MASTERARBEIT | MASTER'S THESIS

Titel | Title

Optimizing Last-Mile Delivery: An Analysis of Hybrid Truck/ Robot Solutions

verfasst von | submitted by

Timo Kemnik BSc

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of

Master of Science (MSc)

Wien | Vienna, 2025

Studienkennzahl lt. Studienblatt | Degree programme code as it appears on the student record sheet:	UA 066 915
Studienrichtung lt. Studienblatt | Degree programme as it appears on the student record sheet:	Masterstudium Betriebswirtschaft
Betreut von | Supervisor:	Ass.-Prof. Dr. Christian Tilk

Abstract

Weltweit befindet sich der Warenverkehr in einem starken Wachstum, für das keine Trendwende in den nächsten Jahren prognostiziert wird. Neben dem starken Wachstum, insbesondere im E-Commerce, welches für die Logistikprozesse eine große Herausforderung sein würde, bieten die Unternehmen in dem stark umkämpften Markt zusätzlich immer weitere Serviceleistungen an, um sich von der Konkurrenz abzugrenzen. Insbesondere die immer kürzer werdenden Lieferzeiten stellen den individuellsten und damit kostenintensiven Teil, der letzten Meile des Logistikprozesses, vor eine große Herausforderung. Um diesen Herausforderungen zu begegnen, werden innovative Methoden entwickelt, wozu das im Rahmen dieser Arbeit untersuchte Modell der sogenannten hybriden Truck/Roboter-Zustellung gehört. Hierfür wird in dieser Arbeit ein heuristisches Modell von Ostermeier et al. (2022) weiterentwickelt, um die Lösungsqualität zu erhöhen und zusätzlich werden die Einflüsse verschiedener Parameter auf die Gesamtkosten untersucht. Im Rahmen dieser Weiterentwicklung zeigt sich, dass die Lösungsqualität im Vergleich zu dem Modell von Ostermeier et al. (2022) deutlich verbessert werden konnte. Die Ergebnisse zeigen zusätzlich, dass größere maximale Transportkapazitäten des Vans nicht zu geringen Gesamtkosten führen müssen und die Zusammensetzung des genutzten Datensatzes einen signifikanten Einfluss auf die Lösung hat. Die Ergebnisse lassen den Entschluss zu, dass in diesem Modell die individuellen zeitabhängigen operativen Kosten für den Van einen stärkeren Einfluss auf die Kostenstruktur haben als die individuellen zeitabhängigen Kosten der Roboter. Die Untersuchung hat zusätzlich ergeben, dass die distanzabhängigen Kosten in der Kostenstruktur einen höheren Einfluss haben als die zeitabhängigen Kosten.

Abstract

The global transport of goods is experiencing strong growth, for which no trend reversal is forecast in the next several years. In addition to the strong growth, particularly in e-commerce, which would be a major challenge for logistics processes, companies in this highly competitive market are also offering an increasing number of additional services in order to set themselves apart from the competition. In particular, decreasing delivery times pose a major challenge for the most customized and therefore cost-intensive part of the logistics process, the last mile. Innovative methods are being developed to meet these challenges, including the so-called hybrid truck/robot delivery model analyzed in this thesis. For this purpose, a heuristic model by Ostermeier et al. (2022) is further developed in this thesis in order to increase the solution quality and, in addition, the influences of various parameters on the total costs are analyzed. This further development shows that the solution quality could be significantly improved compared to the model of Ostermeier et al. (2022). The results also reveal that larger maximum transport capacities of the van do not necessarily lead to lower total costs and that the composition of the data set used has a significant influence on the solution. The results allow the conclusion that in this model the individual time-dependent operational costs for the van have a stronger influence on the cost structure than the individual time-dependent costs of the robots. The study also showed that the distance-dependent costs have a greater influence on the cost structure than the time-dependent costs.

Table of Contents

Chapter 1: Introduction

This thesis examines the hybrid use of vans and robots as a possible approach to address the challenges of last mile delivery. Global retail sales are on the rise and the B2C e-commerce sector in particular is expected to grow significantly this decade from EUR 2.65 trillion in 2020 to EUR 5.92 trillion in 2029, which corresponds to growth of 223% (Statista 2024). This sharp rise in sales in B2C e-commerce will also lead to a significant increase in global shipment volumes. As no reversal of this trend is to be expected, companies and logistics service providers are under pressure to cope with this increasing volume and therefore constantly have to improve and develop their delivery systems and processes in order to enhance efficiency. Especially the last-mile logistic is facing many different challenges, and therefore many approaches for solving these challenges occurred over the last decades. The last-mile is a significant cost factor in the supply chain, especially if the recipient has to be in attendance to accept the delivery (Hübner, Kuhn, and Wollenburg 2016). Some studies even describe the last-mile of the delivery process as a bottleneck for many companies, as it represents the most individualized and consequently the least efficient segment of many supply chains, where the specific needs of the end customer take precedence (Macioszek 2018). According to Vanelslander et al. (2013), the last-mile can constitute up to half of the total logistics costs, depending on the industry. This significant cost burden has therefore driven considerable interest in optimizing this segment of the supply chain.

Chapter 2: Literature Review

2.1 Last-Mile Logistics

According to Lim, Jin, and Srai (2018), the term "last-mile" has its roots in the telecommunications industry, where it is used to describe the final section of a network. Although they vary in detail, existing definitions of last-mile logistics generally agree in referring to the last segment of the delivery process, which in a B2C-delivery process usually means from a distribution center to the final customer. In the context of last-mile logistics in B2C operations, this refers to the stage where the logistics provider directly interacts with the end customer (Bosona 2020). Within the realm of last-mile logistics, various systems are employed. One approach involves distributing from multiple locations to numerous customers, where, for instance, several restaurants or food vendors supply a wide range of customers through a service provider, functioning in a decentralized manner. Conversely, the traditional hub-andspoke model operates from a central location, from which end customers are served (Ulmer and Streng 2019; Dai et al. 2019).

There are many different concepts currently used for classic last-mile delivery. One of these is the traditional delivery by a person who drives a van and hands over the consignments. A comparatively newer concept and variation of this is the use of cargo bikes, where a person also drives the vehicle and hands over the consignment. Another rising concept is the use of parcel lockers, which are supplied by larger vehicles and customers must collect their consignments themselves (Boysen, Fedtke, and Schwerdfeger 2021). However, the classic models offer little scope for improvement in the context of rising consignment numbers and high traffic volumes, particularly in urban areas (Ranieri et al. 2018). Studies have indicated that in urban areas, traditional van-only delivery services are capable of delivering up to 120 parcels within an eight-hour work shift (Bosona 2020). In response to the expanding parcel delivery market, service providers and retailers are increasingly promoting and offering shorter delivery times. This trend, however, necessitates a corresponding increase in traffic to meet these expedited delivery schedules (Savelsbergh and Van Woensel 2016). In addition to the elevated costs, the growing volume of traffic imposes a significant burden on urban environments. This increase in traffic results in higher greenhouse gas emissions, which deteriorate air quality and contribute to increased noise levels. For instance, in the UK, using private cars for shopping at local stores generates, depending on the assumptions, up to twenty-four times more CO2 emissions compared to home delivery by van. This figure, however, fluctuates depending on residential location; in suburban areas, emissions are approximately five times higher than for people that live in city centers. Investigating these relationships is difficult, as other activities are usually also conducted when using private cars, making it a challenge to isolate pure substitution effects. Nevertheless, if such substitutions were to occur, the total kilometers driven per delivered shipment could potentially be reduced by up to 70%. Moreover, most cities are not equipped to handle the growing traffic demand, exacerbating the strain on infrastructure and overall urban sustainability. (Demir et al. 2015; Van Loon et al. 2014; Edwards, McKinnon, and Cullinane 2009; Mokhtarian 2004)

Numerous contemporary approaches are under development to address the complexities of last-mile delivery. An illustration of this is found in crowd shipping, wherein any registered individual can utilize a mobile application to accept deliveries using their preferred mode of transportation (Buldeo Rai et al. 2017). A variation of the parcel locker is, for example, the use of private car boots, which are then filled by couriers and the parcel can be collected at any time (Boysen, Fedtke, and Schwerdfeger 2021). Punakivi, Yrjölä, and Holmström (2001) have tested the concepts of the delivery box and reception box, which, according to their research, have the potential to reduce delivery costs by up to 60% as they eliminate the need for customers to receive deliveries in person. The reception box is fixed to the wall, garage, or a similar structure at the customer's residence, allowing deliveries to be securely placed inside and locked. In contrast, the delivery box is a mobile unit brought by the delivery service and temporarily attached to the customer's home. The customer retrieves the shipment from the box, which is then collected by the delivery service for reuse.

Modern approaches to delivery involve the integration of drones and autonomous robots, each presenting distinct methods. There are similarities to the classic vehicle routing problem (VRP), as the drones or robots usually start their route at a depot, follow a route and return to the depot at the end, which will be illustrated in more detail in this paper (Boysen, Fedtke, and Schwerdfeger 2021). As part of advancing technological development and Industry 4.0, the individual processes and participants in the entire supply chain are becoming increasingly digitalized and more connected. This enhanced data exchange and automation among various stakeholders significantly improves efficiency, which also extends to the last mile of the delivery process. (Bányai, Illés, and Bányai 2018)

2.2 Hybrid Truck/Robot Deliveries

Three different main scenarios have currently been developed for the possible uses of autonomous delivery robots.

The first option involves small robots traveling at walking speed on the pavement to deliver the consignments to their destination. The second scenario is the use of significantly larger autonomous robots that drive on the road at the size of cars, although it is debatable whether these are still classified as delivery robots. The third scenario, which is also the focus of this thesis, is that the small delivery robots are transported by a van and then unloaded to carry out the final stage of the entire delivery process (Hossain 2023; Pani et al. 2020; M. Figliozzi and Jennings 2020). Murray and Chu (2015) first introduced the hybrid concept of combining vans that unload robots, which then deliver the parcels. The interest and relevance of this comparatively new topic has increased significantly in recent years, as evidenced by the fact that over 90% of scientific papers dealing with autonomous delivery robots have been published since 2019 (Srinivas, Ramachandiran, and Rajendran 2022). In scientific research on last-mile delivery, the focus is typically on areas such as cost, sustainability, and effectiveness while the majority of studies addressing the van-based robot routing problem primarily target objectives related to time or cost (Yu, Puchinger, and Sun 2022; Mangiaracina et al. 2019).

The main concept of the majority of hybrid truck-and-robot studies involves using the truck as a so-called mother ship, which is used to load the robots and parcels, transport them to the unloading points and then unload the robots there (Ostermeier, Heimfarth, and Hübner 2022; Boysen, Schwerdfeger, and Weidinger 2018), while some concepts also take into account the collection of the robots after successful delivery (Yu, Puchinger, and Sun 2022). Boysen, Schwerdfeger, and Weidinger (2018) argue in their paper that the slow speed of the robots would result in long waiting times for the van if it had to collect them again. The advantages of the hybrid approach could therefore be maximized if they would return to the depots independently. The aim of using the van as a mother ship is to compensate for the limitations of delivery robots, such as their low speed and limited range, and thus optimize the last-mile (Yu, Puchinger, and Sun 2022). These disadvantages of robots will be discussed later in this paper. Other comparable concepts have also been developed, such as the one by Bakach et al. (2021), which has the approach that the trucks supply larger robot depots and these then operate from their depots and supply the end customers.

A major advantage of using robots for delivery is the significant reduction in personnel costs compared to traditional delivery concepts. Ideally, the robots are able to navigate autonomously and, if permitted by law, several robots can be monitored simultaneously by one employee, which significantly increases the efficiency of the personnel costs used. In addition, the operating costs such as maintenance and energy consumption of the robots are estimated to be comparatively very low (Bakach, Campbell, and Ehmke 2021; Boysen, Fedtke, and Schwerdfeger 2021). Small electric delivery robots are significantly more energy-efficient compared to traditional fuel-powered vehicles or even modern flying drones (M. A. Figliozzi 2017).

As Gerrits and Schuur (2021) have outlined, there is a trend in increasing numbers of urban regions worldwide to establish car-free city centers or high-costs zones for driving in these areas. While these areas generally remain accessible for delivery traffic at present, there is still a risk that this will also change in the future. Taking this development into account, the hybrid concept of van-robot delivery could be an innovative and promising concept for the future.

The research has shown that customers are willing to pay to use autonomous delivery robots, with interest also being highest in urban areas where the nearest shopping facilities are within walking distance. It was also emphasized that the decisive criteria are not necessarily the affordability of the delivery options, but that flexibility and reliability play a crucial role. The combination of a delivery option that is cost-efficient, adaptable and reliable is what accelerates customer satisfaction and acceptance.(Pani et al. 2020; Kapser and Abdelrahman 2020)

The latest generation of delivery robots, such as those from the company Starship, are capable of transporting small parcels weighing up to 10 kilograms at walking speed, mainly using pedestrian routes. According to the manufacturer, these robots can fulfill delivery orders for a whole day with a full battery. As already mentioned, one of the major weaknesses of the robots is their low speed, which has a negative impact on the efficiency and duration of last-mile delivery. (Starship Technologies 2024; Boysen, Fedtke, and Schwerdfeger 2021). An additional challenge in the implementation of robots in logistics processes is that currently, due to the design and functionality of the robots, only home-attended delivery is possible, as a person must remove the parcel from the robot. In addition, the current assumption is that only one recipient can be supplied per delivery process (Boysen, Schwerdfeger, and Weidinger 2018; Ostermeier, Heimfarth, and Hübner 2022).

In addition to the technical limitations, there are also various legal issues to be clarified, as legislation in most countries has not yet regulated the details of how to deal with autonomous robots. In particular, the question of who is liable in the event of accidents when they drive completely autonomously, are monitored by a human or are only partially controlled is often still unresolved. On which roads the robots are allowed to drive also depends on many different factors such as size, weight, speed, safety, etc. (Hoffmann and Prause 2018). However, beyond the legal regulations and technical constraints on the speed of delivery robots, the acceptance among other non-motorized road users, such as pedestrians and cyclists, tends to decrease as the speed of these robots increases. Moreover, increasing the speed of delivery robots would necessitate more complex systems, which would, in turn, drive up costs related to research, development, procurement, and ongoing operations (Plank et al. 2022). In order to be able to use the robots optimally and in a way that maximizes benefits, the properties of the robots must take into account many different factors such as capacities, legal regulations, tasks and the conditions of their environment (Plank et al. 2022).

The two industry giants, FedEx (with its Roxo robot) and Amazon (with its Scout robot), both tested their delivery robots in field trials. However, both companies discontinued the projects without providing deeper reasons for their actions (Garland 2023). Companies like the US-based Daxbot offer their delivery robots as a service to other businesses, allowing them to integrate these robots into their delivery networks. The Daxbot models currently function semi-autonomously, meaning that in certain situations where the robot is unsure how to proceed, a Daxbot employee can remotely assume control (Daxbot 2024).

In cooperation with Mercedes, Starship Technologies has already put this approach into practice and tested it, but it has not yet gone into regular operation (Starship Technologies 2024).

2.3 Drone compared to robots

In addition to autonomous delivery robots, drones are also part of the technological development and another way to optimize the last mile. As there are many different challenges to overcome when introducing robots, with or without a van, this also applies to a possible competing product, the aerial drone (Boysen, Fedtke, and Schwerdfeger 2021). Compared to a ground-based robot, the aerial drone has significant disadvantages, particularly in terms of possible payload, legal restrictions, and the complexity of monitoring. Drones are also limited in terms of possible payload weight and the number of consignments that can be transported and are therefore also rather unsuitable for delivering larger consignments or supplying several customers. Legal regulations, in particular airspace regulations, security and privacy concerns also complicate the use and development of deployment options, especially in urban and densely populated areas. Monitoring the drones during the flight is also much more complex in comparison to ensure a safe procedure, especially when several drones are operating in the same airspace. (Chung, Sah, and Lee 2020; Bosona 2020; Rodrigues et al. 2022)

The characteristics of drones also offer advantages over ground-based robots. Compared to robots, drones are able to reach significantly higher speeds and are not dependent on infrastructure or roads, which in turn results in significantly shorter delivery times. This allows them to travel directly and avoid obstacles such as traffic or construction sites and they are therefore less limited by external influences, which minimizes the risk of delays. These characteristics of drones make them particularly interesting for time-sensitive and high-value consignments.(Boysen, Schwerdfeger, and Weidinger 2018; Boysen, Fedtke, and Schwerdfeger 2021; Murray and Chu 2015)

The approaches of delivery by using flying drones are largely comparable to those of ground-based robots. Some concepts also include logistics depots from which the drones deliver to customers, while others also utilize a hybrid approach in which the drones are launched from a van in the target region to complete the last-mile (Murray and Chu 2015; Agatz, Bouman, and Schmidt 2018; Boysen, Fedtke, and Schwerdfeger 2021). Due to the significantly higher speed of drones compared to robots operating on the sidewalk, in this model the approach of the drones returning to the van's new position after delivery to make further deliveries is a much more viable approach. In addition, battery replacement would also be a possible scenario in order to minimize the limitations of the range and use the drones efficiently (Boysen, Schwerdfeger, and Weidinger 2018; Alfandari, Ljubić, and De Melo Da Silva 2022; Goodchild and Toy 2018). Moeini and Salewski (2020) have combined the two hybrid approaches in their model to solve the ‘Truck-Drone-ATV Routing Problem’ with the help of a genetic algorithm. In this approach, the van is again used as a so-called mother ship to transport and drop off both robots and drones in this case. The unloading of the robots or drones is decided by various factors such as the characteristics of the shipment in order to carry out the final delivery to the customer.

Chapter 3: Methodology

3.1 Research Design

This thesis investigates the optimization of a hybrid truck/robot problem, parts of which are comparable to the classical vehicle routing problem (VRP). In the classical VRP, the route planning of a truck is addressed, but in contrast to the model used in this work, the truck must visit each stop exactly once. Therefore, elements of it are comparable to the route planning of the van in the VRP, but this is a modification and the additional planning of the robots is also added.(Ostermeier, Heimfarth, and Hübner 2022; Eksioglu, Vural, and Reisman 2009; Golden, Magnanti, and Nguyen 1977; Clarke and Wright 1964)

The aim is to investigate the influence of various parameters - in particular the capacities of the transport vehicles and the individual cost structures - on the quality of the solution. A heuristic optimization method is used to address the problem, as the VRP is considered to be an NP-hard problem and more exact methods are not practical for this approach due to excessive computing capacity (Alfandari, Ljubić, and De Melo Da Silva 2022). Instead, a combined Iterated Local Search (ILS) - Variable Local Search (VLS) approach is used in this thesis, as these methods are considered in the literature to be powerful methods for solving complex combinatorial optimization problems (Mladenović and Hansen 1997; Lourenço, Martin, and Stützle 2010).

The algorithm was implemented in C++ and the tested instances are based on the same data sets and cost parameters as in Ostermeier's et al. (2022) model. This ensures that the results obtained can be directly compared with the existing work. A detailed description of the data structure is provided in a later section of this paper. Another important aspect of this work is the analysis of cost structures. Individual time costs for trucks and robots as well as route-dependent costs that are identical for both means of transport are analyzed.

To evaluate the developed model, experiments were analyzed to investigate:

1. how different truck capacities affect the solution,
2. the extent to which individual cost structures affect the solution results,
3. how the results of the various extensions of the model change in comparison to Ostermeier's et a. (2022) original model.

The results are evaluated based on total costs. The aim of this analysis is to examine not only the performance of the extended model, but also to gain a better understanding of the factors influencing hybrid truck/robot delivery networks.

3.2 Problem Description

The work is based on an extension of Ostermeier's et al. (2022) model, which in turn is based on Boysen's et al. (2018) work. While Ostermeier's model makes specific assumptions about time windows and waiting times for delivery by the robots, the version developed here excluded these restrictions.

The model extended in this work is based in parts, mainly in the construction of the truck tour, on the basic principles of the Open Vehicle Routing Problem (OVRP). The main difference between the classic Vehicle Routing Problem (VRP) and the OVRP is that, as in the model of this thesis, the truck does not have to return to the starting point as is usual in the VRP, but the allocation ends at the last customer, in this case the depot/drop-off point. In addition, the planning of the truck tour is subject to the capacity restriction in the number of robots that can be transported, so it can be seen as a variation of the capacitated Open Vehicle Routing Problem (COVRP). Further decisive differences to these known problems are that in the model of this work the vans do not have to use all stops (in this case depots and drop-off points) and the additional planning of the robot assignment. (Ostermeier, Heimfarth, and Hübner 2022; Letchford, Lysgaard, and Eglese 2007; Eksioglu, Vural, and Reisman 2009; Repoussis et al. 2010; Golden, Magnanti, and Nguyen 1977; Clarke and Wright 1964)

The model developed in this thesis describes a network of customers who have to be supplied by vans/delivery robots in a hybrid model and will be examined in detail later in the thesis. Due to the previously explained assumption of hybrid delivery, the concept consists of a van that loads the parcels and robots at the beginning of the tour and then starts the tour. The van unloads the robots for final deliveries either at so-called drop-off points, where the robots are just dropped off and the van moves on, or depots, where new robots can also be loaded, from where they subsequently deliver to the customers. The model is based on the assumption that the robots return to the depot or collection points independently after delivery from the customer and are no longer included in the cost planning. Furthermore, for the simplification of this model, no constraints are placed on the range limitations of the robots, as it is assumed that the current capacities of the robots already available on the market are sufficient for the size of the data set.(Ostermeier, Heimfarth, and Hübner 2022; Boysen, Fedtke, and Schwerdfeger 2021)

To summarize, the model consists of three different steps, which will be explained in more detail later in the thesis. In the first step, the van's route is planned, to which the robots are then assigned to the customers and the respective stops in the second step. In the third step, an ILS - VLS approach is used to further optimize this planning. It is important to emphasize that only the starting point of the tour is defined, and all customers must be served. Any stops at depots or drop-off points can be removed and added and not all of them are required to be integrated into the tour, the solution only needs to be feasible at the end of the process.

The objective of this problem is to minimize the total costs, which include both time-dependent costs for trucks and robots as well as distance-dependent costs that are identical for both vehicle types. The optimization aims to find an efficient routing strategy that balances vehicle capacities and operational constraints while reducing overall expenses.

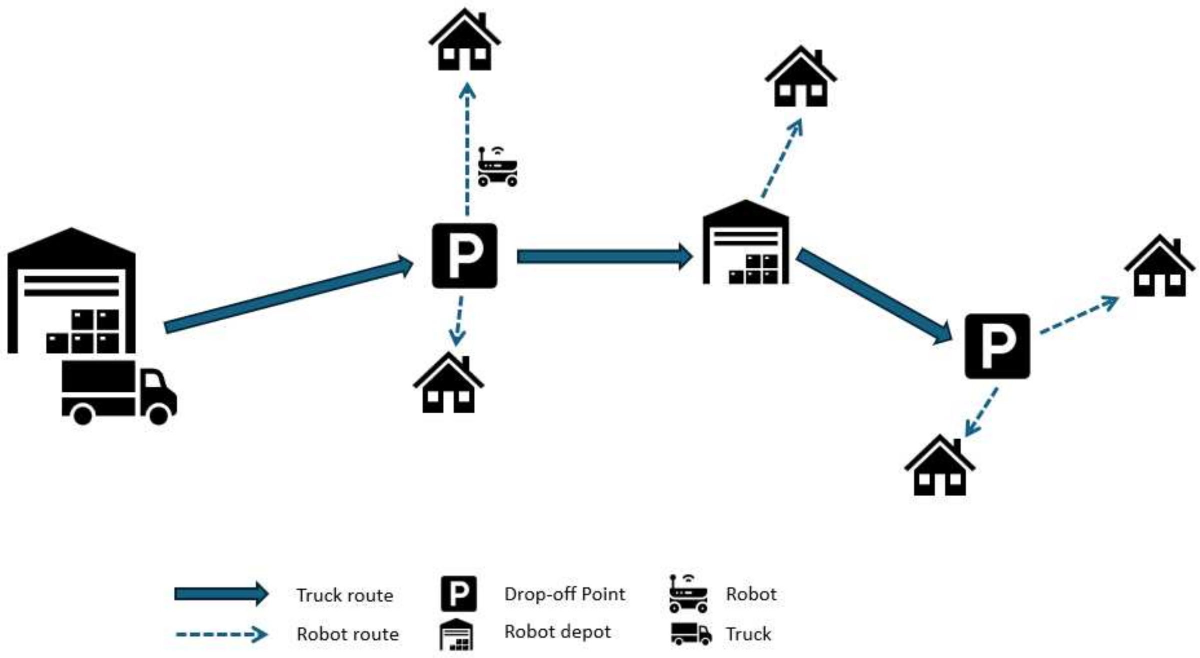

Figure 1: Graphical visualization of the model

3.3 Dataset

The data set used consists of various key elements, such as the customer, a start depot, depots to load new robots and drop-off points where the van only stops to unload assigned robots. The start depot serves as the beginning of the tour for the van where it is loaded with the parcels for the entire tour and the robots up to its maximum capacity. In addition to this, the designated depots along the route serve a dual purpose: they not only function as locations where new robots can be loaded but also act as points from which robots, when accompanied by the van carrying the necessary packages, can be dispatched for delivery. These depots, however, are solely for robot deployment and not warehouses, meaning the van with the packages is essential for the delivery process.

Furthermore, the dataset includes drop-off points where the van makes scheduled stops. At these locations, the robots assigned to specific delivery tasks are unloaded from the van. These robots, once deployed, proceed to make the final deliveries to the customers.

A total of five different instance sizes, each with 20 variants, are tested and analyzed. In each instance data set, the number of robot depots is 25 and the number of drop-off points is 48. In addition, the length of the square area in kilometers is identical for all variants. The five different instance sizes differ in the number of customers to be supplied, with 20 data sets with 25, 50, 75, 100 and 125 customers in each size.

The model contains various adjustable parameters, an important one of them is the maximum capacity of the robots that can be transported simultaneously in the van. Another relevant parameter is the distance-dependent cost, which is calculated per kilometer for the robots and the van and reflects the operational costs and energy consumption. These parameters are customizable and allow different scenarios to be tested and evaluated.

Another parameter is the time-dependent costs, which, in contrast to the distance-dependent costs, vary between the two means of transport, with vans accounting for the significantly higher costs per hour. The Euclidean distance is used to calculate the distances between the various points such as customers, depots and drop-off points, which are calculated between the coordinates stored in the data set.

In the data set used, the travel times for both vans and robots are stored in separate time matrices. These matrices contain the traveling time in seconds between the individual points and enable a realistic representation of the different travel speeds of the two vehicle types. These matrices form the basis for calculating the total delivery times and optimizing route planning in the model.

3.4 Algorithm

In the first step of the algorithm, only the truck tour is planned, whereby the capacity restrictions for the maximum number of transportable robots and the costs for the truck are not taken into account. The truck starts its tour at the initial depot, marking the beginning of the delivery process. All robot depots and drop-off points, i.e. all possible points that can be assigned to customers, are then evaluated for each customer in the current data set. This is done by calculating the delivery costs from the customers to the respective delivery points, whereby the delivery costs consist of the time and distance-dependent costs. For each customer, the point with the lowest delivery costs is identified and saved for later processing. In order to organize the van's route planning as efficiently as possible, the optimum route is calculated from the starting depot. This is done by using the list of optimum stops for each customer and always evaluating and carrying out the van's next stop using a scoring system. The distance from the truck's current position to all stops that have not yet been assigned is calculated and the stop with the highest score, i.e. the shortest distance from the current position, is assigned. After each assignment to the truck tour, the point and also the customer are marked as assigned and removed from the list. The score is then calculated again based on the new position of the truck and assigned until all the stops that are optimal for the respective customers have been assigned.

The next step of the algorithm is the scheduling of the robots on the initial truck tour that was obtained in the first step. The scheduling of the robots is a key component of the algorithm and includes several steps. In the first part, the scheduling examines the feasibility of the solution regarding the capacity constraints of the available robots. As one robot is needed per customer, the number is determined by the number of customers. The capacity is then added up for each depot along the current truck route and checked to see whether the number is theoretically sufficient to serve all customers. In cases where the available robots are insufficient to serve all customers, the function attempts to incorporate yet unused robot depots into the truck’s route to replenish the supply. The depot chosen is the closest one to the final stop of the current truck tour, calculated by the Euclidean distance between the depot and the stop. If no such depot is available, the function terminates, and the solution is deemed infeasible.

Once the overall feasibility has been determined, the next step is to calculate for each customer for which robot depot or drop-off point of the current truck tour the costs are lowest. The most cost-effective stop is then saved for each customer together with the associated costs, but the actual robot availability is not considered. Based on the previously calculated optimal stops for each customer, the customers are assigned to the points. It is then necessary to check whether the robot assignment to the stops is feasible. This is done by matching the available robots to the maximum capacity at the start of the truck route and then checking at each stop along the actual route whether it is a robot depot or a drop-off point. If it is a robot depot, the available robots are reset to the maximum capacity and if it is a drop-off point, the number of assigned robots is subtracted from the available robots. The initial solution is feasible if enough robots are always available to serve all customers and is saved for further optimization.

If there are not enough robots available and the solution is therefore not feasible, a regret insertion method is executed, and the customers are sequentially assigned. For this purpose, the costs for their optimum stop are updated for all customers until they have been assigned and it is checked whether capacity is available at this stop. The regret values, which is the difference in cost between their optimal stop and the second-best option, are then calculated for all customers. However, due to the logic of a consistent truck tour, only the stops that are located after the optimal stop in the truck tour are taken into account when calculating the second-best stop. The regret values are then ranked in descending order and the customer with the highest regret value is to be assigned to their best stop. Before allocation, a verification of the available robots at this stop is performed and only if there is sufficient capacity, the allocation is made, otherwise the customer will not be assigned. If this stop is a drop-off point, the available robots will be reduced by one for all preceding and succeeding drop-off points until a robot depot is encountered. If the customer has been successfully assigned and the capacities have been updated, the regret values are recalculated for the unassigned customers and the next assignment is tried until all customers have finally been assigned to a stop.

In the next part, the approach is to optimize the actual schedule by evaluating if swapping of two customers would lead to a better solution. The optimization process begins by identifying customers who are not assigned to their most optimal delivery stop. This involves iterating over all current stops on the truck tour and evaluating which customers are assigned to it. A comparison is then made for each customer as to which stop they are currently assigned to and what would be the most cost-effective stop for them. The calculation is based on the time-dependent robot costs and the distance costs. Any customer that is not assigned to their most cost-effective stop is stored on a separate list for later evaluation. Once all non-optimally assigned customers have been identified, a swap mechanism is used to try to improve the current solution, while the truck route remains unchanged and only the customers are swapped, which ensures that there are no issues regarding the availability of the robots. For each possible pair of non-optimally assigned customers, the extent to which the solution would change if they both swapped their assigned stops is calculated. If the swap results in lower total costs, the swap is carried out and the current planning is updated. This process is carried out for all possible pairs in this list and repeated until no further improvements can be achieved.

The aim of this procedure is to allocate customers to the stops on the truck tour as efficiently as possible.

Based on the previously created truck tour with the corresponding robot assignment, a local search procedure is then applied to further optimize the current best solution.

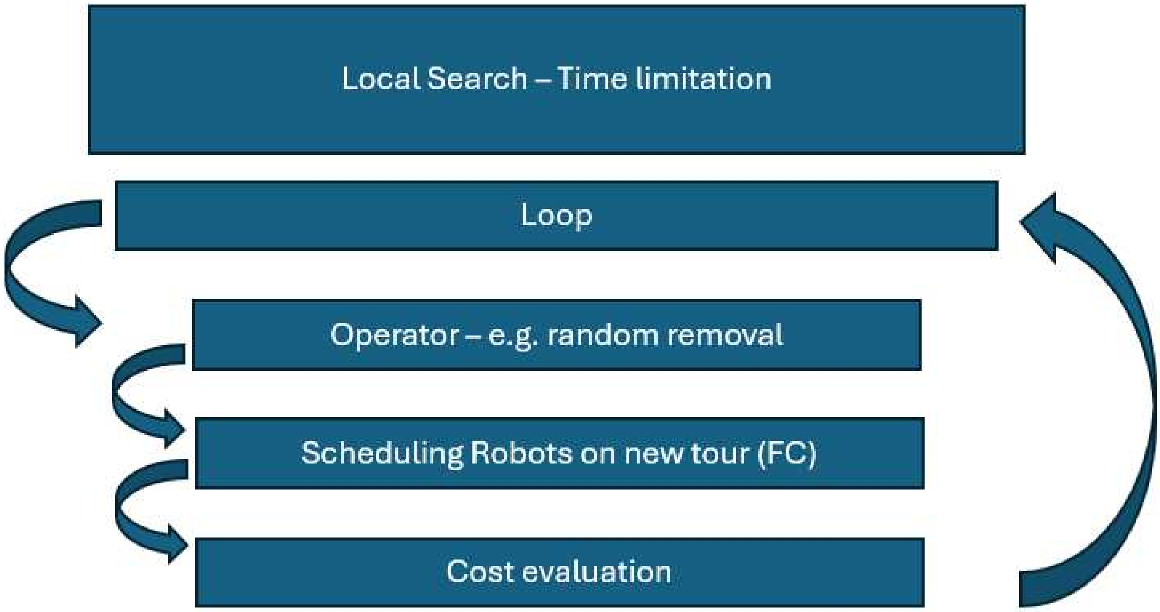

Figure 2: Functionality of the basic local search method

The local search algorithm is based on an outer and an inner loop to obtain the best possible result. The criterion for the duration of the local search procedure was tested with different options, such as a maximum number of iterations or a limited number of iterations without improvement. However, the tests showed that both the results and the comparability were best with time limits, so that each instance runs through the local search procedure for the same amount of time.

In each iteration of the outer loop, the algorithm attempts to explore a new solution by making a random modification, without considering any factors like feasibility or costs improvements, to the current truck route. This outer loop therefore uses an Iterated Local Search approach to explore new neighborhoods without considering the factors mentioned previously in order to avoid remaining in local optima. This swap mechanism only applies from the second iteration of the loop, as in the first iteration, the current best solution calculated in the first steps is to be used as the basis for the inner loop. From the second outer iteration onwards, two random stops of the truck tour are selected and their positions are swapped. The selection of the stop excludes the start depot, as this is the only non-exchangeable point on the truck tour. By swapping the positions of two stops on the truck tour, more diversification is to be integrated into the solution and a new sequence is to be tested. The algorithm also ensures that two different stops are always selected so that there is always a swap and a stop does not have to swap positions with itself, which does not lead to a new solution. This step ensures that the solution continues to evolve and the algorithm can investigate a variety of possible configurations in the search for an optimal or near-optimal solution. Once the truck's route has been modified by the swap operation, the algorithm proceeds to reschedule the robots by using the previously introduced method to schedule the robots on the modified tour and to account for the changes to be used in the inner loops. This swap of the outer loop is performed exactly once, except for the first iteration, before the inner loop is initiated.

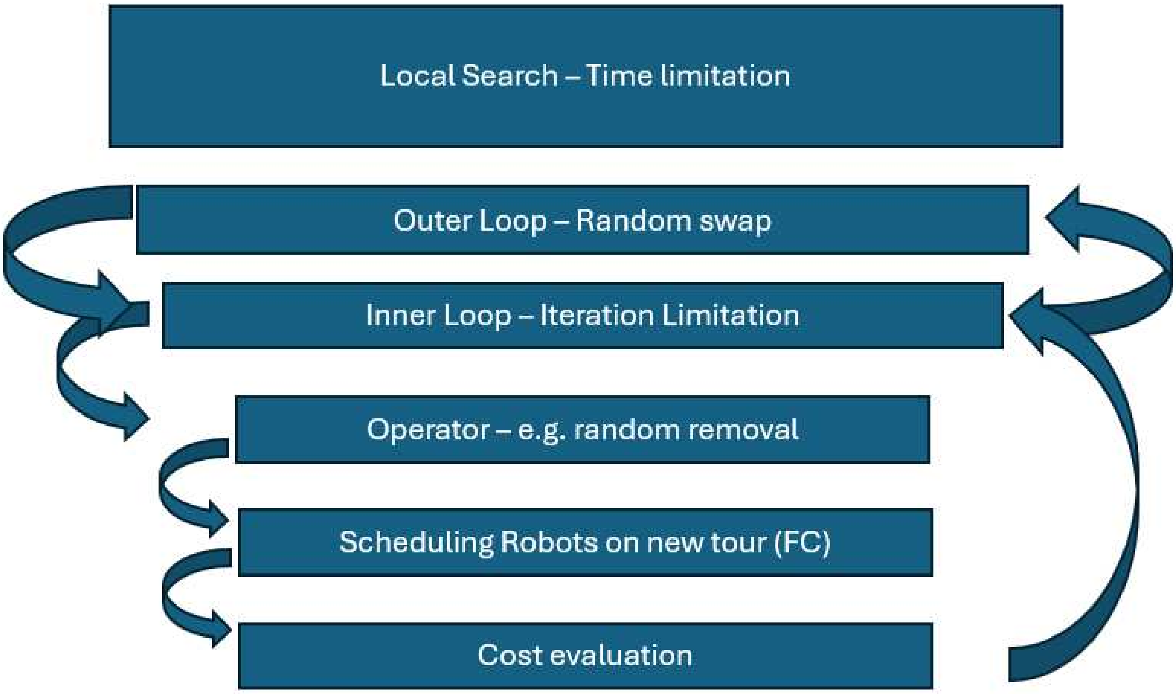

Figure 3: Functionality of the extended local search method

In the inner loop of the local search algorithm, the goal is to optimize the truck and robot schedules progressively by using different operators to improve the solution. This loop with a Variable Local Search approach will test and evaluate different configurations to find optimizations of the overall solution by reducing costs. The runtime of the inner loop is limited by a maximum number of iterations in order to fulfill the trade-off between computing time and solution quality. In addition, the outer loop with the random swap introduces more variability into the calculation, which should lead to a better solution quality. The new tour after the swap then runs through the inner loop again and this process is repeated until the time limit for the entire local search procedure is reached As soon as the time limit is reached, the loop is ended earlier. At the beginning of each inner loop, the current truck tour, the robot schedule and the associated costs are saved to ensure that if the iteration does not lead to an improvement, the solution can be restored for the next iteration.

In each iteration within the inner loop, a so-called operator is randomly selected from a set of five or seven, depending on the configuration, different operators. Each of these operators has different approaches to potentially improve the solution. The random selection of operators is intended to ensure that a large number of different changes are made in different orders to avoid potential local optima. Five of the operators are based on Ostermeier et al. (2022), namely the best-cost removal, the random swap, the random insertion, the random removal, the depot insertion and complemented by two more, the best-cost insertion and the best-cost swap.

If the algorithm selects the best-cost removal operator, it proceeds by attempting to improve the current solution by removing a stop from the truck's route. The basic idea is to remove a stop, reassign its associated robot tasks to other stops, and evaluate whether this results in a lower overall cost. When the first of the operators, the best-cost removal operator, is selected by the random generator, it tries to improve the current solution by removing a stop from the truck's tour. The basic principle of this operator is that a stop is removed and the robots/customers assigned to the stop are assigned to other stops and the impact on the solution is then evaluated.

For this purpose, each stop of the tour is temporarily removed one after the other and the respective robots are saved. To save computing time and increase efficiency, the entire robot scheduling method is not used after the stop has been removed, but the costs are only estimated. The estimated costs are calculated by taking into account the distances and times between the neighboring stops for the truck's new route, which are now directly connected after the removal. The costs of the robots are then only calculated for the two nearest stops.

If the evaluation reveals that the estimated cost of removing a stop results in an improvement, it is marked as a potential improvement and saved for later evaluation. After all stops have been evaluated, the initial route is restored with the associated robot assignments, as the current solution is only adjusted if an improvement in the overall solution is ensured. Now the stop with the highest estimated cost saving is removed and it is checked whether enough robots are available to ensure that this solution is feasible. If the solution is feasible, the robots are reassigned using the detailed method.

If this new solution is feasible and leads to a reduction in total cost, this is the new current solution and will be updated for the next iteration. If either is not the case, the initial solution is restored and the next iteration uses the same solution as a starting point.

The second of the operators adapted by Ostermeier et al. (2022) is the random swap operator, which is intended to introduce random perturbations into the current truck tour. In the first step of the operator, two random stops of this tour are selected by a random generator, whereby the first and last stops are excluded. It is also ensured that the random generator selects two different stops. As soon as the two random stops have been selected for this operator, the two positions within the truck tour are swapped.

As soon as the swap has been performed, the algorithm calculates the robot assignment using the entire robot assignment method, as in this case it only has to be calculated once and therefore the computing capacities need to be utilized less. This also ensures the feasibility of the new tour. In order to quantify the effects of the swap on the tour, the new total costs of this tour are calculated and compared with the previous best solution. If this new tour results in lower total costs, i.e. a better solution, the current solution is updated and this is the new solution that is used for the next steps. If the new tour does not lead to lower total costs, the previous tour is restored and continues to be used as the basis for the next steps.

If the third of the Ostermeier et al. (2022) based operators, the random insertion operator, is selected, a random stop should be selected and added to the current truck tour. This stop is randomly selected by a random generator, but with the condition that it is a stop that does not yet exist in the current tour. This ensures that a new additional stop is added, which can only be a depot or drop-off point. As soon as the valid stop has been identified, a random position within the current truck tour is selected by a random generator, whereby the first position is excluded due to the obligatory start depot. The stop is now added at the selected position and the new truck tour is constructed.

Once the stop has been added to the truck route, the robot scheduling is performed to adjust the robot assignments according to the modified truck route. This step again ensures that the assignment of customers/robots remains feasible and that the robots can be optimally assigned to the newly created tour. The total costs are then calculated and evaluated for this new composition of the truck tour with the associated planning of the robots. If these are lower than those of the previous best solution, a new current best solution is found and updated. If this is not the case, the previous solution is restored and the operator has not led to any improvement.

The fourth of the operators based on Ostermeier et al. (2022) is the random removal operator, which selects a random stop of the tour to remove. The selection by the random removal operator again excludes the first stop to fulfill the conditions. Once the stop has been selected, it is removed with the aim of seeing whether this can improve the overall solution.

As soon as the robots have been assigned and feasibility has been established, the total costs of the new current solution are calculated and evaluated again. These new total costs are again compared with those of the current best solutions. If they are lower, the current best solution is updated and removing the stop has improved the total cost and the operator has been successful. If the total costs are not lower, there has been no improvement and the previously saved best solution is still used as the basis.

The fifth operator based on Ostermeier et al. (2022) is the depot insertion operator. The aim of this operator is to change the overall structure and improve possible assignments, as the insertion of an additional depot is intended to increase robot availability during the truck tour. The position for the insertion is again determined by a random generator, whereby the position of the start depot is excluded as with the previous operators. The next step is to check which of the robot depots in the data set are not yet scheduled in the current tour. The deviation that occurs when this depot is inserted at the position is now calculated for each depot that has not yet been scheduled. This deviation is calculated on the basis of the distance between the predecessor and successor at this position and the depot itself. If the depot currently being evaluated has the best deviation found so far, it is labeled as the current best depot.

As soon as all potential depots have been evaluated, the depot with the highest deviation, in other words the highest cost saving, is selected. If a suitable depot is found, the truck tour is adjusted accordingly and the robots are assigned. If the total costs calculated on this basis are lower than those of the current best solution, a new best solution has been found and the tour and allocation re updated. If this is not the case, the initial solution is restored.

The best-cost insertion operator was developed to improve the performance of the Local Search section. It is an adaptation of the best-cost removal operator. The aim of this operator is to find a depot or drop-off point that is added in one place and reduces the total cost. For this purpose, all stops that are not yet present in the current tour are identified. The cost changes of the truck tour are then calculated for each possible position in the current truck tour for each stop that has not yet been assigned (depot or drop-off point) if this stop is added there. This deviation is determined by comparing the distances and travel times between the current stops before and after the insertion and factoring in the additional cost of including the new stop. Both the time and distance-dependent costs are taken into account for this calculation.

If the calculated deviation for the current stop at the position is less than the current best deviation, the stop and the position are saved and stored as the current best option. As soon as all stops on all positions have been evaluated, the best stop on the corresponding best position is inserted and forms the basis for further evaluation.

The robot schedule is recalculated and updated based on this new truck tour, which optimizes the planning for this tour and at the same time ensures that the tour is feasible. This method is also only applied at the end of the operator in order to keep the algorithm as efficient as possible and the optimizations only take place in the truck tour, on the basis of which the new optimal robot assignment is then made. The costs of this new solution are then recalculated and compared with the current best overall solution. If the total costs are lower using this depot insertion method, a new best solution has been found. If this is not the case, the initial solution from the beginning of the operator is restored and continues to serve as the basis for future calculations as long as the time limit for the local search method has not yet been reached.

The seventh operator, the best-cost swap operator, has also been developed to improve solution quality. It works according to a similar scheme to the best-cost insertion operator. In contrast to this operator, which attempts to reduce the overall costs by adding an additional stop, the best-cost swap operator works exclusively with the stops that are already present in the tour. The aim of this operator is to improve the efficiency of the tour by rearranging the stops rather than adding new ones.

For each possible pair of stops within the current truck tour, a temporary swap is performed and the deviation in time and distance-related costs is calculated. If the evaluation shows that swapping the positions of these two stops results in lower costs for the truck tour than the current best tour, this swap is marked as the current best swap. In this way, all swaps are gradually evaluated and as soon as they have all been evaluated, the best swap is executed.

After this swap has been carried out in the truck tour, the robot planning is carried out again using this tour. The new total costs are now calculated again for the new truck tour with the corresponding robot schedule and compared with the current best solution. If the newly found solution is better and has lower total costs, the current best solution is updated. If the total costs are not lower, the initial solution at the start of the operator is restored and the best solution remains unchanged.

Chapter 4: Results and Discussion

This section presents the results of the different variants and parameters of the investigations. The aim of this section is to evaluate the different scenarios and draw conclusions from them.

In order to improve the results, the local search method was tested to determine which runtime delivers the best ratio between runtime and quality of results. To evaluate the optimal computation time of the local search method, different runtimes were tested. This test covers different time intervals, from short runtimes of 60 seconds, which create the baseline for comparability, to long runtimes of 1800 seconds, and indicates how and whether the solution quality improves as a result of the changes. The change in solution quality was measured as a percentage improvement compared to the reference solution of 60 seconds, with negative values signaling an improvement. This section analyzes the results in detail in order to shed light on the influence of the local search duration on the solution quality and to derive practical recommendations for the choice of runtime.

Extending the runtime to 120 seconds already shows an initial significant improvement in the solution, which improves by -0.858%. This shows that the algorithm is able to identify relevant local minima in the solution space in a short time. The extension to 180 seconds leads to a further significant improvement to -2.801%, which represents an additional progress of about 1.943% compared to 120 seconds.

The results in this time window illustrate that the first few minutes of the local search method are particularly effective. Here, the algorithm utilizes the full capacity of local optimization by quickly refining solutions towards a better area of the solution space. It is particularly noteworthy that the majority of improvements take place within the first three minutes, which indicates a high level of efficiency in this phase.

The improvement in solution quality slowly begins to stagnate in the middle range of the tested runtimes. At a runtime of 240 seconds, a change of -6.022 % is achieved compared to the basic solution of 60 seconds. This is a significant increase compared to the shorter runtimes and shows that the additional time in this range still enables significant improvements. If the run time is increased further to 300 seconds, the solution improves to -6.136 % compared to the basic solution, but with a significantly lower progress of only around 0.114 % compared to the result of 240 seconds.

With an extension to 360 seconds, the change remains constant at -6.136% compared to 60 seconds, which shows that the algorithm does not achieve any further improvements here. This indicates that saturation is reached after 300 seconds, and no significant additional local minima are found. The results in this time frame illustrate the decreasing marginal benefit of longer runtimes. It is particularly important to emphasize that 300 seconds represents a clear limit beyond which further computing time no longer delivers any significant improvements.

With longer runtimes, the improvements are marginal and only justify the additional computing effort in exceptional cases. An extension to 600 seconds leads to a change of -6.292% compared to the basic solution, which is only an improvement of 0.156% compared to 360 seconds. At 900 seconds, the change is -6.362% to the solution of 60 seconds, which represents a minimal additional improvement of 0.070% compared to 600 seconds. Finally, the best solution is achieved at the maximum runtime of 1800 seconds with -6.472 % compared to the baseline solution, but the difference to 900 seconds is also very marginal at 0.11 %.

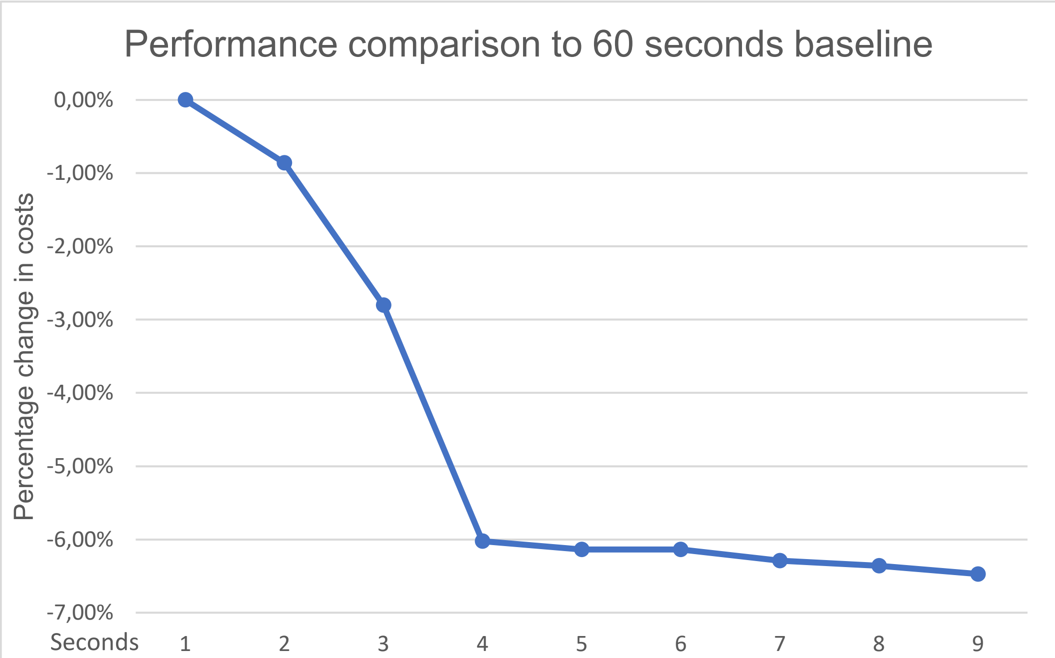

Figure 4: Performance comparison of different computing times

In conclusion, it can be stated that a runtime of 300 seconds is selected as the best runtime, as it guarantees a good solution quality with efficient use of resources. The analysis has shown that the algorithm largely exhausts its optimization potential in this period and achieves a significant improvement in the solution without a further extension of the computing time leading to significant progress. While longer runtimes such as 600 or 900 seconds still enable marginal improvements in solution quality, these minimal advances are not in proportion to the significant increase in computing time required.

Particularly noteworthy is the clear saturation effect that becomes recognizable from a runtime of 300 seconds. In this range, the improvement curve flattens out considerably, which indicates that the algorithm has already carried out the essential local optimizations and that additional computing time no longer brings any significant advantages.

The analysis of the number of iterations in the inner local search loop shows that it has a significant influence on the total costs. Different settings of iteration limits of 75, 150, 225 and 300 were tested, with the total costs at 300 iterations serving as a reference value. It is noticeable that lower iteration limits consistently lead to lower total costs, while the highest total costs occur at the highest number of iterations of 300. This result might be explained by the way the algorithm works, which runs through an outer loop after the iterations in the inner loop, in which random swaps are carried out. These random changes of the solution space seem to help the algorithm escape being locked into local optima and explore new areas of the search space, which improves the quality of the solutions.

Low iteration limits in the design of the algorithm mean that these random swaps are performed more frequently without regard to restrictions or evaluations, as the inner loop is completed faster and thus more outer iterations can be performed within the time limit. This greater weighting of the random element ensures more effective exploration of the solution space and reduces the likelihood of the algorithm getting stuck in local optima. As a result, the overall costs are lower with shorter iteration limits. An iteration limit of 75 results in a total cost reduction of -3.54% compared to the reference of 300 iterations within the inner loop. This effect is favoured by the high frequency of random swaps, which leads to a more intensive exploration of the search space as they do not take any restrictions into account. With 150 iterations, an even greater reduction in total costs of -4.31% is achieved compared to the 300 iterations. This shows that the optimization within the inner loop is given more space, while the outer loop is still called frequently enough to effectively avoid local optima.

With a further increase to 225 iterations within the inner loop, the deviation compared to the total costs at 300 iterations is reduced to -2.89%. This indicates that the longer runtime of the inner loop allows the algorithm to perform deeper local optimizations. At the same time, however, the frequency of random swaps in the outer loop is reduced, which limits the exploration of new solution spaces. This effect becomes particularly clear with the maximum number of iterations of 300: Here, the algorithm setting leads to the highest total costs, as the inner loop takes up the majority of the computing time and the outer loop with its random changes is used less frequently. While the algorithm should theoretically perform the deepest local optimization within 300 iterations and thus lead to a high degree of solution quality, this is counteracted by the rare frequency with which the random perturbation of the outer loop is used. The algorithm remains in local optima for longer and is less likely to find more favorable solutions in other parts of the search space, which leads to higher overall costs.

The results illustrate the trade-off between the depth of optimization in the inner loop and the frequency of random swaps in the outer loop. Shorter iteration bounds encourage greater exploration of the solution space through more frequent random changes, leading to lower overall costs. Longer iteration bounds allow for more detailed optimization within local domains, but limit the random component of the algorithm, which appears to increase the overall costs.

In summary, the analysis indicates that 150 iterations provide the optimal balance between intensive local optimization and the necessary random component for exploring the solution space. This combination enables the algorithm to both optimize efficiently within local areas and to regularly explore new solution spaces. This result emphasizes the importance of carefully balancing the number of iterations to ensure both low overall costs and a robust and effective solution.

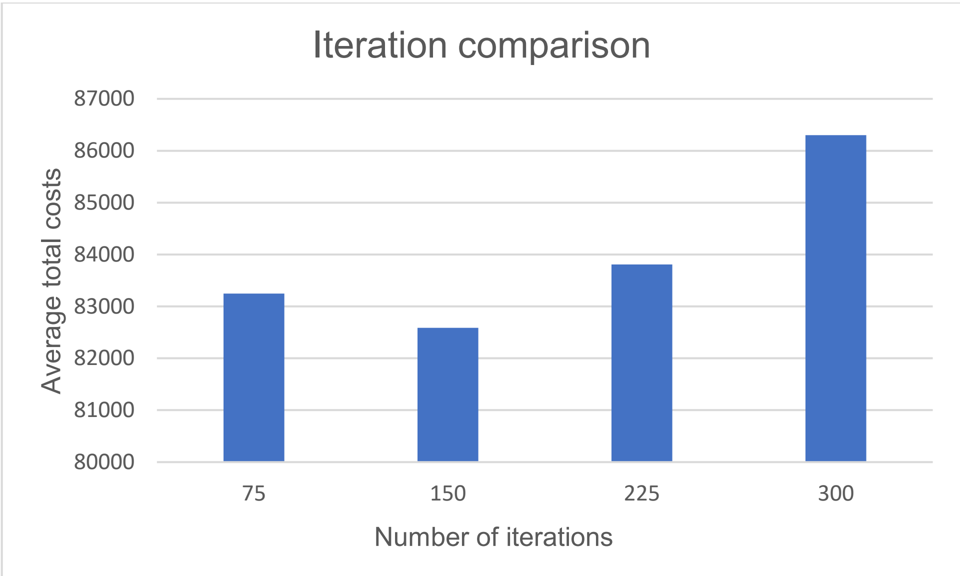

Figure 5: Performance comparison of iteration limits

Instance \ Iterations	75	150	225	300
C25	44218.2	44944.9	44913.6	43976.3
C50	65571	64071.5	65216	65882.8
C75	82382.3	82002.6	82909.4	85457.7
C100	102424.8	101974.3	102719.2	107448
C125	121628.7	119935	123277	128749.8

Table 1: Average costs per instance with different iteration limits

As part of the analysis of the influence of changes in truck and robot costs on total costs, various scenarios were tested in which the costs for the trucks and robots were each adjusted by 10%. The changes considered relate exclusively to the time costs, whereby the costs per kilometer were assumed to be the same for both vehicle types. The aim of this analysis was to evaluate the influence of these adjustments on the total costs.

First, the 10% increase in truck time costs was analyzed. This change led to an average increase in total costs of 1.69%. However, the effects varied depending on the size of the instance: while total costs rose by 0.92% in the smallest instance with 25 customers, the effect was most significant in the instance with 50 customers with an increase of 2.66%. In the larger instances with 75 and 100 customers, the increases were more moderate at 0.98% and 1.36% respectively but rose again in the instance with 125 customers and an increase in total costs of 2.53%.

These results lead to the conclusion that the effects of increasing time-dependent van costs do not increase linearly and evenly across all instances. Especially in the instances with 50 customers, the increasing time-dependent costs of vans seem to have a disproportionately strong effect on the total costs compared to the other sizes. This could indicate that in this scenario the van routes are particularly susceptible to additional time costs, possibly because there is a critical balance between the vans used and the available drop-off points or robot depots.

In the medium-sized instances, which consist of 75 or 100 customers, the increase in total costs is rather moderate. This suggests that the additional time-dependent van costs can be distributed more efficiently in this scenario, which can be explained, for example, by greater flexibility within route planning or more efficient planning options for the robots. Interestingly, the cost increase rises again to 2.53% for the largest instance with 125 customers, which indicates that transport routes become longer and more complex in very large scenarios and therefore an increase in van time costs becomes more significant again. This could be due to the fact that the high number of customers to be supplied with the same number of possible stops offers less scope for cost-efficient routes or that the robots can no longer absorb the rising transporter costs through more efficient planning, so that the rising time-dependent van costs have a more significant impact.

Overall, these results suggest that the sensitivity of total costs to rising truck time costs depends on the specific structure of the instance. In medium scenarios, efficiency gains could partially offset the negative effects, while in smaller and particularly large scenarios the cost increases are more significant.

In contrast, a 10% reduction in truck costs also led to an average reduction in total costs of 1.69%. The analysis of these results shows that a reduction in van time costs does not lead to equal savings in all instances. The comparatively significant average improvement in total costs of 3.39% within the smallest instance with 25 customers is worth noting here. This leads to the conclusion that in small scenarios the explicit van costs account for a comparatively high proportion of the total costs and the focus should be on reducing these. One possible explanation for this could be that in smaller instances, the number of journeys required between the drop-off points and depots is more limited and therefore a direct saving through reduced time costs is possible without major structural adjustments to route planning.

In the medium-sized to larger instances, however, the cost reduction is more moderate. While a significant saving of 2.43% was still recorded with 75 customers, the saving decreases continuously with increasing instance size: with 100 customers it is only 1.69% and with 125 customers only 0.9%. This could lead to the conclusion that as the number of customers increases, the influence of time-dependent van costs on total costs decreases, as the other costs in the cost structure become more relevant. The instance with 50 customers in which total costs remained almost unchanged at -0.02% should be highlighted here. This suggests that a reduction in van time costs hardly plays a role in this scenario because other factors dominate the costs. It is possible that this instance is in a kind of ‘transitional area’ in which the truck routes are already organized relatively efficiently and a reduction in time costs only has a marginal influence. Alternatively, it could be that there is a high degree of flexibility in route planning in this instance, so that the optimization potential through lower time costs has already been exhausted.

Overall, a reduction in van time costs leads to noticeable savings, particularly in small instances, while the effect decreases with increasing instance size. This could indicate that the time-dependent van costs account for a proportionately larger share of the total costs in smaller scenarios, with other factors accounting for a greater proportion of the cost structure in larger instances.

With a 10% reduction in robot time costs, the average saving in total costs was 0.86%, which is less than the effect of the changes in van time costs. The savings were relatively high in smaller instances in particular: total costs were reduced by 1.72% in the instance with 25 customers and by 1.57% in the instance with 75 customers. In the instance with 50 customers, on the other hand, the savings were only 0.6%, while there was virtually no change in the instance with 100 customers (+0.08%). In the largest instance, consisting of 125 customers, the cost savings were comparatively low with an average reduction of -0.47%. This result leads to the conclusion that a reduction in time-dependent robot costs plays a much more significant role in smaller instance sizes, as the relative importance of the robots is much more substantial here. In larger instances with more customers, the overall cost structure appears to be affected more considerably by other factors.

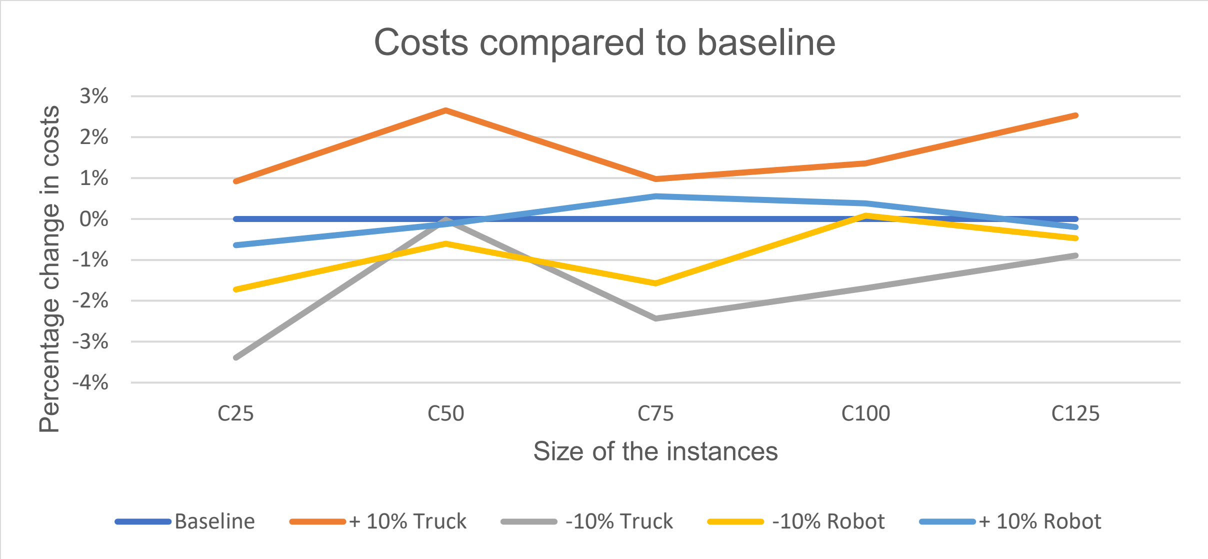

Figure 6: Effects of adjustments to cost parameters

Of particular interest is the result of a 10% increase in robot time costs, which, contrary to expectations, did not lead to a deterioration in the overall solution, but remained almost neutral with a minimal average improvement of -0.01%. The analysis of the individual instances shows that the effects vary below 1% in positive or negative deviation from the basic solution: While in the instance with 75 customers, an increase in robot time costs led to a slight increase in total costs of 0.56%, a reduction in total costs of 0.64% was observed in the smallest instance with 25 customers. The changes were also slightly negative in the instances with 50 and 125 customers (-0.12% and -0.2% respectively).

One possible explanation for this contradictory behavior lies in the limited interchangeability between robots and vans. Since the robots are only used after they have been transported by the vans, their routes and their influence on the total costs are heavily dependent on the positioning of the depots and drop-off points. An increase in robot time costs may force the optimization algorithm to adjust the routes of the transporter in such a way that a minimal increase in efficiency is achieved - even if the effect is only marginal. At the same time, a reduction in robot costs does not appear to offer much scope for improvement, as the existing solutions are already highly optimized and the main cost factors seem to apply by different factors.

To summarize, it can be observed that the robot time costs have a comparatively small influence on the total costs. While a reduction in robot time costs can still achieve moderate savings in smaller instances, this effect decreases significantly with increasing instance size. An increase in robot time costs, on the other hand, has minimal negative impact on overall costs and can even lead to a minimal improvement in some cases. This indicates that the robots in the hybrid delivery model are heavily dependent on the higher-level transport structures and that their costs cannot be considered in isolation.

The limited size of the data sets could also have an impact on the solution approach. If the number of available depots and drop-off points in certain regions is not very high, it is possible that the algorithm has already found a near-optimal solution, which means that a further reduction in robot costs will only have a minimal effect on overall costs. If the existing transfer points have been strategically chosen favorably, the balance between truck and robot use remains largely stable, regardless of a reduction in robot costs.

These factors could explain why a reduction in robot costs only results in a moderate reduction in overall costs, meaning that the positive effect of a cost reduction remains comparatively small, as the existing route planning is already highly optimized and offers only limited scope for further improvements.

The analysis shows that the costs per kilometer have a significantly greater influence on the total costs than the pure time costs of the trucks and robots. While a 10% reduction in robot time costs only resulted in an average saving of -0.86% and a comparable adjustment of time costs for trucks led to a change of ±1.69%, a 10% reduction in kilometer costs resulted in a disproportionately high reduction in total costs of -8.04%. This illustrates that the cost component resulting from the distance traveled has a significant influence on the overall cost structure.

This result is also underlined by the fact that it was tested in combination with how the total costs change if the truck costs are increased by 10% and the robot costs are reduced by 10%. This could be a realistic future scenario, assuming potentially increasing costs for the operation of vans due to environmental regulations or restrictions in urban areas and decreasing costs due to economies of scale or improvements in robot technologies. The evaluation of the results showed that under this assumption, the total costs were on average 2.1% higher across the five different instances.

Interestingly, this increase is higher than the isolated 10% increase in van time costs, which previously only resulted in an average increase in total costs of 1.69%. This indicates that the reduction in robot time costs cannot sufficiently compensate for the additional costs caused by the increased time costs of the vans. Analyzing the results at the instance level shows that this effect is of varying degrees. The highest cost increase was observed in the instance with 50 customers, where total costs rose by 3.06%. This could indicate that the dependency on van costs is particularly high in medium problem sizes, meaning that an increase in the cost of this component is more significant. In contrast, the largest instance with 125 customers shows the smallest relative change with a cost increase of only 1.42%.

In summary, the results show that measures to reduce kilometer costs have a significantly greater leverage effect on total costs than reducing the pure time costs of the vehicles used. While the time costs of the trucks have a slightly greater influence than the robot time costs on the total costs, their effect is still significantly less than that of the costs per kilometer.

Parameter / *Instance*	Baseline	+10% Truck	-10% Truck	+10% Robot	-10% Robot	+10% Truck/ -10% Robot
C25	44944.9	45358.2	43421.1	44656.8	44170.1	45739
C50	64071.5	65772.7	64058.2	63991.9	63685.2	66032.6
C75	82002.6	82804.7	80007.8	82458	80712.3	83677.6
C100	101974.3	103362.3	100252.8	102363.2	102058.2	104244.9
C125	119935	122973.5	118860.8	119695.9	119372.4	121634.8

Table 2: Average costs per instance with costs adjustments

The analysis of the results with different transport capacities of the vans shows that the maximum number of robots that can be transported simultaneously has a significant influence on the total costs.

The solution with a capacity of 12 robots was chosen as the reference point, as it is considered a realistic assumption for the problem. The changes to the maximum capacities have led to both improvements and deteriorations in the solutions. The effects differed significantly between the different instance sizes.

The reduction of the maximum possible capacity to 10 results in a slight increase in the average total costs across all instances compared to the basic solution. This indicates that a slight reduction in capacity has only a minimal impact on the overall efficiency of route planning. However, a closer look at the individual instances shows a possible dependence of the effects on the problem size. While the solution improved by -4.51% in the smallest instance with 25 customers, it deteriorated by +3.61% in the largest instance with 125 customers.

These divergent results could indicate that in smaller instances, reduced capacity leads to better utilization of the depots, as the vans must stop more frequently to load new robots. A possible advantage of this could be that not only are new robots loaded, but in the assumption of the model, robots can also be loaded directly at the depots with the parcels from the van and carry out the final delivery. This leads to the assumption that more frequent stops at depots may be strategically advantageous for smaller instances, as this potentially reduces the total distance traveled by the van. For larger instances with a higher number of customers, however, a reduced capacity could lead to longer transport routes for the vans, as more stops at the depots are required to load all the necessary robots. In addition, the number of drop-off points to be approached is comparatively high, which could lead to more inefficient routes, increasing overall costs.

Increasing the capacity to 14 led to an average improvement in total costs of -0.84% compared to the base solution with a maximum capacity of 12. This shows that a slight increase in the maximum capacity of the van can increases the efficiency of the tours, as the van can transport more robots at the same time, which in turn reduces the number of mandatory stops at robot depots and thus allows significantly more flexibility in tour planning. The effect also differs depending on the size of the instances. While the smallest instance with 25 customers has a 3.07% more expensive solution compared to the basic solution, the solution in the largest instance with 125 customers is improved by -2.28%. These results indicate that increased capacity is particularly beneficial in large instances, where a larger number of robots are required to ensure efficient delivery to customers. The more efficient utilization of existing drop-off points due to the increased flexibility could have a significant impact here, as fewer mandatory stops at fixed depot locations are required to bring all robots to their most optimal destinations.

Increasing the capacity further to 16 did not lead to an additional improvement, but to an average deterioration in total costs of 0.87% compared to the basic solution. This indicates that above a certain capacity limit, no further increases in efficiency can be achieved and that capacity extensions can even have detrimental effects. The second smallest instance size with 50 customers shows the largest deviation from the base solution with a significant deterioration and average cost increase of 4.5%. These results indicate that a further increase in the maximum possible capacity seems to lead to more inefficient route planning in the algorithm used in this thesis, as it seems to change the structure of route planning significantly. Especially in small and medium-sized instances, it can lead to a decrease in the number of routes traveled by the van, but a longer distance is covered, which affects the flexibility of the robot planning.

To summarize, it can be stated that the influence of the maximum capacity of the van has a significant impact on the overall costs of the solution, but that these effects are highly dependent on the size of the instance. Whilst a slight increase in maximum capacity to 14 can lead to an improvement in overall costs, the results show that a further increase can be inefficient. Similarly, reducing the capacity can be beneficial in smaller instances, while leading to a worse solution in larger instances.

Instance \ Capacity	10	12	14	16
C25	44944.9	45358.2	43421.1	44656.8
C50	64071.5	65772.7	64058.2	63991.9
C75	82002.6	82804.7	80007.8	82458
C100	101974.3	103362.3	100252.8	102363.2
C125	119935	122973.5	118860.8	119695.9

Table 3: Average costs per instance at different capacities of the truck

As part of this thesis, four different model variants were examined to analyze the efficiency and performance of the hybrid delivery concept with robots and vans. The approach was built on the Ostermeier et al. (2022) model, which was used as a basis and comparative solution. Building on this, step-by-step extensions were made in order to evaluate the influence of additional optimization mechanisms on the quality of the solution.

As outlined previously, the first extension consisted of adding two additional operators to the basic model with the aim of improving the solution quality through an extended search strategy. In a further variation, an outer loop was integrated into the local search algorithm, whereby the original five initial operators were initially retained. Finally, this method was also applied to the extended variant with the additional operators to analyze whether the combination of the outer loop with an extended search space offers further potential for optimization.

In the first step, Ostermeier's et al. (2022) reference solution, without an outer loop within the local search optimization and with the five initial operators (best-cost removal, random swap, random insertion, random removal, depot insertion), was compared with the solution of the extended model with the two additional operators (best-cost insertion, best-cost swap) without any further changes.

Although the two new operators are designed as best-cost operators and are therefore designed to achieve the greatest possible cost reduction for the current setting within the scope of their application, their integration nevertheless led to an average increase in total costs of 1.67% compared to the basic model according to Ostermeier et al. (2022). This contradicts the expectation that targeted optimization through best-cost strategies should lead to better solutions. One possible explanation for this could be that although the new operators make locally optimal decisions, this could increase the risk of getting locked into local optima.

A closer look at the individual instances shows that the deterioration in the first four problem sizes remains relatively constant. In the smallest instance with 25 customers, total costs increased by 1.86%, with 50 customers by 1.15%, with 75 customers by 1.18% and with 100 customers by 1.56%. These values are all in a comparably similar range, which indicates that the additional operators in these instances do not cause any drastic negative effects, but only slightly increase costs.

Only the largest instance with 125 customers shows a significantly higher cost increase of 2.61%. This could indicate that the additional operators in more complex problems increasingly lead to inefficient changes that impede the optimization process. This could be due to the fact that the original model already enables a balanced solution design in which different operators jointly develop a sensible structure of routes and allocations. The new best-cost operators could disrupt this existing balance by bringing about optimal but strategically unfavorable changes in the short term.

The type of new operator could also have an impact to the quality of the solution. Since best-cost insertion and best-cost swap specifically select the most cost-effective insertion or swap step, this may limit exploratory changes. The original model could possibly allow a greater variance in the changes, whereby solutions are found that are better in the long term, even if they initially appear less cost-efficient in individual steps.

To summarize, it can be stated that the integration of the two additional best-cost operators did not lead to an improvement in overall costs but instead caused higher costs on average. It is particularly noticeable that the cost increases are relatively even in most instances, but are particularly pronounced in the largest instance with 125 customers.

Variant / *Instance*	Ostermeier et al. (2022)	Two added Operator	Percentage change
C25	50939.3	51886.55	1.86%
C50	76498.2	77374.7	1.15%
C75	92726.4	93823.8	1.18%
C100	112936.1	114697.75	1.56%
C125	131136.5	134561.35	2.61%

Table 4: Average costs per instance Ostermeier et al. (2022) / two added operator variant

The extension of the model by introducing an outer loop with a random swap in an iterated local search / variable neighborhood search approach without any restrictions after 150 iterations led on average to a significant improvement of the solution compared to Ostermeier's et al. (2022) original model. Across all instance sizes, total costs were reduced by an average of -9.95%, indicating that this additional mechanism significantly improves solution quality and expands the solution space.

The more detailed analysis of the respective instance sizes shows that the most significant savings of -15.66% on average in total costs were achieved in the instance with 50 customers. The improvement was also significant in the instance with 75 customers, with a saving of -10.49%. In smaller and also the larger instances, a slightly lower performance could be identified compared to the medium-sized instances with 50 and 75 customers respectively, although significant savings could still be achieved compared to Ostermeier's et a. (2022) initial model. In the smallest instance that contains 25 customers, the total costs were reduced by -8.99% on average while the reduction in the largest instances containing 100 and 125 customers was -7.78% and -6.84% respectively.

These results indicate that the introduction of the iterated local search approach of the outer loop with random perturbation has a significant positive effect on the quality of the solution, which was most pronounced for medium-sized instances. This is probably since the original search strategy according to Ostermeier et al. (2022) tends to stagnate early in local optima in these scenarios. The random swap introduced, on the other hand, specifically promotes the diversification of the search and thus makes it possible to overcome inefficient solution structures. As a result, potentially better solutions can be found that would otherwise remain unattainable due to early convergence.

Although the analysis of the results shows that the outer loop of the local search algorithm with the random swap leads to a significant increase in performance in all instance sizes, the improvement is lower for the smallest and largest instances compared to the medium instances. This could be explained by the better performance of the solution quality of the reference model of Ostermeier et al. (2022) in these edge instances, which reduces the effectiveness of the additional exploratory search step. One possible hypothesis is that the solution space in the middle instances has a greater variety of potential solutions and thus offers more optimization potential.

To summarize, it can be said that the extension of the Ostermeier et al. (2022) model by an outer loop with an iterated local search approach using a random swap led to a significant improvement in solution quality. The cost savings were particularly significant for medium-sized instances, while they were still significant but lower for the smallest and largest instances. This indicates that the outer loop effectively prevents the solution from remaining in local optima and promotes the diversification of the solution space, although the effect varies depending on the instance size.

Instance \ Variant	Ostermeier et al. (2022)	Additional Outer Loop	Percentage change
C25	50939.3	46359.05	-8.99%
C50	76498.2	64516.35	-15.66%
C75	92726.4	82995.85	-10.49%
C100	112936.1	104150.2	-7.78%
C125	131136.5	122164.7	-6.84%

Table 5: Average costs per instance Ostermeier et al. (2022) / Outer Loop variant

In the extended variant of the local search algorithm with the outer loop including the random swap, the operators best-cost insertion and best-cost swap were also implemented. Compared to Ostermeier's et al. (2022) basic solution, there is a significant improvement in solution quality and an improvement compared to all previous variants. An average total cost reduction of -11.02% was achieved across all instances. This examines that the integration of the additional operators significantly increases the efficiency and effectiveness of the algorithm.

A detailed analysis of the individual instances illustrates that the greatest cost savings were again achieved in the instance with 50 customers, this time with a significant reduction of -16.86% compared to the basic solution. Significant improvements in the solution compared to the reference model were also achieved in the instances with 75, 100 and 125 customers respectively: A cost reduction of -11.17% was recorded for the instance with 75 customers, a reduction of -8.85% for the instance with 100 customers and a reduction of -9.23% for the instance with 125 customers. The smallest instance with 25 customers comprised an average cost reduction of -8.97% compared to the basic solution.

In the first of the tested variants, the addition of the two best-cost operators without the outer loop led to a deterioration in the solution quality. But the combination with the outer loop and the iterated local search approach containing the random swap reversed this effect and led to a significant improvement. This can be analyzed as follows: The outer loop with the random swap prevents the persistence in local optima favored by the best-cost operators by encouraging the algorithm to further explore the solution space through regular diversification of the search. The random swap of the outer loop without taking other factors into account compensates for the disturbance of the solution structure caused by the best-cost operators, so that the best-cost operators can once again maximize their advantages in the diversified search areas. The combination enables a more balanced exploration of the search space. The outer loop mitigates the negative effects of the best-cost operators in complex scenarios and improves the stability of the solution. The many iterations of the outer loop enable the best-cost operators to repeatedly optimize new, diversified solution areas in addition to the operators already given.

It has been proven that the combination of the outer loop with the additional best-cost operators is an effective method for significantly optimizing the solution quality of the entire algorithm. This combination enables a more efficient search of the solution space and leads to significant cost savings. The fact that the total costs could be reduced by an average of -11.02% compared to Ostermeier's et al. (2022) reference model emphasizes the effectiveness of this combination. In particular, the addition of the outer loop with the iterated local search approach of the random swap allows to overcome local optima and to explore the solution space. This seems crucial for the performance of this extended variant. The results confirm that this variation is a robust and powerful strategy for optimizing complex problem instances in the context of this hybrid truck/robot problem.

Variant / *Instance*	Ostermeier et al. (2022)	Additional Outer Loop + two Operator	Percentage change
C25	50939.3	46368.7	-8.97%
C50	76498.2	63600.45	-16.86%
C75	92726.4	82369.7	-11.17%
C100	112936.1	102944.4	-8.85%
C125	131136.5	119037.75	-9.23%

Table 6: Average Costs per instance Ostermeier et al (2022) / Additional Outer Loop + two operator variant

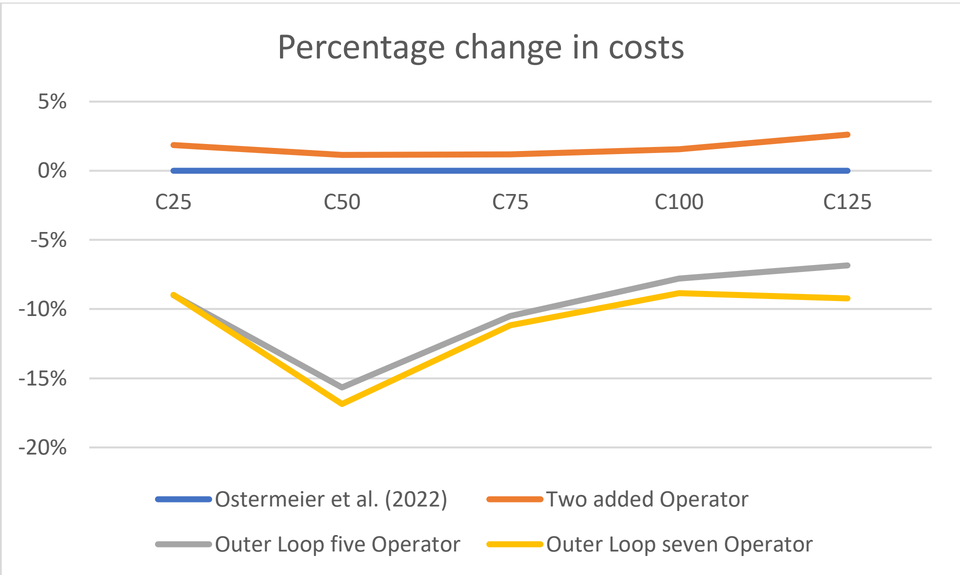

Figure 7: Performance comparison of model variants

Chapter 5: Conclusion

The last mile represents both a planning and a financial challenge for the logistics industry. Due to the constantly growing B2C e-commerce sector and increasing service requirements, the optimization of these processes will become even more important in the future. (Statista 2024; Hübner, Kuhn, and Wollenburg 2016; Vanelslander, Deketele, and Van Hove 2013)

In this thesis, a hybrid delivery model consisting of vans and autonomous delivery robots was analyzed. Based on Ostermeier's et al. (2022) model, an Iterated Local Search / Variable Local Search (ILS-VLS) approach was used to improve the solution quality. This heuristic algorithm was implemented in C++ and the effects of various parameters and approaches were evaluated in numerical experiments. The extent to which the different runtimes of the local searches affect the solution quality is shown and the trade-off between solution quality and runtime is illustrated. In addition, it is shown to what extent the number of iterations within the inner loop of the variable local search approach until the perturbation of the outer loop is performed with the iterated local search approach has an influence on the solution quality. Both investigations were used to improve the quality of the solution and thus reduce the total costs of the objective.

The key results of the numerical experiments can be summarized as follows:

- The effects of the individual parameters vary depending on the size of the data sets analyzed
- Time-dependent costs for vans have a stronger influence on the total costs than the time-dependent costs of the robots
- Distance-dependent costs have a stronger influence on the total costs than time-dependent costs
- Increasing the transport capacities for robots in vans does not necessarily lead to an improved solution; the effects are strongly dependent on the instance size
- The extended model of this work obtains on average 11.02% better results than the original basic model of Ostermeier et al. (2022)

These results illustrate the complexity of hybrid delivery networks and show that precise parameterization is essential for the efficiency of such a system.

This work extends the existing research approaches in the comparatively newer area of hybrid delivery concepts, in this case hybrid van/robot delivery. The use of an ILS-VLS approach to solve the problem shows that these heuristic methods are well suited for complex optimization problems and can also be used promisingly for hybrid delivery systems and improve the solution qualities. In addition, the results on the cost structure and influences of the van's capacities provide implications regarding which factors could be focussed on in the context of practical application.

The results of this thesis are subject to limitations that must be taken into account when interpreting them. Due to the lack of time windows for the final delivery to the customer and the lack of consideration of the process after the delivery of the robots, practical challenges of the process and the cost structure may not be fully mapped.

In addition, it can be assumed that the assumption of the same distance-related costs for robots and vans does not correspond to practice and could therefore provide scope for more practical results in future studies.

For future research, the use of a more complex data set could also be considered, in which, for example, real road maps are taken into account and possibly also include dynamic traffic situations in which the robots could utilize their advantages in urban areas.

It would also be interesting to investigate the extent to which the autonomous driving of both the vans and the robots has an influence on the problem. The research approach with the additional combination with drones such as (Moeini and Salewski 2020) is also a promising approach.

Appendix

Instance	Ostermeier et al. (2022)	Ostermeier el al. (2022) + 2 Operator	Outer Loop 5 Operator	Outer Loop 7 Operator
C100_000	129267	121955	112770	116808
C100_001	116710	117107	105985	104868
C100_002	103375	104422	101697	97132
C100_003	113014	112629	103804	101536
C100_004	101563	99029	94809	90817
C100_005	116490	124440	107671	106475
C100_006	109889	109744	109321	105246
C100_007	101417	110520	98955	94967
C100_008	117332	126926	93021	89768
C100_009	120538	126150	109490	111825
C100_010	110268	111152	105439	104082
C100_011	118099	114956	92578	92868
C100_012	103369	104390	101515	102928
C100_013	121800	120366	109060	112730
C100_014	112122	113856	111326	108077
C100_015	104997	122515	98728	97820
C100_016	119823	112966	107596	106223
C100_017	124666	109685	102945	103818
C100_018	105466	110842	109040	105441
C100_019	108517	120305	107254	105459
C125_000	133049	133792	122917	118385
C125_001	128567	134900	119954	119375
C125_002	130460	126595	119892	118648
C125_003	131208	151407	132694	131494
C125_004	123653	150474	120024	118705
C125_005	131153	133663	114109	111565
C125_006	130362	124925	125823	121077
C125_007	137859	133878	125606	120631
C125_008	137933	134998	128216	129388
C125_009	114419	114402	112269	110082
C125_010	126426	127455	104952	104392
C125_011	120572	120385	112077	107725
C125_012	131822	139242	131021	126114
C125_013	126927	139474	131525	126248
C125_014	137696	145264	122414	119471
C125_015	124887	130956	121779	118646
C125_016	145706	144074	132614	128537
C125_017	137870	137690	103306	99716
C125_018	131505	135139	131969	127182
C125_019	140656	132514	130133	123374
C25_000	57350	56249	51303	49490
C25_001	51774	55965	49825	49029
C25_002	50359	50217	44892	45061
C25_003	54431	54559	48839	49736

C25_004	53066	53219	46190	47006
C25_005	40032	49548	37033	36405
C25_006	52365	52748	41573	41224
C25_007	56428	56125	50798	50970
C25_008	53877	50894	44885	46337
C25_009	34944	44018	34191	34191
C25_010	51365	51365	51569	51365
C25_011	48409	48409	37498	40895
C25_012	52988	52383	50170	50050
C25_013	49841	49660	46411	46411
C25_014	50993	51144	50849	50190
C25_015	53378	55018	52240	51967
C25_016	59666	59666	55302	54996
C25_017	49552	49378	40432	40141
C25_018	46494	45944	45875	46207
C25_019	51474	51222	47306	45703
C50_000	75301	77434	61476	60696
C50_001	78397	82820	64545	62874
C50_002	80074	72139	68480	61370
C50_003	73338	74686	59364	59241
C50_004	80909	80867	64810	64869
C50_005	75351	82062	68845	65456
C50_006	76088	77797	67045	68029
C50_007	70410	72806	62566	62664
C50_008	80303	77510	63732	61400
C50_009	80936	82815	74427	74578
C50_010	77088	81064	69958	67700
C50_011	70128	75400	62687	61711
C50_012	83657	78045	59915	59915
C50_013	87572	74869	70303	68733
C50_014	76289	76033	78138	76700
C50_015	67197	69612	62813	63000
C50_016	80363	90269	63587	63256
C50_017	61135	67001	42772	46530
C50_018	74228	72981	66912	65856
C50_019	81200	81284	57952	57431
C75_000	95154	95703	73260	73815
C75_001	84645	81580	82693	81684
C75_002	92662	96005	81734	81432
C75_003	89433	89413	78889	78505
C75_004	93026	105211	84321	85783
C75_005	91665	85858	81466	83349
C75_006	91104	95585	94670	90682
C75_007	86218	100828	86907	85817
C75_008	103656	94180	96163	90815
C75_009	84808	87105	67738	68144
C75_010	91308	88089	83034	81458

C75_011	98441	85704	77227	75856
C75_012	92681	106631	91004	90689
C75_013	96387	101793	76201	75744
C75_014	83011	72242	67078	67470
C75_015	92971	94200	92216	91120
C75_016	96594	90373	89461	87747
C75_017	108485	107815	83272	85682
C75_018	85548	94339	81609	81916
C75_019	96731	103822	90974	89686

Table 7: Results of all instances in each tested model variant

List of figures

List of tables

References

Agatz, Niels, Paul Bouman, and Marie Schmidt. 2018. ‘Optimization Approaches for the Traveling Salesman Problem with Drone’. *Transportation Science* 52 (4): 965–81. https://doi.org/10.1287/trsc.2017.0791.

Alfandari, Laurent, Ivana Ljubić, and Marcos De Melo Da Silva. 2022. ‘A Tailored Benders Decomposition Approach for Last-Mile Delivery with Autonomous Robots’. *European Journal of Operational Research* 299 (2): 510–25. https://doi.org/10.1016/j.ejor.2021.06.048.

Bakach, Iurii, Ann Melissa Campbell, and Jan Fabian Ehmke. 2021. ‘A Two-tier Urban Delivery Network with Robot-based Deliveries’. *Networks* 78 (4): 461–83. https://doi.org/10.1002/net.22024.

Bányai, Tamás, Béla Illés, and Ágota Bányai. 2018. ‘Smart Scheduling: An Integrated First Mile and Last Mile Supply Approach’. Edited by Dimitris Mourtzis. *Complexity* 2018 (1): 5180156. https://doi.org/10.1155/2018/5180156.

Bosona, Techane. 2020. ‘Urban Freight Last Mile Logistics—Challenges and Opportunities to Improve Sustainability: A Literature Review’. *Sustainability* 12 (21): 8769. https://doi.org/10.3390/su12218769.

Boysen, Nils, Stefan Fedtke, and Stefan Schwerdfeger. 2021. ‘Last-Mile Delivery Concepts: A Survey from an Operational Research Perspective’. *OR Spectrum* 43 (1): 1–58. https://doi.org/10.1007/s00291-020-00607-8.

Boysen, Nils, Stefan Schwerdfeger, and Felix Weidinger. 2018. ‘Scheduling Last-Mile Deliveries with Truck-Based Autonomous Robots’. *European Journal of Operational Research* 271 (3): 1085–99. https://doi.org/10.1016/j.ejor.2018.05.058.

Buldeo Rai, Heleen, Sara Verlinde, Jan Merckx, and Cathy Macharis. 2017. ‘Crowd Logistics: An Opportunity for More Sustainable Urban Freight Transport?’ *European Transport Research Review* 9 (3): 39. https://doi.org/10.1007/s12544-017-0256-6.

Chung, Sung Hoon, Bhawesh Sah, and Jinkun Lee. 2020. ‘Optimization for Drone and Drone-Truck Combined Operations: A Review of the State of the Art and Future Directions’. *Computers & Operations Research* 123 (November):105004. https://doi.org/10.1016/j.cor.2020.105004.

Clarke, G., and J. W. Wright. 1964. ‘Scheduling of Vehicles from a Central Depot to a Number of Delivery Points’. *Operations Research* 12 (4): 568–81. https://doi.org/10.1287/opre.12.4.568.

Dai, Hongyan, Jiawei Tao, Hai Jiang, and Weiwei Chen. 2019. ‘O2O On-Demand Delivery Optimization with Mixed Driver Forces’. *IFAC-PapersOnLine* 52 (13): 391–96. https://doi.org/10.1016/j.ifacol.2019.11.156.

Daxbot. 2024. ‘Dax Delivery Robots’. 5 September 2024. https://daxbot.com/dax-delivery-robots.

Demir, Emrah, Yuan Huang, Sebastiaan Scholts, and Tom Van Woensel. 2015. ‘A Selected Review on the Negative Externalities of the Freight Transportation: Modeling and Pricing’. *Transportation Research Part E: Logistics and Transportation Review* 77 (May):95–114. https://doi.org/10.1016/j.tre.2015.02.020.

Edwards, J.B., A.C. McKinnon, and S.L. Cullinane. 2009. ‘Carbon Auditing the “Last Mile”: Modelling the Environmental Impacts of Conventional and Online Non-Food Shopping’. Logistics Research Centre School of Management and Languages Heriot-Watt University.

Eksioglu, Burak, Arif Volkan Vural, and Arnold Reisman. 2009. ‘The Vehicle Routing Problem: A Taxonomic Review’. *Computers & Industrial Engineering* 57 (4): 1472–83. https://doi.org/10.1016/j.cie.2009.05.009.

Figliozzi, Miguel A. 2017. ‘Lifecycle Modeling and Assessment of Unmanned Aerial Vehicles (Drones) CO 2 e Emissions’. *Transportation Research Part D: Transport and Environment* 57 (December):251–61. https://doi.org/10.1016/j.trd.2017.09.011.

Figliozzi, Miguel, and Dylan Jennings. 2020. ‘Autonomous Delivery Robots and Their Potential Impacts on Urban Freight Energy Consumption and Emissions’.

Transportation Research Procedia 46:21–28. https://doi.org/10.1016/j.trpro.2020.03.159.
Garland, Max. 2023. 'FedEx and Amazon Still Haven't Figured out Sidewalk Delivery Robots. Will Mass Adoption Ever Come?' 12 March 2023. https://www.supplychaindive.com/news/fedex-roxo-amazon-scout-delivery-robots-mass-adoption/646702/.
Gerrits, Berry, and Peter Schuur. 2021. 'Parcel Delivery For Smart Cities: A Synchronization Approach For Combined Truck-Drone-Street Robot Deliveries'. In *2021 Winter Simulation Conference (WSC)*, 1–12. Phoenix, AZ, USA: IEEE. https://doi.org/10.1109/WSC52266.2021.9715506.
Golden, B. L., T. L. Magnanti, and H. Q. Nguyen. 1977. 'Implementing Vehicle Routing Algorithms'. *Networks* 7 (2): 113–48. https://doi.org/10.1002/net.3230070203.
Goodchild, Anne, and Jordan Toy. 2018. 'Delivery by Drone: An Evaluation of Unmanned Aerial Vehicle Technology in Reducing CO 2 Emissions in the Delivery Service Industry'. *Transportation Research Part D: Transport and Environment* 61 (June):58–67. https://doi.org/10.1016/j.trd.2017.02.017.
Hoffmann, Thomas, and Gunnar Prause. 2018. 'On the Regulatory Framework for Last-Mile Delivery Robots'. *Machines* 6 (3): 33. https://doi.org/10.3390/machines6030033.
Hossain, Mokter. 2023. 'Autonomous Delivery Robots: A Literature Review'. *IEEE Engineering Management Review* 51 (4): 77–89. https://doi.org/10.1109/EMR.2023.3304848.
Hübner, Alexander Hermann, Heinrich Kuhn, and Johannes Wollenburg. 2016. 'Last Mile Fulfilment and Distribution in Omni-Channel Grocery Retailing: A Strategic Planning Framework'. Edited by Neil Towers and Herbert Kotzab. *International Journal of Retail & Distribution Management* 44 (3). https://doi.org/10.1108/IJRDM-11-2014-0154.
Kapser, Sebastian, and Mahmoud Abdelrahman. 2020. 'Acceptance of Autonomous Delivery Vehicles for Last-Mile Delivery in Germany – Extending UTAUT2 with Risk Perceptions'. *Transportation Research Part C: Emerging Technologies* 111 (February):210–25. https://doi.org/10.1016/j.trc.2019.12.016.
Letchford, A N, J Lysgaard, and R W Eglese. 2007. 'A Branch-and-Cut Algorithm for the Capacitated Open Vehicle Routing Problem'. *Journal of the Operational Research Society* 58 (12): 1642–51. https://doi.org/10.1057/palgrave.jors.2602345.
Lim, Stanley Frederick W.T., Xin Jin, and Jagjit Singh Srai. 2018. 'Consumer-Driven e-Commerce: A Literature Review, Design Framework, and Research Agenda on Last-Mile Logistics Models'. *International Journal of Physical Distribution & Logistics Management* 48 (3): 308–32. https://doi.org/10.1108/IJPDLM-02-2017-0081.
Lourenço, Helena R., Olivier C. Martin, and Thomas Stützle. 2010. 'Iterated Local Search: Framework and Applications'. In *Handbook of Metaheuristics*, edited by Michel Gendreau and Jean-Yves Potvin, 146:363–97. International Series in Operations Research & Management Science. Boston, MA: Springer US. https://doi.org/10.1007/978-1-4419-1665-5_12.
Macioszek, Elżbieta. 2018. 'First and Last Mile Delivery – Problems and Issues'. In *Advanced Solutions of Transport Systems for Growing Mobility*, edited by Grzegorz Sierpiński, 631:147–54. Advances in Intelligent Systems and Computing. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-62316-0_12.
Mangiaracina, Riccardo, Alessandro Perego, Arianna Seghezzi, and Angela Tumino. 2019. 'Innovative Solutions to Increase Last-Mile Delivery Efficiency in B2C e-Commerce: A Literature Review'. *International Journal of Physical Distribution & Logistics Management* 49 (9): 901–20. https://doi.org/10.1108/IJPDLM-02-2019-0048.
Mladenović, N., and P. Hansen. 1997. 'Variable Neighborhood Search'. *Computers & Operations Research* 24 (11): 1097–1100. https://doi.org/10.1016/S0305-0548(97)00031-2.
Moeini, Mahdi, and Hagen Salewski. 2020. 'A Genetic Algorithm for Solving the Truck-Drone-ATV Routing Problem'. In *Optimization of Complex Systems: Theory, Models, Algorithms and Applications*, edited by Hoai An Le Thi, Hoai Minh Le, and Tao Pham

Dinh, 991:1023–32. Advances in Intelligent Systems and Computing. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-21803-4_101.
Mokhtarian, Patricia L. 2004. ‘A Conceptual Analysis of the Transportation Impacts of B2C E-Commerce’. *Transportation* 31 (3): 257–84. https://doi.org/10.1023/B:PORT.0000025428.64128.d3.
Murray, Chase C., and Amanda G. Chu. 2015. ‘The Flying Sidekick Traveling Salesman Problem: Optimization of Drone-Assisted Parcel Delivery’. *Transportation Research Part C: Emerging Technologies* 54 (May):86–109. https://doi.org/10.1016/j.trc.2015.03.005.
Ostermeier, Manuel, Andreas Heimfarth, and Alexander Hübner. 2022. ‘Cost-optimal Truck-and-robot Routing for Last-mile Delivery’. *Networks* 79 (3): 364–89. https://doi.org/10.1002/net.22030.
Pani, Agnivesh, Sabya Mishra, Mihalis Golias, and Miguel Figliozzi. 2020. ‘Evaluating Public Acceptance of Autonomous Delivery Robots during COVID-19 Pandemic’. *Transportation Research Part D: Transport and Environment* 89 (December):102600. https://doi.org/10.1016/j.trd.2020.102600.
Plank, Martin, Clément Lemardelé, Tom Assmann, and Sebastian Zug. 2022. ‘Ready for Robots? Assessment of Autonomous Delivery Robot Operative Accessibility in German Cities’. *Journal of Urban Mobility* 2 (December):100036. https://doi.org/10.1016/j.urbmob.2022.100036.
Punakivi, Mikko, Hannu Yrjölä, and Jan Holmström. 2001. ‘Solving the Last Mile Issue: Reception Box or Delivery Box?’ *International Journal of Physical Distribution & Logistics Management* 31 (6): 427–39. https://doi.org/10.1108/09600030110399423.
Ranieri, Luigi, Salvatore Digiesi, Bartolomeo Silvestri, and Michele Roccotelli. 2018. ‘A Review of Last Mile Logistics Innovations in an Externalities Cost Reduction Vision’. *Sustainability* 10 (3): 782. https://doi.org/10.3390/su10030782.
Repoussis, P.P., C.D. Tarantilis, O. Bräysy, and G. Ioannou. 2010. ‘A Hybrid Evolution Strategy for the Open Vehicle Routing Problem’. *Computers & Operations Research* 37 (3): 443–55. https://doi.org/10.1016/j.cor.2008.11.003.
Rodrigues, Thiago A., Jay Patrikar, Natalia L. Oliveira, H. Scott Matthews, Sebastian Scherer, and Constantine Samaras. 2022. ‘Drone Flight Data Reveal Energy and Greenhouse Gas Emissions Savings for Very Small Package Delivery’. *Patterns* 3 (8): 100569. https://doi.org/10.1016/j.patter.2022.100569.
Savelsbergh, Martin, and Tom Van Woensel. 2016. ‘50th Anniversary Invited Article—City Logistics: Challenges and Opportunities’. *Transportation Science* 50 (2): 579–90. https://doi.org/10.1287/trsc.2016.0675.
Srinivas, Sharan, Surya Ramachandiran, and Suchithra Rajendran. 2022. ‘Autonomous Robot-Driven Deliveries: A Review of Recent Developments and Future Directions’. *Transportation Research Part E: Logistics and Transportation Review* 165 (September):102834. https://doi.org/10.1016/j.tre.2022.102834.
Starship Technologies. 2024. ‘Starship Robots – Your Local, Community Helpers. Accessed On’: 14 August 2024. https://www.starship.xyz/the-starship-robot/.
Statista. 2024. ‘Umsatz Im Markt Für E-Commerce Weltweit in Den Jahren 2020 Bis 2029’. May 2024. https://de.statista.com/prognosen/484763/prognose-der-umsaetze-im-e-commerce-markt-in-der-welt.
Ulmer, Marlin W., and Sebastian Streng. 2019. ‘Same-Day Delivery with Pickup Stations and Autonomous Vehicles’. *Computers & Operations Research* 108 (August):1–19. https://doi.org/10.1016/j.cor.2019.03.017.
Van Loon, P., A.C. McKinnon, L. Deketele, and J. Dewaele. 2014. ‘The Growth of Online Retailing: A Review of Its Carbon Impacts’. *Carbon Management* 5 (3): 285–92. https://doi.org/10.1080/17583004.2014.982395.
Vanelslander, Thierry, Lieven Deketele, and Dennis Van Hove. 2013. ‘Commonly Used E-Commerce Supply Chains for Fast Moving Consumer Goods: Comparison and Suggestions for Improvement’. *International Journal of Logistics Research and Applications* 16 (3): 243–56. https://doi.org/10.1080/13675567.2013.813444.

Yu, Shaohua, Jakob Puchinger, and Shudong Sun. 2022. 'Van-Based Robot Hybrid Pickup and Delivery Routing Problem'. *European Journal of Operational Research* 298 (3): 894–914. https://doi.org/10.1016/j.ejor.2021.06.009.