn randomly from these combinations. For datasets with 64 and 128 customers, the random choice is repeated and the sequences randomly shuffled accordingly, resulting in a set of 50 unique customer decision scenarios. This, again, results in 2,500 instances, over which the expected value is taken as a overall result of the upper bound.

For the lower bound strategy, no customer decision scenarios are computed, since customers are given no choice over their appointment-day. Hence, the lower bound strategy is evaluated on only 50 different instances.

6 Results

The numeric performance results for each strategy, including the upper and lower bounds, are displayed in Table 2. The table reports the expected value of total distance traveled in kilometers per strategy, combining all final routes that are built for each dataset and customer decision scenario and displays each strategy's percentage difference in distance compared to the upper bound method to enhance comparability between strategies. Overall, taking the expected value across all tested datasets and customer decision scenarios, the best performing strategies, sorted from best to worst performance, are the *CVRP-TSP, CVRP-Skip*, followed by the *Kmeans-TSP* and finally the *KmeansTSP-Skip* strategy. The upper- and lower bounds perform as expected, showing the highest and lowest expected distance traveled, respectively. Section 6.1 compares the results of different strategies for datasets with 64 customers. This comparison aims to explore distance-based differences across provinces and evaluate the performances while controlling for the number of customers. In section 6.2, the results for only the Utrecht province will be presented, comparing the performances of final routes across different numbers of customers involved.

	Utrecht32	Utrecht64	S-Holland64	Overijssel64	Utrecht128	Overall
CVRP-Skip	299.35	468.42	610.70	800.43	794.20	594.62
	(-20.6%)	(-18.3%)	(-19.0%)	(-17.8%)	(-14.7%)	(-17.6%)
CVRP-TSP	294.20	463.21	603.52	794.84	784.59	588.07
	(-22.0%)	(-19.2%)	(-20.0%)	(-18.4%)	(-15.7%)	(-18.6%)
KmeansTSP-Skip	299.79	494.42	634.40	829.18	827.00	616.96
	(-20.5%)	(-13.8%)	(-15.9%)	(-14.9%)	(-11.2%)	(-14.5%)
Kmeans-TSP	292.17	486.52	621.53	814.34	810.80	605.07
	(-22.5%)	(-15.2%)	(-17.6%)	(-16.4%)	(-12.9%)	(-16.2%)
Upper Bound	377.04	573.56	754.14	974.19	931.14	722.01
Lower Bound	227.60	374.00	491.50	643.80	671.40	481.66
	(-39.6%)	(-34.8%)	(-34.8%)	(-33.9%)	(-27.9%)	(-33.3%)

Table 2: Absolute performances in kilometers and relative performances compared to upper bound

6.1 Comparative Analysis Across Provinces

Figure 9 summarizes the strategy performance results, measured by the total expected distance of final routes, for the datacases *Utrecht64, South-Holland64* and *Overijssel64*.

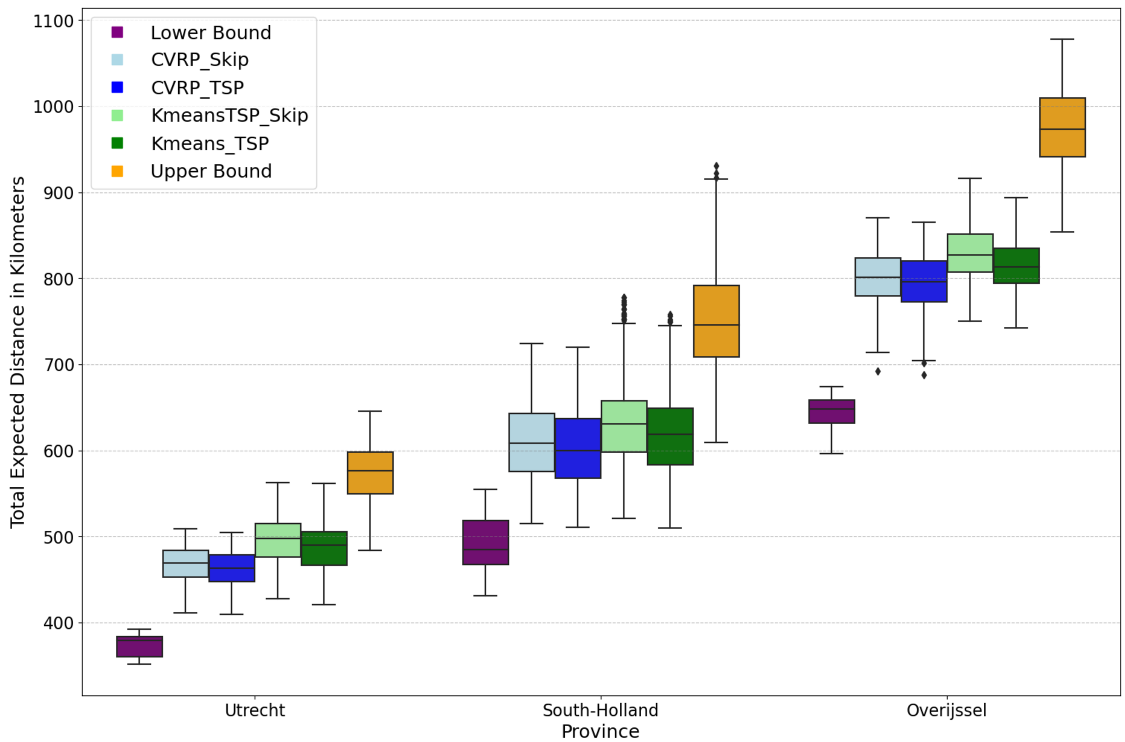

Figure 9: Strategy performances visualized for datasets with 64 customers by province

The results displayed in Figure 9 depict that, on average, the *CVRP-TSP* strategy performs best across all provinces for datasets with 64 customers. As a second best strategy, the *CVRP-Skip* method outperforms both *KmeansTSP-Skip* and *Kmeans-TSP*, among which the *Kmeans-TSP* is the better choice in minimizing the expected travel distance. In particular, it is observable from the result for each province, that with respect to the first stage method, i.e. either *stage1-CVRP* or *stage1-Kmeans*, the reoptimization strategy on average performs better than the Skip approach in stage 2, which is further explored in more depth in section 6.1.2.

While the relative performance of each strategy remains consistent in the hierarchical order of performances for the 64 customer cases, it is observable that the expected total travel distance increases from Utrecht to South-Holland to Overijssel. This pattern likely results from systematic differences in the distance matrices computed for each dataset across these provinces. In other words, customer locations are, on average, the most spread out in Overijssel, followed by South-Holland and Utrecht, as was shown by Figure 8 in section 5.3, therefore naturally resulting in differences in the total expected travel distance. Building on this, the boxplots of Figure 9 depict that the spread in the solution values changes depending on the province and the method. In particular, the span of values is visibly larger for the South-Holland datasets, compared to the Utrecht cases, with 64 customers respectively. The impact of distance-related differences between provinces is further investigated in the following part.

6.1.1 Impact of Distance Ranges on Strategy Performance

Further analysis of province-specific differences in strategy performances reveals that the probability of a strategy being the best performing method fluctuates with the average distance of a dataset's distance matrix, previously addressed in subsection 5.3. Table 3 illustrates this finding by comparing the best strategies over three different average distance ranges of datasets' distance matrices, with a *low range* of below or equal to 20 kilometers, *mid range* of between 20 and 30 kilometers, and a *large range* of beyond 30 kilometers.

Strategies	**Low Range**	**Mid Range**	**Large Range**
CVRP-TSP (%)	85.8%	78.8%	68.8%
Kmeans-TSP (%)	11.2%	18.4%	25.8%
CVRP-SKIP, *CVRP-TSP* (%)	2.8%	1.0%	4.8%
CVRP-TSP, *Kmeans-TSP* (%)	0.2%	1.6%	0.6%
CVRP-SKIP, *CVRP-TSP*, *Kmeans-TSP* (%)	0.0%	0.2%	0.0%
Total (%)	100%	100%	100%

Table 3: Likelihood different strategies showing the best performance by distance range

Table 3 shows that the likelihood of the *CVRP-TSP* approach being the best performing strategy decreases with a higher distance range. At the same time, the likelihood of *Kmeans-TSP* showing the best performance increases with an increase in the distance range. Furthermore, there are instances, in which *CVRP-Skip* and *CVRP-TSP* share the best performance with an equal total travel distance on these instances. In these cases, which amount to 1% to 4.8% of the instances with 64 customers, depending on the distance range, reoptimizing the a priori routes by solving a TSP in stage 2 does not improve the performance of the final routes compared to employing the skip approach in stage 2. Similarly, in rare instances, the *CVRP-TSP* and *Kmeans-TSP* perform equally well in the expected travel distance. This can either suggest that the first stage outcomes of both methods, i.e. how customers are partitioned into groups, are equal, or that the outcomes of stage 1 are different but result in the same expected distance in the final route after solving the TSP in stage 2. Further investigation shows the latter, as none of the instances display the same customer grouping for the K-means algorithm as for the CVRP. Additionally, in only one of the 1,500 instances involving 64 customers, the strategies *CVRP-Skip*, *CVRP-TSP* and *Kmeans-TSP* are jointly the best performing methods.

6.1.2 Impact of Stages on Strategy Performance

To further investigate the performances of different strategies and understand potential root causes of these differences, examining and comparing differences across stages is crucial. Figure 10 shows 4 scatterplots, each of which visualizes two methods' resulting distances against each other for instances of 64 customers. In each of the sub-figures, the blue dots represent the total travel distances computed for *Utrecht64* datasets, while the orange dots display travel distances for *South-Holland64* and the green dots show the results for *Overijssel64*. The dotted line represents equal performance of the strategies displayed on both axes.

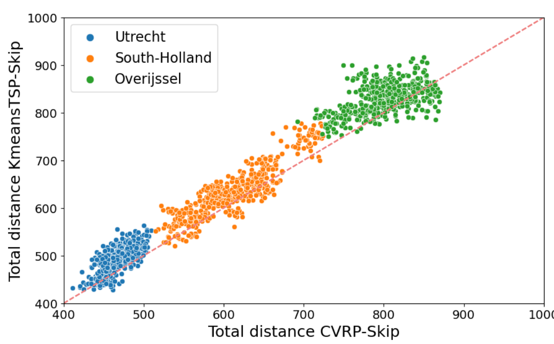

(a) CVRP-Skip vs. KmeansTSP-Skip

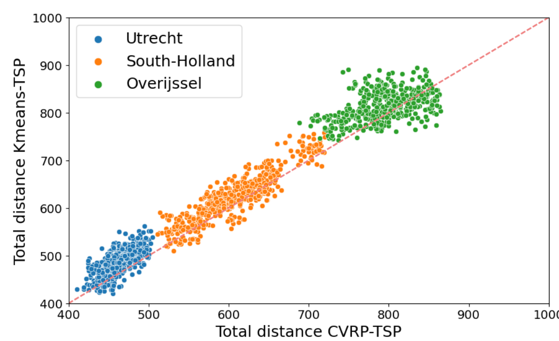

(b) CVRP-TSP vs. Kmeans-TSP

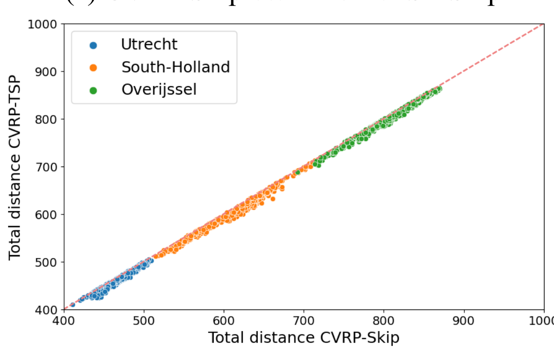

(c) CVRP-Skip vs. CVRP-TSP

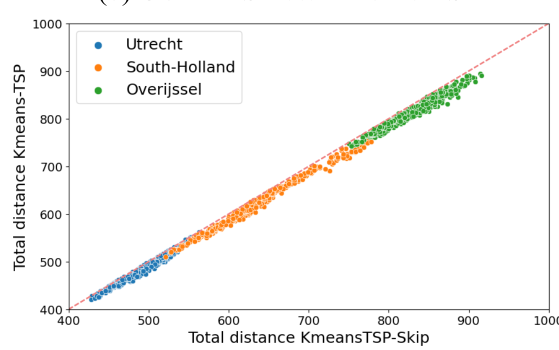

(d) KmeansTSP-Skip vs. Kmeans-TSP

Figure 10: Stage 1 method (a,b) and stage 2 method (c,d) distance comparison for 64 customers

As previously shown in Figures 8 and 9, Figure 10 supports that on average, Utrecht is characterized by the highest customer density, showing the lowest results in total distance traveled, followed by South-Holland and Overijssel. Sub-figures (a) and (b) compare the first stage approach while controlling for the second stage approach, and sub-figures (c) and (d) compare the second stage solution approach while controlling for the first stage approach. The scatterplots paint a clear picture – the plots depicting the first stage comparison between the two skip strategies (*CVRP-Skip* and *KmeansTSP-Skip*) as well as the two TSP strategies (*CVRP-TSP* and *Kmeans-TSP*) in sub-figures (a) and (b), show a much higher spread in the dots than scatterplots (c) and (d), comparing the stage 2 approaches. In other terms, the chosen stage 1 solution approach gives room for much more improvement than the stage 2 solution approach. To support this observation numerically, Table 4 in Appendix A reports the mean percentage and absolute differences in kilometers of distance between the strategies compared in the scatterplots.

Moreover, sub-figures (c) and (d) clearly illustrate that the *stage2-TSP* approach consistently matches or outperforms the performance of the *stage2-Skip* approach across all datasets, when both strategies are controlled for the same stage 1 approach. This is evident from the dots consistently staying on, but mostly below, the dotted line of equal performance, indicating that the *stage2-TSP* approach always yields a lower or equal total distance compared to the *stage2-Skip* approach.

Figure 11 provides a comparison of solving *stage1-CVRP* and *stage1-KmeansTSP* on a dataset with 64 customers located in Overijssel, represented by the dots, and one depot, represented by the red star. It illustrates that the route-based partitioning shown in Figure 11 (a), obtained via solving *stage1-CVRP*, leads to a different grouping outcome, compared to solving the *stage1-KmeansTSP* approach, depicted in Figure 11 (b). Further inspection reveals that in none of the instances, the CVRP and K-means algorithms grouped customers in the same way. For the displayed instance in Figure 11, with the same chosen customer decision scenario for both strategies respectively, the *KmeansTSP-Skip* approach outperforms the *CVRP-Skip* approach by 80 kilometers. In particular, the total distance for *CVRP-Skip* is of 865 kilometers, for *CVRP-TSP* of 859 kilometers, for *KmeansTSP-Skip* of 785 kilometers and for *Kmeans-TSP* of 774 kilometers. In other words, both approaches involving *stage1-Kmeans* outperform the approaches using *stage1-CVRP* in that instance, which underscores the importance of the stage1 approach as suggested by Figure 10.

Further analysis reveals that the performance of a strategy not only depends on how well the chosen approach in stage 1 groups customers together, but also on how convenient the respective customer decision scenario is for combining the customers in stage 2, based on the a priori routes. Hence, assessing performance by the expected value across customer decision scenarios. As an example, for the instance visualized in Figure 11, in 37 of the 50 customer decision scenarios, the *KmeansTSP-Skip* approach performs better than the *CVRP-Skip* approach, which in turn performs better in the other 13 customer decision scenarios.

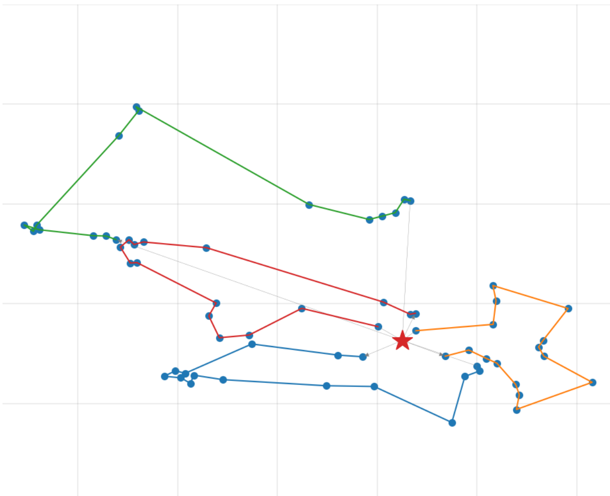

(a) *stage1-CVRP*

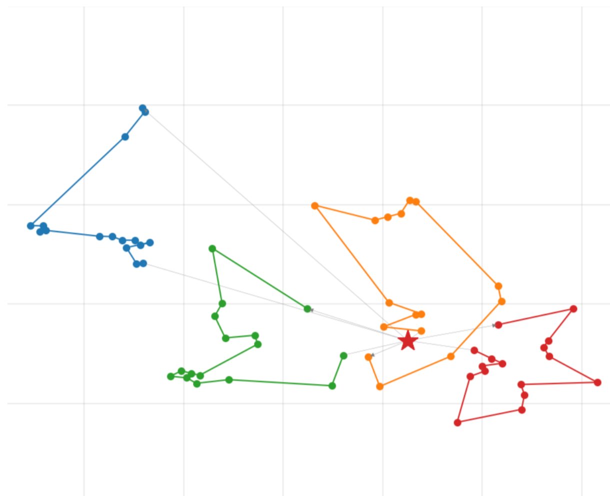

(b) *stage1-KmeansTSP*

Figure 11: Difference between *stage1-CVRP* and *stage1-KmeansTSP* on an Overijssel dataset

6.2 Comparative Analysis Across Different Dataset Sizes

Additional insight into the performance of different strategies is provided by comparing the strategy performances on the province of Utrecht when increasing the number of customers to visit. In that, the performances for each strategy, including the upper- and lower bound, on the datasets for 32, 64 and 128 customers in Utrecht are displayed in Figure 12. In this comparative analysis, not the total expected travel distance over each instance is computed as in Figure 9 and Table 2, but the mean travel distance per truck per appointment-day, as it is of higher interest for this analysis by ensuring comparability between datasets with a different number of customers. Hence, the distances displayed are the distances that one truck, on average, travels to visit all customers on one route. The expected distances matching Figure 12, including the percentage decreases of each strategy's expected distance traveled compared to the upper bound result, are reported in Table 5 in Appendix A.

The best performance by decreasing order follows the sequence of *CVRP-TSP, CVRP-Skip, Kmeans-TSP* and *KmeansTSP-Skip* for the cases with 64 customers, as it was detailed in subsection 6.1. This hierarchy, however, changes with an increase and decrease in size of the dataset. For

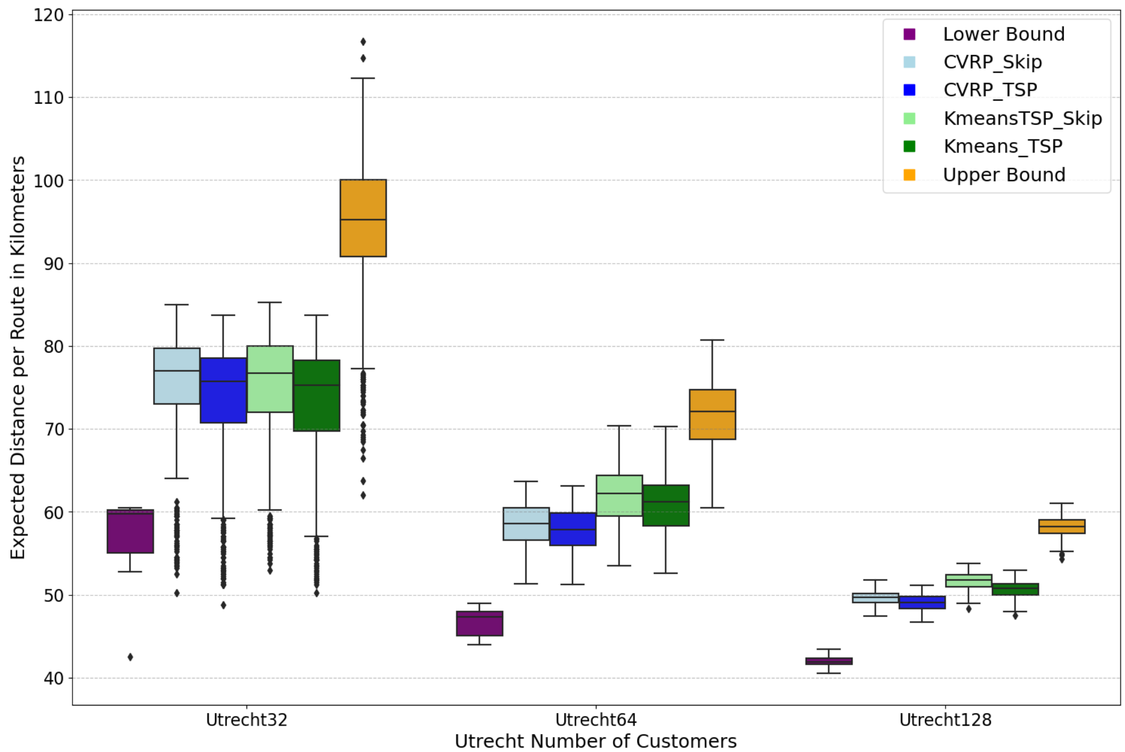

Figure 12: Expected distance per day/truck for Utrecht32, Utrecht64 and Utrecht128 by strategy

cases with 32 customers, the *Kmeans-TSP* shows the best performance, followed by the *CVRP-TSP* approach. The third and fourth best strategies are the *CVRP-Skip* and *KmeansTSP-Skip* respectively. As for increasing the number of customers to 128, the best performance is again given by the *CVRP-TSP* approach, closely followed by the *CVRP-Skip* approach. The third and fourth best strategies for 128 customers are the *Kmeans-TSP* and *KmeansTSP-Skip*. These findings suggest that the performance of each individual strategy relative to the others highly depends on the number of customers involved in the studied instance. The results moreover suggest that, with an increasing number of customers, the stage 2 approach becomes less important compared to the stage 1 method, in which the CVRP solution approach appears to yield the best results. In particular, since in both scenarios of 64 and 128 customers, the two best performing strategies adopt a *stage1-CVRP* approach, the stage 1 method seems to gain importance the more customers are involved in the dataset. In particular, the results point to *stage1-CVRP* as the better choice with an increasing amount of customers.

Moreover, another insight provided by Figure 12 is that, across all strategies, the average distance per route decreases with every incremental increase in the number of customers for the Utrecht area. This further suggests that customers can be grouped more efficiently with a higher number of customers. This is also supported by the decreasing size of boxes in Figure 12, illustrating that the expected travel distance per route becomes more consistent across all instances as the number of customers increases.

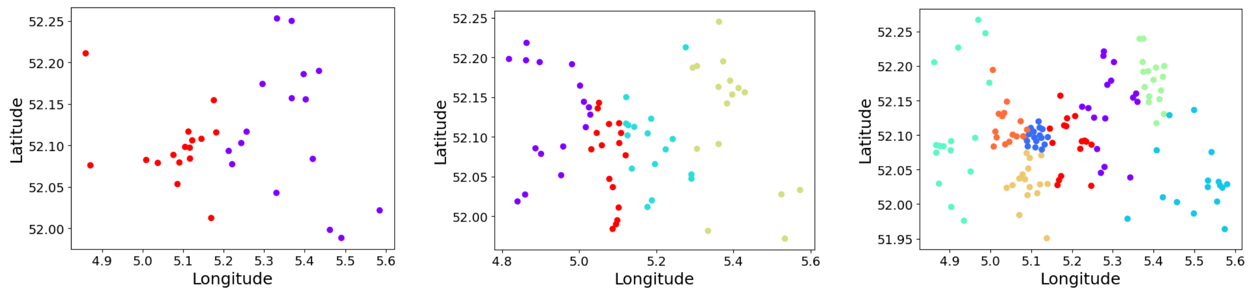

(a) Utrecht32 (b) Utrecht64 (c) Utrecht128

Figure 13: K-means algorithm on Utrecht datasets with 32, 64 and 128 customers

Figure 13 visualizes this by the example of applying the K-means algorithm on a dataset for Utrecht with 32 customers (sub-figure (a)), 64 customers (sub-figure (b)) and 128 customers (sub-figure (c)), respectively. In each plot of Figure 13, differently colored dots represent distinct clusters of 16 customer locations as an outcome of stage 1. Figure 13 showcases that the average within-cluster distance decreases with an increase in the number of customers, supporting the observation of overall decreasing travel distances per day/truck, demonstrated by Figure 12. In fact, in datasets with 32 customers, the average distances between customers within clusters are of 11.2 kilometers, while the average distance reduces to 10.44 kilometers and 7.58 kilometers for datasets with 64 and 128 customers, accordingly.

7 Discussion

In this section, the results are discussed, providing both managerial implications and academic embedding. After recommendations are offered, the study's limitations are detailed, followed by ideas for future research.

7.1 Conclusions

This study provides valuable insights into the performance of different two-stage strategies. In the first stage, the strategies experiment with two different approaches with which the provider groups customers together, with the aim of assigning appointment-day offerings to each group. In stage 2, the final routes have to be determined by the company, where either a skip- or reoptimization approach is adopted. Overall, the results show that the best performing strategies, in decreasing order, are *CVRP-TSP*, *CVRP-Skip*, *Kmeans-TSP*, and *KmeansTSP-Skip*, thereby in part contradicting the expectations regarding efficiency formulated in section 5.2. With that, the results highlight the importance of the approach chosen in stage 1, in which CVRP-based strategies generally outperform strategies that use the K-means algorithms, which is in line with Bührmann and Bruwer (2021) and Rasmussen et al. (2012), contending that clustering may reduce routing performance.

These insights, however, are to be set in a certain context. While the dominance of *CVRP-TSP* is generally high, it shrinks with a higher spread of customer locations. In other words, the more customers are spread out on a geographic space, the better *stage1-Kmeans* performs relative to *stage1-CVRP* in grouping customers. Combining this insight with the observation that the *Kmeans-TSP* strategy outperforms the other strategies on Utrecht datasets with 32 customers suggests that this strategy can be a viable, less computationally demanding (Nallusamy et al., 2009, p. 131), alternative to the *CVRP-TSP* strategy specifically for a smaller number of customers and customer landscapes with a relatively low customer density. Hence, the characteristics of customer locations are found to play a pivotal role in deciding upon which strategy to choose, which is why one strategy cannot be determined as the inherently superior strategy. This is further the case because customer decision scenarios strongly impact a strategy's performance.

While the first-stage method is found to have a substantial impact on the overall performance of a strategy, the second-stage method, which decides upon whether to apply the skip approach or to reoptimize by solving TSPs for each final route, offers less potential for improvement. Generally speaking, the *stage2-TSP* approach always matches or slightly outperforms the *stage2-Skip* approach, thereby providing support for Bertsimas et al. (1990, p. 1020). When controlling for the stage 1 method, the expected impact of reoptimization ranges between roughly 1% and 2% (please refer to Table 4 in Appendix A). This suggests that AHS providers should streamline their efforts into choosing the better stage 1 method, appropriate to their situation, rather than reoptimizing routes, since the latter involves vast computational resources and yields little to no benefit to the final route's total distance, which is in line with the reasoning presented in Schalekamp (2007, p. 1)

Another insight provided by this study is that across all strategies, more efficient routes are produced with an increasing number of customers involved in the instance. In particular, this study finds that the first-stage approaches are able to group customers into more efficient clusters when more customers are involved in each cluster. In support of this finding, Bührmann and Bruwer (2021) find that travel cost increases with the number of clusters. As a consequence, this study demonstrates that introducing a higher number of customers to each cluster or group of customers, on average, leads to efficiency gains for each final route and more robustness in resulting distances.

This study additionally contributes to academic literature by offering support for the trade-off between customer-service level and routing efficiency, previously pointed out by Agatz et al. (2011), Bühler et al. (2016), Côté et al. (2024), Visser and Savelsbergh (2019), and Zhan et al. (2021). In that, the results indicate that a higher level of customer service, represented by more appointment-day choices offered to customers, leads to a decrease in routing performance, as supported by the performance results of the upper- and lower bound.

7.2 Recommendations

To translate these conclusions into managerial implications, a nuanced perspective on companies' decision-making should be offered. AHS providers that aim to strike a balance between *efficiency, complexity and customer service* should acknowledge that a higher customer service promotes inefficiencies in the routing of vehicles, as reflected by the (*free-choice*) upper bound strategy results. Hence, for their first-stage decision making, companies should choose a limited number of appointment-day offerings to balance the *efficiency - customer service* trade-off.

Given a reasonable level of customer service, such as the case of two appointment-day offerings assumed in this thesis, companies are overall well advised with grouping customers with a CVRP, as this approach overall outperforms K-means-based methods in this study, particularly in scenarios with a larger number of densely located customers. However, if a company's customer landscape is rather small and highly spread out, K-means based strategies grow to more competitive alternatives.

AHS providers are further advised to minimize their efforts on second-stage reoptimization, as it offers minimal route efficiency improvements, and instead prioritize selecting the optimal first-stage grouping method. Companies should therefore experiment with the *stage2-Skip* approach on a priori routes, as it generally provides competitive results at low planning complexity.

Relative to their planning horizon, providers should further include a large number of customers into the same planning period, as it enhances first-stage grouping efficiency and overall robustness in minimizing the total distance traveled.

7.3 Limitations & Future Research

This study is subject to limitations. First, restrictions are imposed by computational limitations. In order to obtain results in a reasonable amount of time, only a limited number of datacases, datasets and customer decision scenarios are considered in this thesis, thereby potentially compromising the reliability of results.

Another limitation of this thesis is the usage of the Haversine distance. While being a simplifying assumption for this study, it can have substantial drawbacks in its interpretation when applying the strategies to real-world cases, since it does not consider any road infrastructure but returns the shortest distance between two points using latitude and longitude coordinates. As a consequence, the real-world application of the studied strategies could yield different results than the ones presented in this thesis. Hence, in future research efforts, this study could be extended to a setting, in which existing route infrastructure and potentially traffic information are considered.

Additionally, future research could be conducted on experimenting with different clustering algorithms for the proposed strategies *KmeansTSP-Skip* and *Kmeans-TSP*. Similar to Bührmann and Bruwer (2021), where K-medoids yields better results than K-means, improvements to the cluster-based strategies in this study could be made by employing a different clustering algorithm.

This study only takes the location of customers as information for their grouping in the first stage of a strategy into account. Future research could incorporate other aspects useful to the grouping of customers, such as value- and risk profiles, as suggested by Beheshti et al. (2015, p. 404), to create strategies which are more robust to customer decisions. Additionally, leveraging historical data to predict customer preferences could make first-stage grouping decisions more accurate and further boost the efficiency of final routes.

Finally, future studies could extend this thesis by combining customer locations across provinces and by considering multiple depot locations. Considering that most AHS providers, just as the company studied in this thesis, operate across a larger geographic area than the ones assumed in this work, incorporating this extension would most likely provide a more representative scenario, in which differences between the performances of strategies may become more important.

Appendix A

Methodology

Haversine distance calculation (GeeksforGeeks, 2022):

$$d = 2r \sin^{-1}\left(\sqrt{\sin^2\left(\frac{\phi_2 - \phi_1}{2}\right) + \cos(\phi_1)\cos(\phi_2)\sin^2\left(\frac{\lambda_2 - \lambda_1}{2}\right)}\right)$$

where:

d : Distance between the two points

r : Radius of the sphere

ϕ_1, ϕ_2 : Latitudes of point 1 and point 2

λ_1, λ_2 : Longitudes of point 1 and point 2

Data

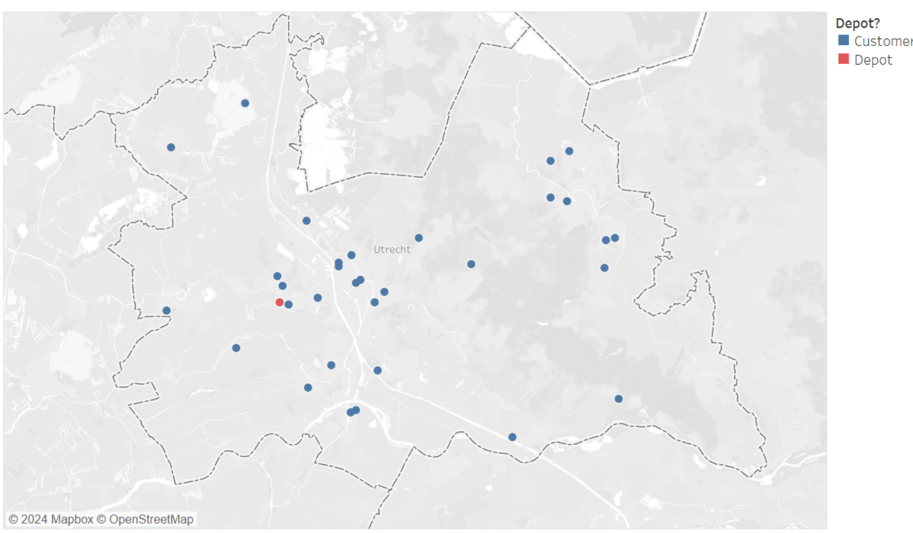

Figure 14: Synthetic depot & 32 customer locations in the Utrecht province

Figure 14 shows a sample dataset of 32 customers, represented by the blue dots, and one depot, shown in red, located in the province of Utrecht. In this datacase (*Utrecht32*), 1 vehicle with capacity $c = 8$ is required to visit all customers across the full planning period of 4 days.

Results

	Mean % Difference	Absolute Difference (km)
CVRP-Skip vs. KmeansTSP-Skip	-4.05%	-26.15
CVRP-TSP vs. Kmeans-TSP	-3.26%	-20.27
CVRP-Skip vs. CVRP-TSP	1.02%	5.99
KmeansTSP-Skip vs. Kmeans-TSP	1.84%	11.87

Table 4: Comparison of the mean percentage difference and the absolute difference in total distance traveled for instances of 64 customers

	Utrecht32	Utrecht64	Utrecht128
CVRP-Skip	74.83	58.55	49.64
	(-20.6%)	(-18.3%)	(-14.7%)
CVRP-TSP	73.55	57.90	49.04
	(-22.0%)	(-19.2%)	(-15.7%)
KmeansTSP-Skip	74.95	61.80	51.69
	(-20.5%)	(-13.8%)	(-11.2%)
Kmeans-TSP	73.04	60.82	50.68
	(-22.5%)	(-15.2%)	(-12.9%)
Upper Bound	94.26	71.70	58.20
Lower Bound	56.90	46.75	41.96
	(-39.6%)	(-34.8%)	(-27.9%)

Table 5: Comparison of distances per route in kilometers for each strategy across Utrecht32, Utrecht64 and Utrecht128, including their relative performance compared to the upper bound

Appendix B

Algorithm 1 High-Level Pseudo Code of Two-Stage Strategies

```
for c in data_cases do
    generate different customer location datasets
    store generated datasets as datasets[c]
    for i in datasets[c] do
        for a in stage1_approaches do
            solve stage 1
            store as solution[c][i][a] in stage1_solutions
            for solution in stage1_solutions[a] do
                for s in customer_decision_scenarios do
                    for d in scenario s decisions do
                        create split datasets as datasets_split[c][i][a][s][d]
                        for j in datasets_split[c][i][a][s][d] do
                            for b in stage2_approaches do
                                solve stage 2
                                store as solution[c][i][a][s][d][j][b] in stage2_solutions
                            end for
                        end for
                    end for
                end for
            end for
        end for
    end for
end for
```

Algorithm 2 Stage 1 Pseudo Code

```
for c in data_cases do
    for i in datasets[c] do
        solve CVRP with capacity (c × T)
        store as solution in CVRP_solutions[c][i]

        apply K-means with cluster of equal size (c × T)
        store as cluster in Kmeans_solutions[c][i]
        for k in Kmeans_solutions[c][i] do
            solve TSP
            store as solution in KmeansTSP_solutions[c][i][k]
        end for
    end for
end for
```

Algorithm 3 Stage 2 Pseudo Code

```
for r in CVRP_solutions[c][i] do
    for s in customer_decision_scenarios do
        for d in scenario s decisions do
            split a priori routes as split_datasets
            for j in split_datasets do
                total distance = \sum_{z∈Z} d[r_z][r_{z+1}] following sequence in r, for z in j
                save total distance as CVRP_Skip_results[c][i][s][d][j]

                solve TSP
                save total distance as CVRP_TSP_results[c][i][s][d][j]
            end for
        end for
    end for
end for
for r in KmeansTSP_solutions do
    for s in customer_decision_scenarios do
        for d in scenario s decisions do
            split a priori routes as split_cluster_datasets
            for j in split_cluster_datasets do
                total distance = \sum_{z∈Z} d[r_z][r_{z+1}] following sequence in r, for z in j
                save total distance as KmeansTSP_Skip_results[c][i][s][d][j]

                solve TSP
                save total distance as Kmeans_TSP_results[c][i][s][d][j]
            end for
        end for
    end for
end for
```

References

Agatz, N., Campbell, A., Fleischmann, M., & Savelsbergh, M. (2011). Time slot management in attended home delivery. *Transportation Science*, *45*(3), 435–449. https://doi.org/10.1287/trsc.1100.0346

Becker, L. (2023, August). Enhancing Technician Efficiency: Comparison of Appointment Scheduling Methods. Business Analytics & Management. http://hdl.handle.net/2105/69619

Beheshti, A. K., Hejazi, S. R., & Alinaghian, M. (2015). The vehicle routing problem with multiple prioritized time windows: A case study. *Computers & Industrial Engineering*, *90*, 402–413. https://doi.org/10.1016/j.cie.2015.10.005

Bertsimas, D., Jaillet, P., & Odoni, A. R. (1990). A Priori Optimization. *Operations Research*, *38*(6), 1019–1033. https://doi.org/10.1287/opre.38.6.1019

Bräysy, O., & Gendreau, M. (2005a). Vehicle routing problem with time windows, Part I: Route construction and local search algorithms. *Transportation Science*, *39*(1), 104–118. https://doi.org/10.1287/trsc.1030.0056

Bräysy, O., & Gendreau, M. (2005b). Vehicle routing problem with time windows, Part II: Metaheuristics. *Transportation Science*, *39*(1), 119–139. https://doi.org/10.1287/trsc.1030.0057

Bruck, B. P., Castegini, F., Cordeau, J.-F., Iori, M., Poncemi, T., & Vezzali, D. (2020). A Decision Support System for Attended Home Services. *INFORMS Journal on Applied Analytics*, *50*(2), 137–152. https://doi.org/10.1287/inte.2020.1031

Bühler, D., Klein, R., & Neugebauer, M. (2016). Model-based delivery cost approximation in attended home services. *Computers and Industrial Engineering*, *98*, 78–90. https://doi.org/10.1016/j.cie.2016.05.014

Bührmann, J. H., & Bruwer, F. (2021). K-Medoid Petal-Shaped Clustering for the Capacitated Vehicle Routing Problem. *The South African Journal of Industrial Engineering*, *32*(3). https://doi.org/10.7166/32-3-2610

Campbell, A. M., & Savelsbergh, M. W. P. (2005). Decision Support for Consumer Direct Grocery Initiatives. *Transportation Science*, *39*(3), 313–327. https://doi.org/10.1287/trsc.1040.0105

Campbell, A. M., & Thomas, B. W. (2008). Challenges and advances in a priori routing. In B. Golden, S. Raghavan, & E. Wasil (Eds.), *The vehicle routing problem: Latest advances and new challenges* (pp. 123–142). Springer US. https://doi.org/10.1007/978-0-387-77778-8_6

CBS. (n.d.). *Numerieke postcode (pc4), 2022* [Licensed under CC BY 4.0. https://creativecommons.org/licenses/by/4.0/]. Retrieved March 12, 2024, from https://www.cbs.nl/nl-nl/dossier/nederland-regionaal/geografische-data/gegevens-per-postcode

Cordeau, J. F., Iori, M., & Vezzali, D. (2023). A survey of attended home delivery and service problems with a focus on applications. *4OR-Q J Oper Res*, *21*(4), 547–583. https://doi.org/10.1007/s10288-023-00556-2

Côté, J.-F., Mansini, R., & Raffaele, A. (2024). Multi-period time window assignment for attended home delivery. *European Journal of Operational Research*, *316*(1), 295–309. https://doi.org/10.1016/j.ejor.2024.01.021

Dantzig, G. B., & Ramser, J. H. (1959). The Truck Dispatching Problem. *Management Science*, *6*, 80–91. https://doi.org/https://doi.org/10.1287/mnsc.6.1.80

GeeksforGeeks. (2022, September 5). *Haversine formula to find distance between two points on a sphere*. Retrieved June 11, 2024, from https://www.geeksforgeeks.org/haversine-formula-to-find-distance-between-two-points-on-a-sphere/

Holland, C., Levis, J., Nuggehalli, R., Santilli, B., & Winters, J. (2017). UPS optimizes delivery routes. *Interfaces*, *47*(1), 8–23. https://doi.org/10.1287/inte.2016.0875

Iswari, T., & Asih, A. M. S. (2018). Comparing genetic algorithm and particle swarm optimization for solving capacitated vehicle routing problem. *IOP Conference Series: Materials Science and Engineering*, *337*(1), 12004. https://doi.org/10.1088/1757-899X/337/1/012004

Jaillet, Patrick. (1988). A priori solution of a travelling salesman problem in which a random subset of the customers are visited. *Operations Research*, *36*(6), 929–936. https://doi.org/10.1287/opre.36.6.929

Korayem, L., Khorsid, M., & Kassem, S. S. (2015). Using grey Wolf algorithm to solve the capacitated vehicle routing problem. *IOP Conference Series: Materials Science and Engineering*, *83*(1). https://doi.org/10.1088/1757-899X/83/1/012014

Laporte, G., Gendreau, M., Potvin, J. Y., & Semet, F. (2000). Classical and modern heuristics for the vehicle routing problem. *International Transactions in Operational Research*, *7*(4-5), 285–300. https://doi.org/10.1111/j.1475-3995.2000.tb00200.x

Laporte, G. (2007). What you should know about the vehicle routing problem. *Naval Research Logistics*, *54*(8), 811–819. https://doi.org/10.1002/nav.20261

Liu, R., Yuan, B., & Jiang, Z. (2019). A branch-and-price algorithm for the home-caregiver scheduling and routing problem with stochastic travel and service times. *Flexible Services and Manufacturing Journal*, *31*(4), 989–1011. https://doi.org/10.1007/s10696-018-9328-8

Molina, J. C., Salmeron, J. L., & Eguia, I. (2020). An acs-based memetic algorithm for the heterogeneous vehicle routing problem with time windows. *Expert Systems with Applications*, *157*, 113379. https://doi.org/10.1016/j.eswa.2020.113379

Nallusamy, R., Duraiswamy, K., Dhanalakshmi, R., & Parthiban, P. (2009). Optimization of Multiple Vehicle Routing Problems using Approximation Algorithms. *International Journal of Engineering Science and Technology*, *1*(3), 129–135. https://doi.org/10.48550/arXiv.1001.4197

Nicolas Rincon-Garcia, B. J. W., & Cherrett, T. J. (2018). Requirements from vehicle routing software: perspectives from literature, developers and the freight industry. *Transport Reviews*, *38*(1), 117–138. https://doi.org/10.1080/01441647.2017.1297869

Rasmussen, M. S., Justesen, T., Dohn, A., & Larsen, J. (2012). The Home Care Crew Scheduling Problem: Preference-based visit clustering and temporal dependencies. *European Journal of Operational Research*, *219*(3), 598–610. https://doi.org/10.1016/j.ejor.2011.10.048

Rowell, M., Gagliano, A., Wang, Z., Goodchild, A. V., Sage, J. L., & Jessup, E. L. (2012). *Improving Statewide Freight Routing Capabilities for Sub-National Commodity Flows* (tech. rep.). Washington State Department of Transportation.

Schalekamp, F. (2007). *Some Results in Universal and A Priori Optimization* [Doctoral dissertation]. Retrieved March 17, 2024, from https://ecommons.cornell.edu/server/api/core/bitstreams/b962d8e0-e12a-4c1e-8f37-d23b74c13a38/content

Spliet, R., & Gabor, A. F. (2015). The time window assignment vehicle routing problem. *Transportation Science*, *49*(4), 721–731. https://doi.org/10.1287/trsc.2013.0510

Thangiah, S. R., Nygard, K. E., & Juell, P. L. (1991). Gideon: A genetic algorithm system for vehicle routing with time windows. *[1991] Proceedings. The Seventh IEEE Conference on Artificial Intelligence Application*, *1*, 322–328. https://doi.org/10.1109/CAIA.1991.120888

Vidal, T. (2022). Hybrid genetic search for the CVRP: Open-source implementation and SWAP* neighborhood. *Computers and Operations Research*, *140*. https://doi.org/10.1016/j.cor.2021.105643

Vidal, T., Crainic, T. G., Gendreau, M., & Prins, C. (2013). A hybrid genetic algorithm with adaptive diversity management for a large class of vehicle routing problems with time-windows. *Computers and Operations Research*, *40*(1), 475–489. https://doi.org/10.1016/j.cor.2012.07.018

Visser, T., & Savelsbergh, M. (2019, January). *Strategic Time Slot Management: A Priori Routing for Online Grocery Retailing*. Econometric Institute, EUR. http://hdl.handle.net/1765/115051

Wang, Y., Ma, X., Xu, M., Wang, Y., & Liu, Y. (2015). Vehicle routing problem based on a fuzzy customer clustering approach for logistics network optimization. *Journal of Intelligent and Fuzzy Systems*, *29*(4), 1427–1442. https://doi.org/10.3233/IFS-151578

Waßmuth, K., Köhler, C., Agatz, N., & Fleischmann, M. (2023). Demand management for attended home delivery—A literature review. *European Journal of Operational Research*, *311*(3), 801–815. https://doi.org/10.1016/j.ejor.2023.01.056

Wouda, N. A., Lan, L., & Kool, W. (2024). PyVRP: A High-Performance VRP Solver Package. *INFORMS Journal on Computing*, (0). https://doi.org/10.1287/ijoc.2023.0055

Yurchak, R., & Casares, J. (2021, June). *Postal-codes-data: Nl_full.txt* [Licensed under CC BY 4.0. https://creativecommons.org/licenses/by/4.0/]. Retrieved March 13, 2024, from https://github.com/symerio/postal-codes-data/blob/master/data/geonames/NL_full.txt

Zhan, Y., Wang, Z., & Wan, G. (2021). Home service routing and appointment scheduling with stochastic service times. *European Journal of Operational Research*, *288*(1), 98–110. https://doi.org/10.1016/j.ejor.2020.05.037

Zhang, H., Ge, H., Yang, J., & Tong, Y. (2022). Review of Vehicle Routing Problems: Models, Classification and Solving Algorithms. *Archives of Computational Methods in Engineering*, *29*(1), 195–221. https://doi.org/10.1007/s11831-021-09574-x