MSc Industrial Engineering and Management Thesis

Design of an autonomous last-mile delivery system for small parcels

Distribute

UNIVERSITY OF TWENTE.

Design of an autonomous last-mile delivery system for small parcels

A master's Thesis in the field of Industrial Engineering and Management

How to design an autonomous B2C last mile delivery system, using the Campus of the University of Twente as a case study?

Author

M.H. Schotman (Marijn) – s2175541
MSc Industrial Engineering & Management

University of Twente

Drienerlolaan 5
7522 NB, Enschede
The Netherlands

Supervisors University of Twente

First Supervisor: Prof. Dr. Ir. M.R.K. Mes (Martijn)
Faculty of Behavioural Management and Social Sciences
Dep. Industrial Engineering and Business Information Systems

Second Supervisor: Dr. Ir. M. Koot (Martijn)
Faculty of Behavioural Management and Social Sciences
Dep. Industrial Engineering and Business Information Systems

Distribute

Hengelosestraat 527
7521 AG, Enschede
The Netherlands

Company Supervisor

Dr. Ir. B. Gerrits
Director of Distribute

Preface

Dear reader,

In front of you lies my master's thesis: "*Design of an autonomous last-mile delivery system for small parcels"*. This research was conducted at Distribute in Enschede and marks the conclusion of my master's studies in the field of industrial engineering and management.

First and foremost, I want to thank my company supervisor Berry Gerrits, for the opportunity to execute this research at Distribute and for his continuous support throughout the process. From start to finish, I was given the freedom to shape the direction of this thesis according to my own interests, and I could always count on Berry's guidance.

I also want to thank former employee Robert Andringa for his helpful input and assistance. Additionally, I am grateful to my fellow graduate students at Distribute, Yoran Nijenhuis and Maarten Ipskamp, for the constructive peer feedback and interesting discussions we shared along the way.

Furthermore, I would also like to thank my UT supervisor Martijn Mes for his significant help throughout this research. I could always turn to him for questions and help. I also want to thank my second supervisor, Martijn Koot, for improving my thesis with his valuable feedback towards the end of my thesis.

Finally, I would like to thank my friends and family for their support over the last year, both personally and academically.

I hope you enjoy reading this thesis

Marijn Schotman

Weerselo, 2025

Management Summary

This thesis is conducted at Distribute and investigates the design of an autonomous business-to-consumer last-mile delivery system, using the campus of the University of Twente as a case study. The current last-mile delivery system is approaching a point that it will no longer be sustainable due to various economic, social, and environmental challenges. Autonomous delivery presents a potential solution by eliminating delivery driver personnel costs, achieving more accurate delivery time windows through local depots and further reducing emissions. However, despite this potential, it is unclear how to implement such a system. This is the main goal of the thesis which is why we formulated the following main research question:

How to design an autonomous Business-2-Consumer last mile delivery system, using the Campus of the University of Twente as a case study?

The research starts by conducting an extensive literature review. We study the current state of last-mile logistics, types of autonomous delivery vehicles (ADVs), existing models and optimization techniques, simulation methods, and key performance indicators (KPIs) used for evaluation.

Building on the knowledge we gained from literature, we develop a generic design framework for designing autonomous last-mile delivery systems. This supports system designers in selecting and combining design elements. This framework does not attempt to cover every operational detail. Instead, it focuses on strategic, system level design choices that must be addressed before implementation. Key design choices included geographical and infrastructure factors, demand estimation methods, depot setup, and fleet compositions. This stage will result in a set of feasible delivery system alternatives.

To answer this question, we introduce a structured decision-making framework that supports system designers in selecting and combining design elements. This framework does not attempt to cover every operational detail. Instead, it focuses on strategic, system level design choices that must be addressed before implementation.

We develop a simulation model to test the performance of different system configurations under varying levels of demand, time windows, and fleet compositions. Key Performance Indicators are defined in the economic, environmental, and social domains, and experiments are conducted to evaluate the system's behavior.

Finally, the experimental results are analyzed, allowing us to draw conclusions about the performance of the proposed system and to develop recommendations for future use.

Key Findings Case Study

There are many examples of autonomous delivery vehicles, but we found that these can be divided into three general categories (See Figure 1):

- **Unmanned Aerial Vehicles (UAVs)**, often referred to as drones, are small, unmanned aircrafts without a pilot on board. Due to their limited size, they are not capable of delivering large, heavy or multiple packages (at once). However, their advantage lies in their ability to travel through the air, allowing them to bypass traffic. This makes them suitable for fast deliveries and are an excellent supplement for delivering parcels once the demand increases.
- **Sidewalk Autonomous Delivery Vehicles (S-ADVs)** are small, ground-based autonomous robots with a low maximum speed and a limited capacity (typically one). They drive preferably on sidewalks or through small pedestrian areas so they can bypass the public road. This makes them suitable for short-distance deliveries to customers, especially when these customers are inaccessible for the other vehicles.
- **Road Autonomous Delivery Vehicles (R-ADVs)** are larger, ground-based autonomous robots that function as a mobile parcel locker. With their significantly higher capacity and often ability to drive on public roads, they are suitable for delivering large volumes.

Figure 1: Examples Autonomous Delivery Vehicles: R-ADV (Express Robot) on the left, S-ADV (Starship Robot) in the center and UAV (Zipline Drone) on the right

A series of experiments were conducted to test the system under different configurations, including variations in fleet composition (homogeneous vs heterogeneous), customer demand levels (43, 106, 192), and delivery constraints (09:00-17:00 delivery, 09:00-21:00 and strict morning/afternoon time windows of 09:00-13:00 & 13:00-17:00).

Recommendation for the University of Twente Campus

For campus deliveries, we recommend not to implement deliveries with strict time windows. Since it would be way too costly to achieve this compared to delivery without time windows. For an autonomous delivery system with full day (09:00-17:00) delivery to customers at their home, we recommend using a heterogeneous fleet of **one R-ADV and two UAVs**. This heterogeneous fleet can deliver up to around 150 parcels a day between 09:00-17:00 (covering days with low demand, average demand and even slightly higher than average demand).

On a lower demand day, the single R-ADV or the two UAVs can deliver the parcels homogeneously (each with around 55% utilization). On an average demand day, a single R-ADV and UAV can deliver these packages with 91% utilization. On days of high demand, we recommend using an option of evening delivery, which allows this fleet to deliver 192 parcels with a utilization of 83%. If evening delivery is not an option, the heterogeneous fleet needs an additional **two drones** to cover the high demand days.

Contents

Preface ii
Management Summary iii
List of Figures viii
List of Tables ix
Chapter 1 Introduction and Methodology 1
1.1 Problem Description 1
1.1.1 Problem Context and Research Motivation 2
The Assignment 2
1.2 Research Design 4
1.2.1 Research Scope 4
1.2.2 Research objectives and research questions 4
Chapter 2 Literature Review 7
2.1 Parcel Delivery Process in the Netherlands (PostNL) 7
2.2 Advantages and Disadvantages of the current landscape 8
2.3 Autonomous Delivery Solutions 10
2.3.1 What is an automated delivery vehicle (ADVs) 10
2.3.2 Classification of autonomous delivery Vehicles 10
2.3.3 Legal Restrictions 14
2.4 Autonomous Delivery Systems 15
2.5 Modelling and Problem Formulations of Autonomous systems 16
2.5.1 Foundational routing problems 16
2.5.2 S-RADVs Models 17
2.5.3 UAV Models 18
2.5.4 R-RADVs models 19
2.6 Approaches to solve vehicle routing problems 19
2.7 Evaluating the solution 21
2.8 Conclusion 23
Chapter 3 Framework for Design of an Autonomous delivery System 24
3.1 Geographical and Infrastructure Considerations 27
3.1.1 Defining the area 27
3.1.2 Defining Delivery Locations 28
3.1.3 The delivery network 29
3.2 Demand of the area 30
3.3 Depot Characteristics and Placement 30
3.4 Operational Logistics 31

3.4.1 Vehicle characteristics 32
3.4.2 Vehicle allocation: Homogeneous vs. Heterogeneous Fleet 32
3.4.3 Routing 33
3.5 Evaluating the model 34
3.5.1 The model 34
3.5.2 Conceptual Model 35
3.5.3 Assumptions of the model 37
3.6 Conclusion 38
Chapter 4 Application of the framework to the UT campus 39
4.1 Geographical and Infrastructure Considerations 39
4.1.1 Defining the area 39
4.1.2 The delivery locations 39
4.1.3 The delivery network 40
4.2 Customer demand 41
4.3 Depot Characteristics and Placement 42
4.4 Operational Logistics 43
4.4.1 Vehicle Characteristics (Fixed) 43
4.5.2 Vehicle allocation (Experimental) 44
4.5.3 Routing 44
4.6 KPI Selection 46
4.7 Evaluating the model 46
4.7.1 Inputs of the simulation model 47
4.7.1 Validity of the simulation model 48
4.8 Conclusion 51
5 Experiments 52
5.1 Experimental Settings 52
5.1.1 Experimental factors 52
5.1.2 Ranges and Combinations 52
5.1.3 Warmup & Replications 53
5.2 Experimental Results 53
5.2.1 Phase 1 53
5.2.2 Phase 2 55
5.2.3 Phase 3 58
5.2.4 Phase 4 60
5.3 Conclusion 63
5.3.1 Economic, Social and Environmental Assessment 64

5.3.1 Recommendations for the University of Twente Campus 65
Chapter 6 Conclusion and Recommendations 67
6.1 Most important Findings and Recommendations 67
6.2 Limitations 69
6.3 Future Research 70
Bibliography 71
Appendix A) Detailed KPI selection 81
Economic dimension 81
Environmental dimension 81
Social/Societal Dimensions 82
Appendix B) Latitude and Longitude polyline campus 83
Appendix C) Extra demand calculations 85
Appendix D) Vehicle Characteristics 88
Appendix E) Model Documentation 89
Appendix F) Detailed Experiments Section 96
Phase 1 96
Phase 2 100
Phase 3 102
Phase 4 104

List of Figures

Figure 1: Examples Autonomous Delivery Vehicles: R-ADV (Express Robot) on the left, S-ADV (Starship Robot) in the center and UAV (Zipline Drone) on the right iv
Figure 2: Schematic Representation of the First, Middle and Last Mile in the Supply Chain (Maersk, 2025).... 1
Figure 3: Map of the University of Twente (Campus | Universiteit Twente, n.d.). 3
Figure 4: 6 Levels of Autonomy (Bruneteau & Bruneteau, 2022) 10
Figure 5: Sidewalk ADVs (Starship robots), taken from 11
Figure 6: Nuro R3 (Left), Neolix (middle) and Express Robot in Duinrell (left); retrieved from: (Wessling, 2022), (Neolix, n.d.) and (MacroStep BV, 2025).... 12
Figure 7: Udelv transporter, retrieved from: (Transporter | Udelv, n.d.) 12
Figure 8: Left: Amazon Prime's delivery UAV (Chen, 2023b), Right: Ziplines UAV (Zipline Fact Sheet | Zipline UAV Delivery & Logistics, n.d.).... 13
Figure 9: Classification of routing literature. Adjusted from Srinivas et al. (2022).... 16
Figure 10: Traveling Salesman Problem (a) and Vehicle Routing Problem (b)--an illustration. Retrieved from (Liu et al., 2014) 17
Figure 11: Overview of heuristics, taken from (F. Liu a et al., 2023) 21
Figure 12: Examples of polygon (left) and multipolygons (right), taken from (Understanding Difference Between Polygon and Multipolygon for Shapefiles in QGIS?, n.d.).... 27
Figure 13: Residential buildings (blue) and campus facility buildings (red), with delivery network plotted on the right 40
Figure 14: Location of new Depot/"old Garage" (Yellow).... 42
Figure 15: Algorithm Solomons Nearest Neighbor.... 45
Figure 16: Distribution of waiting times 48
Figure 17: Automated Order Screenshot pt1 49
Figure 18: Automated Order Screenshot pt2. 49
Figure 19: Delivery Route to 20 random customers (left) vs Delivery route to 100 random customers (right).... 50
Figure 20: Different Delivery Route example 51

List of Tables

Table 1: Overview of possible KPIs 22
Table 2: Design Choice: Defining the area 28
Table 3: Design Choice: Delivery Locations 28
Table 4: Design choice: Delivery network 29
Table 5: Design Choice: Demand of the area 30
Table 6: Design Choice: Depot Strategy 31
Table 7: Example of ADV characteristics 32
Table 8: Design Choice: Fleet configuration 32
Table 9: Design Choice: Routing strategy 33
Table 10: Median, Average and Range of parcels delivered per day 42
Table 11: Light Electric Vehicle Characteristics 43
Table 12: Final KPIs 46
Table 13: Results 20 vs 100 customers 50
Table 14: Experimental Factors 52
Table 15: Initial experiment homogeneous fleet R-ADVs (average demand) 54
Table 16: Number of R-ADVs required for average demand 54
Table 17: Low demand R-ADVs (regular day, no time windows) 54
Table 18: High demand R-ADVs (regular day, no time windows) 54
Table 19: Conclusion Phase 1 55
Table 20: Heterogeneous Fleet Low Demand (Regular Day, No TimeWindows) 56
Table 21: Heterogeneous Fleet Average Demand (Regular Day, no TimeWindows) 56
Table 22: Heterogeneous Fleet High Demand (Regular Day, no TimeWindows) 57
Table 23: Performance electric van regular days 57
Table 24: Performance low Demand Heterogeneous & Homogeneous fleet with night delivery 58
Table 25: Performance Average Demand Heterogeneous and Homogeneous fleet (Night Delivery, No TW) 58
Table 26: Performance High Demand Homogeneous and Heterogeneous Fleet (Night Delivery, No TW) 59
Table 27: Performance Electric Van Night Delivery 59
Table 28: Conclusion Phase 3 (Number of ADVs) 59
Table 29: Performance Low Demand Homogeneous (Regular Day, TW) 61
Table 30: Performance Average Demand Homogeneous (Regular Day, TW) 61
Table 31: Performance High Demand Homogeneous (Regular day, TW) 61
Table 32: Conclusion Phase 4 (Number of ADVs) 62

Chapter 1 Introduction and Methodology

In this thesis, we investigate how autonomous delivery concepts might be able to address the problems with the existing last-mile delivery system. Section 1.1 provides a description of the problem, which entails a definition of the term last-mile delivery, accompanied by some background information, a description of the assignment, the problem context and the research motivation. Section 1.2 outlines the research design, which entails the research scope and the research objective with the devised research questions. Additionally, the deliverables, data collection methods and contributions of the study are described. Finally, the chapter ends with Section 1.3 containing the thesis outline and readers' guide.

1.1 Problem Description

Logistics covers the complete process of storing, coordinating, and transporting resources to their destination. In logistics management, the last-mile delivery refers to the final stage of the delivery process wherein a parcel is transported from a distribution hub to its destination. Some definitions are more precise than others. According to Boysen et al. (2020), the last mile delivery refers to "logistics activities associated with delivering shipments to private customer households in urban areas," while Vakulenko et al. (2019) defines last mile delivery as "delivery from the final upstream point of shipment to the end consumer". Ha et al (2022) encompasses the need for a more rounded definition of last mile delivery, emphasizing the delivery element while encompassing every form of delivery:

"The last transportation of a consignment in a supply chain from the last dispatch point to the delivery point where the consignee receives the consignment."

To give some extra context to this definition, despite being labeled as the "last mile," it may not necessarily be a mile. In fact, it can encompass various distances and utilize any mode of transportation, including bicycles, buses, cars, scooters, and even Unmanned Aerial Vehicles (UAVs). The final dispatch point will most likely be a distribution center or warehouse, but it could also be a store that directly ships the products. The delivery point could be the customer, but it can also include reception points or collection points. Figure 2 shows an example (simplified) overview of an entire supply chain and its components (Ha et al., 2022).

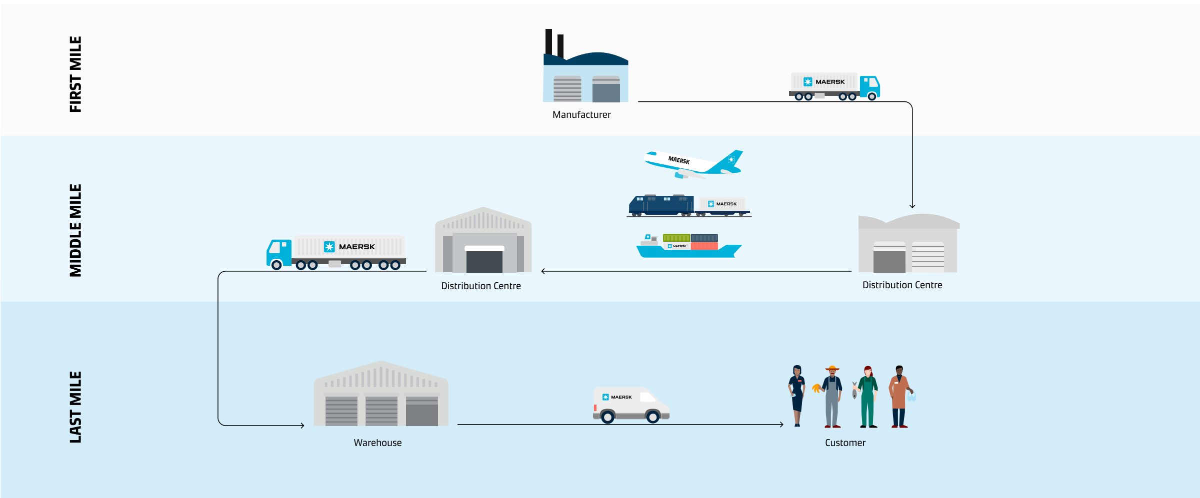

Figure 2: Schematic Representation of the First, Middle and Last Mile in the Supply Chain (Maersk, 2025)

1.1.1 Problem Context and Research Motivation

Allen et al. (2018), Buldeo et al. (2022a), and Kahl (2020) highlight that the last-mile delivery is under increasing pressure due to various challenges. These include increased customer expectations, seasonal peaks in demand, reduced lead times, meeting delivery time windows, a high rate of first-time delivery failure rates, and a growing number of product returns. Some aspects of last-mile delivery are stochastic, such as seasonal peaks, product returns and traffic congestion on roads. Other factors add more constraints to deliveries like tight delivery time windows and lowering energy consumption. Projections indicate that without significant change, delivery-related emissions in the top 100 global cities will increase by 31% in 2030, while traffic congestion could rise by 21% (World Economic Forum, 2020). Simultaneously, the number of delivery vehicles may increase by 36% (Kahl, 2020).

Centraal Bureau voor de Statistiek (2022a) projects that especially large and medium-sized cities will continue to grow in the Netherlands, as will several municipalities around the major cities. The World Economic Forum (2020) claims that urban areas already account for approximately 70% of global emissions, and delivery vehicles represent a disproportionately high share compared to personal vehicles. In Amsterdam, for example, one in every eight vehicles in the inner city is a truck or delivery van. According to Femke Halsema, the Mayor of Amsterdam, urban deliveries cause structural problems in the city of Amsterdam: Many old bridges and quays are not designed for the weight of vehicles and intensive use these days. Ultimately, urban last-mile delivery will increase by 78% by 2030. Mobility experts argue that the limiting factor in the future of urban mobility will not even be affordability, but land.

In addition to environmental and logistical pressures, human labor remains a major bottleneck. Last-mile delivery is relatively labor-intensive, with much of the driver's time spent driving rather than on value-adding tasks (Thomas & Tokar, n.d.). Moreover, human-operated delivery is restricted by infrastructure limitations, such as traffic regulations, parking access, working hours and, not to forget, human error. These constraints make it difficult to scale the current delivery system to meet future demands.

The Assignment

The current last-mile delivery system is under growing pressure due to mobility, sustainability and labor challenges mentioned in the previous subsection. Urban areas are experiencing increased congestion and rising delivery demand. At the same time, there is a clear need to reduce emissions and improve the efficiency of logistics systems. These challenges highlight the need to rethink existing last-mile logistics models and investigating alternative solutions.

Autonomous delivery systems, whether on the ground or in the air, offer a potential solution. By removing the dependency on human drivers and introducing new forms of mobility, these technologies could improve efficiency, reduce emissions, and reduce the pressure on the urban infrastructure. However, despite their potential, the question remains: *How should autonomous delivery systems be designed and evaluated to ensure they are suitable and effective in real-world settings?*

The goal of this thesis is to provide a structured approach for designing and assessing the suitability of autonomous last-mile delivery systems, with a focus on urban environments. While the case study provides a concrete example, the approach and insights are designed with a broader applicability in mind.

By selecting the University of Twente campus as a case, the research can examine autonomous delivery within a realistic and contained environment. The campus is located between the cities of Enschede and Hengelo in the Twente region of the Netherlands. Offering as much as a small town with 3,000 student houses and flats with plenty of facilities, such as a supermarket, bar, gym, general practitioner and even a hairdresser (*Campus | Universiteit Twente*, n.d.). The following list of reasons is why we chose for the campus as a case:

- Like a small town, the campus has people living and working in the same area.
- A university has a much higher chance of agreeing to potential testing of autonomous delivery, making it much more worthwhile to design an autonomous delivery system.
- Since the campus shares characteristics with a small town, insights gained from research here could be scaled up to larger urban applications

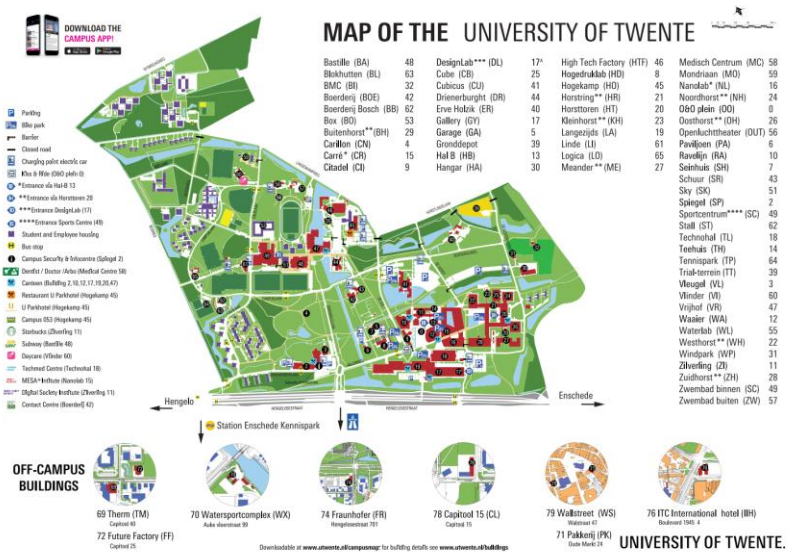

Figure 3: Map of the University of Twente (Campus | Universiteit Twente, n.d.).

This master thesis was conducted at a company called Distribute. Distribute is a young, relatively small, University of Twente Spin-off. Distribute designs and simulates distributed planning and control systems for the logistics and transport sector. Distribute is keen to keep up with developments in technology, aiming to better understand the current last-mile delivery infrastructure and the potential autonomous options for last mile delivery.

This essentially comes down to addressing and solving the following problem:

The current last mile delivery system is approaching a point where it will no longer be sustainable due to various challenges. Autonomous delivery presents a potential solution, but it remains unclear what approach we should use to design and implement this, particularly when considering a case like the University of Twente Campus.

This thesis has two goals:
1) Develop a generic framework for the design and evaluation of autonomous delivery systems
2) Demonstrate the application of this approach through a real-world case
While the outcomes are based on the University of Twente campus, the resulting framework could be applicable to other campuses or controlled urban environments. Future research is needed to further explore its broader applicability.

1.2 Research Design

In this section, we present the design of our research. We discuss the research scope, after which we describe the research objective and research questions. We then outline the data collection methods and contributions of the study. The subsection ends with a thesis outline for the reader.

1.2.1 Research Scope

In this thesis, we explore the field of last-mile delivery with a focus on societal, environmental and business factors. Furthermore, this thesis focuses on the case of the campus of the University of Twente. The campus is used as a base that should provide insight into how autonomous delivery can best be applied to (similar) areas such as the campus. Last-mile delivery can be split up into Business to Consumer (B2C) and Business to Business (B2B). The scope of this thesis lays in the B2C model rather than the B2B model.

Consumers demand increasingly faster, more reliable, and environmentally friendly delivery options. B2C also has a much larger range of different places and is in general less standardized. This causes more and more problems for the B2C model, while the B2B model is much more focused on 'fixed' retail orders with a regular partner. This process could also be more efficient and environmentally friendly in many ways, but for this thesis we still focused on the B2C model.

1.2.2 Research objectives and research questions

The main research objective of this thesis can be formulated as follows:

How to design and evaluate an autonomous B2C last mile delivery system, using the Campus of the University of Twente as a case study?

Several research questions are devised to help us answer the objective of the thesis. For each question, a small description is given about the approach.

Stage 1. Literature review

The thesis starts with a literature review to provide a strong foundation for designing an autonomous last-mile delivery system. The review focuses on understanding the current state of last-mile logistics, the potential of autonomous delivery technologies, and relevant modeling and simulation methods. The theoretical foundation is essential for developing and evaluating delivery system configurations in later stages of the research.

RQ. What does existing literature reveal about the technologies, models and methods available for designing and evaluating autonomous last-mile delivery systems?

- **a)** How are the first, middle, and last-mile stages of parcel delivery organized in practice in the Netherlands?
- **b)** What does the literature say about the current last-mile delivery landscape and its advantages and disadvantages?
- **c)** What types of autonomous delivery vehicles exist, and what are their legal restrictions, opportunities and limitations?
- **d)** What autonomous delivery systems and mathematical models are proposed in the literature?
- **e)** What approaches exist in the literature to solve vehicle routing problems?
- **f)** What simulation methods are used to analyze last-mile logistics systems?
- **g)** What Key Performance Indicators (KPIs) are used in the literature to assess the performance of a delivery system?

The answers to these sub-questions are given in Chapter 2 and form the theoretical foundation of the thesis.

Stage 2. Design Framework with Simulation Model
Building on the knowledge gathered from the literature, the second stage of the study focuses on creating an autonomous last-mile delivery system and presents the simulation model developed for evaluation. The framework is structured as a “menu of choices”, helping system designers make informed decisions about service area boundaries, demand patterns, depot placement, and operational logistics. These design choices are not only theoretically grounded but also serve as direct inputs to the simulation model. The chapter also outlines the structure and capabilities of the simulation tool used to assess delivery performance.

RQ. How can an autonomous B2C last-mile delivery system be designed, using a set of configurable design choices, based on the characteristics of a specific environment?

- **a)** What geographical and infrastructural factors must be considered?
- **b)** How can delivery demand be estimated and integrated into the system design?
- **c)** What depot configuration best supports the expected delivery flows?
- **d)** How should operational logistics be structured?
- **e)** How can simulation be used to assess and validate the performance of a proposed design?

Chapter 3 presents a structured framework for designing autonomous delivery systems and describes how this framework is operationalized within a simulation environment.

Stage 3. Case study: Applying the Framework to the University of Twente
In this stage, the system design framework and simulation model are applied to the specific use case of the University of Twente campus. Based on campus characteristics and available data, a feasible autonomous delivery system is configured. This includes decision-making in the service area, demand estimations, depot configuration and operational logistics. These choices are shaped by both theoretical insights and practical constraints, and they serve as the inputs for the simulation experiments conducted in the next stage.

RQ. How can the autonomous last-mile delivery system be configured for the University of Twente campus using the developed framework and simulation model?

- a) Which design options from the framework best fit the characteristics of the UT campus?
- b) Which system configurations are promising for experimentation?
- c) How can these design choices be implemented in the simulation model?

Chapter 4 presents the system design for the UT case study and identifies potential experimental configurations that will be discussed in Chapter 5.

Stage 4. Experiments and Evaluation
This stage focuses on conducting simulation experiments using the previously selected configurations. The goal is to evaluate how the autonomous delivery system performs under different design scenarios.

RQ. How does the autonomous delivery system perform under different design configurations, and which setup offers the best operational performance?

- a) What experiments should be conducted to assess different delivery configurations?
- b) How do key performance indicators vary across setups?
- c) Which configuration provides the most balanced and effective performance for campus-wide implementation?

Chapter 5 presents the simulation results and provides insights into vehicle-specific performance, cost-efficiency, energy usage, and practical implementation recommendations for the University of Twente campus.

Stage 5. Drawing conclusions and providing recommendations

After experimenting with our autonomous delivery system, we can analyze the results and draw conclusions. This includes evaluating whether the designed autonomous delivery system meets performance expectations, interpreting the implications for the University of Twente campus, and outlining directions for future research or system deployment. The results are discussed in terms of both theoretical insights and practical relevance. This is all done in Chapter 6.

Chapter 2 Literature Review

Autonomous delivery for the last mile delivery is a promising, but also still distant concept. In this chapter, we aim to broaden our knowledge of about three aspects before continuing our research (i) general last-mile delivery (characteristics) (ii) autonomous delivery vehicles and (iii) the accompanying delivery systems. This will give us the theoretical background that we need to properly think about designing an autonomous delivery system. In this chapter, we aim to answer the research question:

What does existing literature reveal about the technologies, models and methods available for designing and evaluating autonomous last-mile delivery systems?

- **Section 2.1** briefly explains how the first, middle and last-mile delivery works in practice
- **Section 2.2** analyses the advantages and disadvantages of the current general last-mile delivery landscape.
- **Section 2.3** outlines the different types of autonomous technologies which could be used for autonomous delivery. It provides specific examples of technologies produced by certain companies as well as a discussion about the opportunities and limitations per type.
- **Section 2.4** describes the different types of delivery systems based on the technology found in Section 2.3
- **Section 2.5** focuses on modeling these systems, including the mathematical formulation of delivery problems
- **Section 2.6** reviews the solution approaches to solve such models
- **Section 2.7** discusses how to evaluate autonomous delivery systems
- **Section 2.8** concludes the literature review

2.1 Parcel Delivery Process in the Netherlands (PostNL)

From *Autoriteit Consument & Markt Dashboard* (2024), we obtained the following data. In 2024, a total of 606 million parcels were delivered in the Netherlands. Out of those, 494 million are delivered B2C (Business to consumer), while 95 million were B2B (business to business) and 16 million were C2X (consumer to business/consumer). The national parcel market share is distributed as follows:

- **PostNL:** 45-50%
- **DHL:** 40-45%
- **Other carriers** (e.g., DPD, UPS): 5 – 10%

Since PostNL has the largest market share, we will mainly focus on PostNL. Parcel Delivery by PostNL (to the campus) works as follows (Een Postgeschiedenis (n.d.); *Feiten En Cijfers Postbezorging* (n.d.)):

First Mile – From webshop to Sorting Center

After a customer places an order online, the retailer packages and labels the parcel. PostNL collects the parcel from the retailer or drop-off point and transports it to one of the 6 large regional sorting centers, for the campus this will be Zwolle.

Middle Mile – Transport between sorting centers and depots

Parcels are sorted by destination zip code and shipped across the national network. For deliveries to Enschede, parcels are forwarded to the Hengelo depot, which serves the Enschede area including the University of Twente Campus.

Last Mile – Final Delivery to the customer

At the Hengelo depot, parcels are loaded into delivery vans and distributed along predefined routes of which one includes the campus.

2.2 Advantages and Disadvantages of the current landscape

As previously said, last-mile delivery is one of the most polluting and costly stages of the delivery supply chain. Aside from these two critical concerns, the existing form of last-mile delivery presents several additional challenges. Ha et al. (2022) categorizes these concerns into five groups:

- Operational challenges
- Infrastructure challenges
- Delivery challenges
- Logistical challenges
- Environmental Challenges.

Below, we will explain these challenges individually to provide the reader with a basic understanding of the problems of the last mile delivery and therefore the need for innovative solutions. These challenges come from Ha et al. (2022).

Operational challenges

The first challenge is about operations, specifically the time required for loading and unloading parcels. Regardless of the mode of transportation, a certain amount of time is needed for parcel loading. Similarly, when the parcels reach their destination, some unloading time is needed. In high-density areas, this unloading time can contribute to other challenges such as traffic congestion, air pollution and noise pollution. Another operational challenge is the cost associated with maintenance, fuel, and labor for last mile delivery. Employing human delivery drivers incurs labor costs, and vehicles running on gasoline or diesel also incur fuel costs. Electric Vehicles (EVs) running on electricity is cheaper compared to cars with Internal Combine Engines (ICEs). It is worth noting that EVs typically have lower maintenance and repair costs compared to gasoline or diesel cars (Electric Vs. Gas Cars: Is It Cheaper to Drive an EV?, 2023).

Infrastructure challenges

In addition to operational challenges, there are also various infrastructure challenges. Example of infrastructure challenges are traffic congestion and there are not enough parking facilities in densely populated areas. With traffic congestion, it goes both ways. This means that it could be negatively impacted by traffic congestion, but it could be negatively impacting the traffic congestion (making it worse). In densely populated areas with a lot of traffic congestion, parcel delivery (especially when dealing with time windows) could run into major problems because it cannot deliver on time. The delivery vehicles could also cause traffic congestion because they may need to stop their vehicle at inconvenient spots due to the lack of parking spaces. Furthermore, the current IT systems being utilized for delivery face their own challenges.

Delivery of Parcels

The delivery of parcels itself presents additional challenges. The growing trend of fast and diverse delivery options, such as same-day delivery and specific time windows, means that parcels no longer follow a simple first-come, first-served (FCFS) approach. Delivery services and drivers must now navigate multiple customer preferences and priorities. Additionally, this comes with problems from order cancellations, incorrect or invalid customer addresses, and customer unavailability. Lastly, there is always the risk of parcel theft at the customer's doorstep.

Logistical challenges
Parcel delivery also means various logistical challenges. For instance, the most efficient delivery route may not always be used due to unexpected circumstances such as road construction. Moreover, the vehicles used for delivery often do not optimize their available space

Environmental challenges
Lastly, environmental challenges are a significant concern. Specifically, the use of diesel or gasoline vehicles for delivery contributes to greenhouse gas emissions, as well as air and noise pollution in urban areas. One contributing factor to this issue is the relatively inefficient fuel-to-mile ratio

Besides these disadvantages, there are also a couple of noteworthy advantages, with a particular focus on the human deliverer.

- Flexibility: Human deliverers could adapt to changes in delivery routes and schedules based on real-time conditions, such as traffic congestion or customer preferences.’
- Customer interaction: Human deliverers provide a personalized touch by interacting directly with customers, addressing any concerns or questions they may have about their deliveries. These issues encountered during the delivery process, can then be quickly resolved. Human deliverers also build trust and rapport with customers through consistent and reliable delivery service, leading to customer loyalty and repeat business
- Versatility and adaptability: Human deliverers can handle unexpected situations, such as locating difficult-to-find addresses or navigating complex delivery environments like apartment buildings or gated communities. They could also predict, based on personal experience, whether certain routes may or may not be more efficient (contradicting the available data). They can also handle a wide range of parcel sizes and types, including fragile or perishable items, ensuring safe and secure delivery. In addition, deliverers can handle special delivery requests, such as specific delivery time slots or instructions for leaving parcels in secure locations.

2.3 Autonomous Delivery Solutions

The second subsection focuses on the concept of Autonomous Delivery Vehicles. The concept of Autonomous Delivery Vehicles (ADVs) is first explained with the different levels of autonomy. Next, it outlines the main types of autonomous delivery vehicles with their opportunities and limitations. Finally, the subsection ends with a general reflection on the limitations of autonomous vehicles.

2.3.1 What is an automated delivery vehicle (ADVs)

Autonomous delivery refers to self-driving vehicles, robots, or drones that are used to transport goods and packages without the help of human intervention or control (Vivatechnology, 2025). They use technologies such as sensors, cameras, and artificial intelligence to scan their surroundings and make decisions based on their observations (Vivatechnology, 2025). The level of autonomy can vary, with some vehicles requiring human intervention, when necessary, while others are fully autonomous and do not require human input. The levels of autonomy for ADVs can be classified using the Society of Automotive Engineers (SAE) autonomous driving levels. These levels range from 0 to 5 (See Figure 4), with 0 indicating no automation and 5 indicating full automation (Inoiță, 2017).

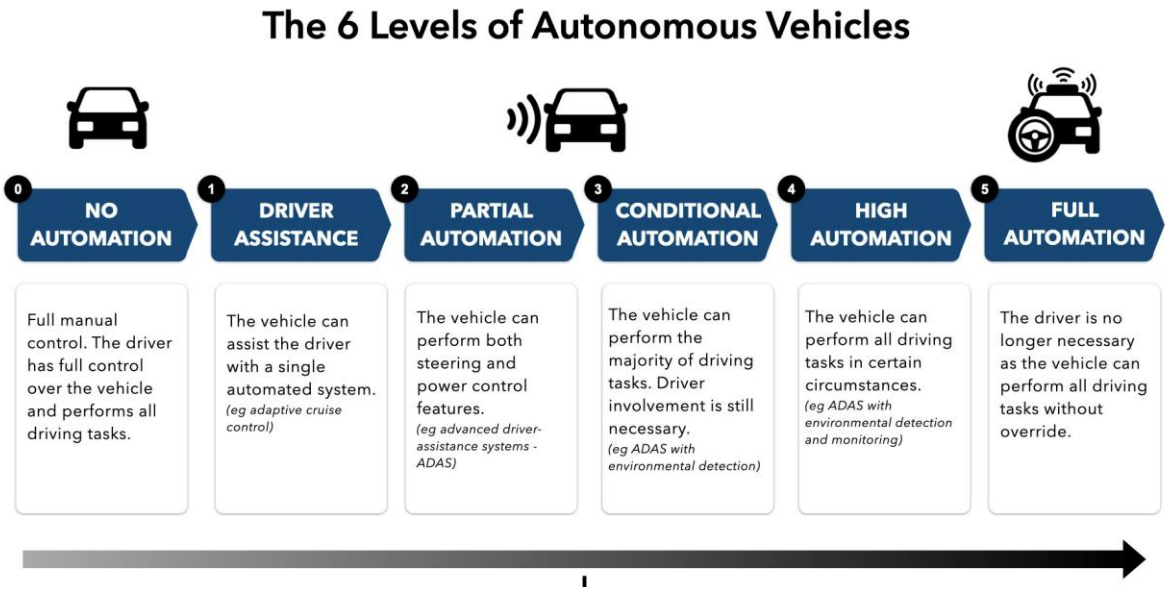

Figure 4: 6 Levels of Autonomy (Bruneteau & Bruneteau, 2022)

2.3.2 Classification of autonomous delivery Vehicles

Now that we've discussed the concept of autonomous delivery, let's broaden our view to explore the different categories that exist within this field. This section examines the varied designs and functionalities that characterize ADVs. The literature is not always consistent in the terminology for describing the technologies for autonomous delivery (A lot of times different terms are used to describe the same thing). In this thesis, we use the term ADVs for the entire set of autonomous vehicles/robots/drones that are used to transport goods and packages. Most of the papers online use this as well as an umbrella term. A distinction frequently made in the literature is between ADVs based on their mode of operation (ground or air) and their vehicle size or carrying capacity. The three main types of ADVs are:

- Sidewalk ADVs: small delivery vehicles designed to drive on the sidewalk/pavements

- Road ADVs: larger delivery vehicles designed to drive on public roads

- UAVs: self-operating flying vehicles that transport through air

Sidewalk ADVs (sometimes referred to as **S-ADRs** by Engesser et al (2023), or simply **ADR** by Alverhed et al. (2024)): These are compact, lightweight delivery robots designed to operate on sidewalks and pedestrian walkways. They travel at a pedestrian speed of around 6 km/h, are relatively small and have a limited carrying capacity (one or two small parcels). S-ADVs use sensors and cameras to navigate safely among pedestrians and deliver parcels to residential or commercial buildings. Starship (Figure 5) is the market leader (Autonomous Delivery Vehicles Market Insights, n.d.) in this form of autonomous delivery, furthermore, many competitors in this field seem to have been discontinued such as amazon (Soper & Day, 2022).

Figure 5: Sidewalk ADVs (Starship robots), taken from

The main advantage according to Starship Technologies (2024b) of S-ADVs is their ability to drive at, like the name says, sidewalks and other places (pedestrian zones) where normal cars or vehicles are not allowed. This allows them to reach additional places (customers), which could otherwise not have been visited. The vehicles can also go up curbs, allowing them to travel across the street. They come with a specially designed, insulated lining. This keeps the food at the desired temperature for the duration of this journey and are also designed to withstand almost all-weather conditions. However, they may face obstacles such as parked cars, construction sites, and even uneven surfaces, which can hinder their smooth navigation. Moreover, S-ADVS may have limited battery life, reducing their operating range and delivery capacity.

Additionally, Gherke et al. (2023) demonstrate in their study that areas with crossings and limited roads are the centers of attention of moderate and dangerous conflicts, with no distinction made as to what space travelers should occupy. It will be difficult for practitioners to safely introduce S-ADVs onto existing crowded sidewalks and urban paths, possibly demanding creative solutions related to S-ADV route planning and urban infrastructure reform. According to observational research from Weinberg et al. (2023), many individuals stop because they are interested in the vehicles, obstructing the path of the vehicle and forcing it to stop. Naturally, this will be less likely in the future as these vehicles become less unique. However, it is highly likely that a S-ADV will have to stop frequently on a busy sidewalk because it could collide with people/animals.

S-ADVs also must face the issue of requiring an actual pavement. They are unable to or are not permitted to drive on public roads (because of slow speed for example). As a result, these vehicles must be perfectly up to date on the current infrastructure and understand that if there is no pavement or road closures, they cannot use these routes.

Road ADVs (also referred to as **R-ADRs** by Engesser et al (2023) and sometimes called **mobile parcel lockers** or **Autonomous Guided Vehicles (AGVs)** by Chen et al. (2021a)) are larger delivery vehicles designed to operate on public roads. They are equipped with sensors and navigation systems to navigate through traffic and deliver parcels to their destinations. Even including external airbags for pedestrian protection in case of a collusion. However, because these vehicles share infrastructure with conventional traffic, they are exposed to a much broader and more complex set of challenges (such as unpredictable human behavior). R-ADVs are suitable for larger parcel sizes and longer distances compared to sidewalk ADVs. These R-ADVs are developed in many different shapes and sizes. A couple of the main examples given in research papers are Nuro, Udelv and Neolix (Buldeo et al., 2022b; Jennings, 2020; Srinivas et al., 2022).

In January 2022, Nuro announced their newest level 4 autonomy model R3 (Figure 6). This new generation will even fit more cargo (24 bags of groceries) and has modular compartments to keep meals hot and drinks cold. At the end of 2024, Nuro announced that they will expand by testing their R3 vehicles after they changed their business strategy by licensing its technology to automakers (Bellan, 2024). Neolix's Autonomous Vehicle X3 Plus has received the world's first Level 4 autonomous driving system international safety standard certification. Their vehicles have been successfully implemented in application situations in 12 countries ranging from Norway and Saudi Arabia to Australia and Japan (*Neolix*, n.d.). One of the key advantages of the Neolix autonomous vehicles is their modular cargo container, which allows for scenario customization, such as food delivery. They also have the capability of replacing its battery in 30 seconds, allowing the vehicles to run continuously for 24 hours. A Dutch company called Macrostep, has also developed their own autonomous delivery vehicles called Express robot. These robots have already been tested in holiday park Duinrell (see Figure 6).

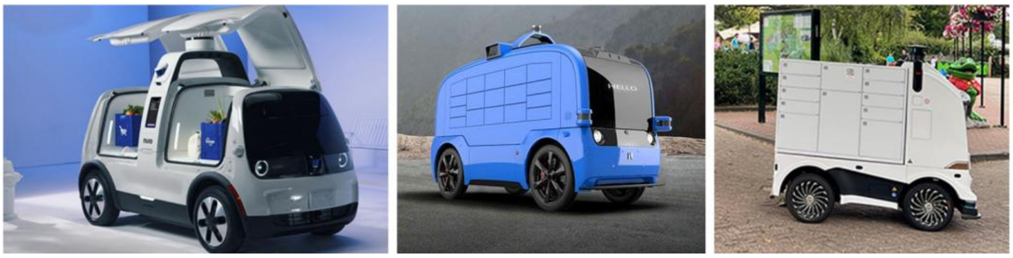

Figure 6: Nuro R3 (Left), Neolix (middle) and Express Robot in Duinrell (left); retrieved from: (Wessling, 2022), (Neolix, n.d.) and (MacroStep BV, 2025)

A larger example of a Road ADV is the Udelv transporter (Figure 7). The Udelv Transporter (level 4 autonomy) is a bus shaped autonomous vehicle built to deliver as many as 80 parcels. The storage space on the bus can be adjusted to accommodate parcels of all shapes and sizes. With its maximum speed of 112 km/h, potentially, the transporter could also drive outside of urban areas on the roads. This vehicle is still in development and the company aims to have 50,000 transporters on public roads by 2028 *(Transporter | Udelv, n.d.)*. Some papers, like Srinivas et al. (2022), use the term ADV to describe these larger road-based vehicles. However, since they still also fall under the category of being road-based delivery vehicles, it makes sense to place them under the category road ADVs.

*Figure 7: Udelv transporter, retrieved from: (*Transporter | Udelv*, n.d.)*

The other category of ADVs is the **Unmanned Aerial Vehicles (UAVs)**, more commonly known as drones. UAVs are self-operating flying vehicles that have the capability to transport parcels through the air. These vehicles don’t require a large landing surface because they can take off and land vertically. UAVs uses a so called “sense and avoid” technology, which enables it to detect and evade obstacles. It starts off in vertical takeoff and then switches to horizontal flight. Short-distance, light-weight package transportation is a good application for UAVs.

Amazon has its own UAV delivery service called Prime Air, which has been deployed since 2023. This program is utilized for delivery in three locations within the United States, as well as cities in Italy and the United Kingdom. Zipline, another UAV company, also delivers packages to locations in America (*Zipline Fact Sheet*, n.d.). Amazons Prime Air UAV can carry parcels weighing up to five pounds and is designed to guarantee quick delivery to customers, usually in an hour or less (Iddenden, 2022). Customers must have a designated “marker” in their garden or fenced area (in case of an apartment building). The UAV uses these markers to locate the drop-off location. In case of Amazons Prime Air UAV, the parcels are released from a height of about 3.7 meters (Chen, 2023b), Zipline uses a mini ‘droid’ which is lowered to the drop-off point (*Zipline Fact Sheet*, n.d.).

Figure 8: Left: Amazon Prime's delivery UAV (Chen, 2023b), Right: Ziplines UAV (Zipline Fact Sheet | Zipline UAV Delivery & Logistics, n.d.)

The biggest advantage of UAVs is that they do not have to deal with traffic, only with other UAVs in the sky (and objects/buildings). By navigating around traffic and skipping traditional road infrastructure, delivery times are significantly reduced. They can use almost straight delivery routes, which allows for very fast deliveries. Another advantage is that UAVs are relatively cheap compared to the other ADVs.

Despite their advantages, UAVS also have significant drawbacks according to Bahabry et al. (2019). One major limitation is their battery life, which restricts their range and continuous operation. This can have negative consequences if the battery runs out, such as application failure or even a crash. Another challenge in urban areas is the presence of tall buildings/trees, which can obstruct the direct routes of UAVs, especially when flying at low altitudes. In other words, the mobility of UAVS is constrained by different permanent and temporary obstacles, which must be considered for optimal path selection and improved navigation. Additionally, tall buildings in urban areas reduces the available airspace for UAVs, thereby limiting the number of UAVs that may be used. It is important to consider the prevention of UAV collisions in urban areas when implementing UAVS. This becomes even more crucial if the number of UAVs in operation increases significantly

UAVs are available in various sizes and configurations, each with its own restrictions on weight and range. In UAV technology, there is a tradeoff between the capacity to carry cargo and the duration of flight. Generally, commercially available UAVs have limited payload capabilities (Singhal et al., 2018). For instance, Amazon's Prime Air UAV can transport parcels weighing up to 2.26 kg (Chen, 2023), which is like the typical payload capacity of a small UAV, around 2kg (Young, 2023), (UAVs, 2021).

The MK30 is quieter and will be able to fly in more diverse weather conditions, meaning customers can receive super speedy deliveries even in situations like light rain and hotter or colder temperatures. However, many countries, such as the Netherlands, experience more extreme weather conditions, particularly heavy rain or snowfall in the winter and fall. Considering climate change and its effects on extreme weather conditions, it is suggested that these events will occur more frequently, always raising the question about the availability of UAV delivery.

The use of UAVs raises significant privacy issues. Many recreational UAVs come with advanced camera systems that enable owners to view and record footage. However, this capability also raises concerns about potential misuse, as individuals may use UAVs to observe places where they should not have access. Previous incidents, like the one in Australia, have highlighted the negative consequences that can arise from such misuse (Yahoo Is Part of the Yahoo Family of Brands, n.d.).

2.3.3 Legal Restrictions

One of the biggest issues with the implementation of autonomous (delivery) vehicles are the legal restrictions. Of course, it should be very important that the implementation of this technology must be at least as safe as or even safer than the current (delivery) vehicles. In this section, we discuss the legal restrictions of autonomous vehicles in the Netherlands. Further on in this research, we will neglect the legal restrictions on autonomous (delivery) vehicles. If we must assume all restrictions now, it is basically not possible to come up with a good autonomous delivery solution. It must be assumed that the use of autonomous delivery will eventually be possible.

Because the implementation of autonomous cars is still in an early development phase, it is not always allowed to use or test this technology. The Dutch approved a bill back in 2019 that allowed self-driving vehicle testing to begin without the presence of a driver. The ‘Experimenteer wet zelfrijdende auto’ (the law governing the experimental use of self-driving automobiles) eliminates legal obstacles. As a result, manufacturers will have more opportunities to undertake self-driving vehicle tests. This law will allow companies to seek permission to perform experiments on public roads using autonomous cars, with a human on hand to take command via remote control (Ministerie van Infrastructuur en Waterstaat, 2023).

In November 2021, clevon mobility, tested their autonomous delivery vehicle in the city of Eindhoven during a 5-day pilot project carried out by DPD Netherlands and Clevon. This marks the start of bringing autonomous delivery services to the public streets of the Netherlands (Railway-News, 2021). Besides this pilot project, there have been other small tests around the Netherlands, mainly on college campuses (Ecommerce News Europe, 2020). However, larger scale testing has not been done yet.

Besides road autonomous vehicles, there are also restrictions on UAVS. The Dutch government divides UAVs into three categories: open, specific, and certified, and they are all subject to flight rules. Low-risk UAV flights are under the open category, which has restrictions including a weight limit of 25 kg, a maximum altitude of 120 meters, and constant visual contact with the UAV. UAVs with greater flying risks, such as those that exceed the previously stated conditions, fall under this category. The last category, which is not yet completely defined, is for “high” risk UAV flights. UAVs longer than three meters that are used to transport person or dangerous goods, or that fly above crowds, fall under this category. (Netherlands Enterprise Agency, RVO, 2023).

2.4 Autonomous Delivery Systems

Using the technologies discussed in 2.3 Autonomous Delivery Solutions (S-ADVs, R-ADVs and UAVs), this section briefly discusses the different autonomous delivery systems known in the literature. Engesser et al. (2023) and Srinivas et al. (2022) discuss the different autonomous delivery options available right now. In section 2.4, we will discuss these systems again in detail and explain how they can be modeled. In general, there are two options: centralized systems or decentralized systems.

In **centralized systems**, the ADVs deliver the packages directly from a single depot/warehouse to the customer(s). The advantage of a centralized system is that it simplifies planning and coordination. Fleets can be homogeneous (for example only UAVs) or heterogeneous (mix of S-ADV, R-ADV, and/or UAV). The logistical processes of homogeneous fleets are easier to manage, compared to a heterogeneous fleet. However, heterogeneous fleets offer greater delivery flexibility (for example, UAVs for remote zones) but this comes with more complex (operational) planning.

In **decentralized systems**, multiple mobile or fixed nodes are used across the service area. ADVs may be relocated throughout the day or mobile depots are used to transport the ADVs to another location. Although this method improves responsiveness and flexibility, it requires more complex planning and coordination. The decentralized systems can be further categorized into three categories:

- Two-Tier Model: conventional trucks transport several parcels to smaller local hubs. From these hubs, the parcels are delivered to the customers.
- UAV-aided Model: A delivery truck equipped with UAV(s) leaves a distribution center to deliver parcels. The delivery truck serves as a mobile depot for the UAV(s).
- S-ADV aided Model: A delivery truck equipped with S-ADV(s) leaves a distribution center to deliver parcels. The delivery truck serves as a mobile depot for the S-ADV(s).
- R-ADV aided Model: In case the R-ADVs are not allowed to drive to a part of service area This is where the platoon model comes in, in which the R-ADVs follow one manually operated vehicle to an AV-friendly zone where they can deliver themselves

Besides these models, the literature also discusses other variants of the delivery model focusing on a specific aspect like reserving lanes or dropping of delivery personnel. Reed et al. (2022) suggest a model in which delivery workers are dropped off in autonomous vans near customer locations. At the drop-off locations, the delivery personnel can deliver the parcels on foot and are picked up later. This solves the problem of having to search for parking spaces and customers having to actively pick up the parcels from the vehicle, but it does come with extra personnel costs.

There are also a lot of papers focusing on one specific option in a predetermined area. Bakach et al. (2022), for example, investigates a robot-based last-mile delivery problem considering path flexibility given the presence of zones with varying pedestrian Level of Service (LOS). While Gherke et al. (2023) observes the sidewalk autonomous delivery robot interactions with pedestrians and bicyclists.

Besides autonomous delivery options, there are also other options proposed in the literature. Akeb et al. (2018) provides an idea to encourage citizens in the same neighborhood (known as “Neighbor Relays”) to collect and deliver parcels to the end consumer when the consumer is not at home. Consumers will be notified and can contact their neighbor using mobile apps to plan parcel delivery. This “Neighbor Realy” earns money in exchange. The high population density in urban areas makes this option appealing to both neighboring relays (in terms of money) and transporters (in terms of cost).

2.5 Modelling and Problem Formulations of Autonomous systems

This subsection builds upon Section 2.4 Autonomous Delivery Systems by focusing on the modelling and problem formulations. The goal of this subsection is not to present a systematic literature review about the different systems, but to try and classify the different problems. Engesser et al. (2023) and Srinivas et al. (2022) discuss the different autonomous delivery options available right now. Figliozzi (2020) examines the efficiency of self-driving (driverless) air and ground delivery vehicles in terms of vehicle miles, energy consumption, and carbon emissions. Jennings (2020) studies the potential impacts on freight efficiency and travel of road autonomous delivery Vehicles.

Figure 9 shows an overview of the different autonomous delivery systems as discussed earlier in 2.4 Autonomous Delivery Systems. This overview is adjusted from Srinivas et al. (2022), with additions that include UAVs and a distinction between S-ADRs and R-ADRS (originally grouped together as ADRs). The literature contains many more different variants than are present below, this should only ensure that we have a general overview of the possible systems. In doing so, we have chosen those systems that are autonomous as much as possible, unless otherwise not possible. The meaning of "as autonomous as possible" indicates that options requiring human intervention (like autonomous vans that send delivery workers) are being ignored. Also, for convenience, we have left out the specific models (focussing on load dependent flight or lane reservation), as they do not so much use a different system but create a different solution.

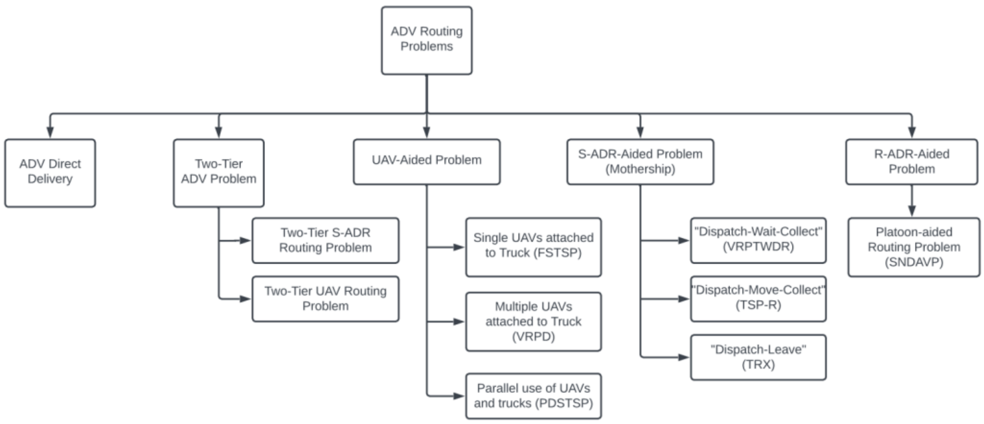

Figure 9: Classification of routing literature. Adjusted from Srinivas et al. (2022)

2.5.1 Foundational routing problems

Two classical routing problems form the basis of the modern delivery models: the Traveling Salesman Problem (TSP) and the Vehicle Routing Problem (VRP).

The TSP, first introduced around the 1930s (Pop et al., 2023), is one of the oldest and most well-known problems in combinatorial optimization. It can be stated as:
"*If a traveling Salesman wishes to visit exactly once each of a list of m cities (where the cost of traveling from city i to city j is c_{ij}) and then return to the home city, what is the least costly route the traveling salesman can take?*" (Hoffman & Padberg, 2001).
See Figure 10 for an example.

The VRP was introduced later, in 1959 (Laporte et al., 2013) and generalizes to TSP where the objective is to efficiently plan multiple routes for a fleet of vehicles from a single starting point (depot/warehouse/distribution center) to visit all specified locations. Each route starts at the starting point, visits a subset of the nodes, and returns to the starting point. Figure 10 shows an example.

The TSP and the VRP come in different variations. Take for example the Vehicle Routing Problem with Time Windows (VRPTW) or the Capacitated Vehicle Routing Problem (CVRP). VRPTW considers delivery time windows, and the CVRP considers the capacity of the delivery vehicles (Baldacci et al., 2012). It is even possible to combine both options into one model, the Capacitated Vehicle Routing Problem with Time Windows (CVRPTW).

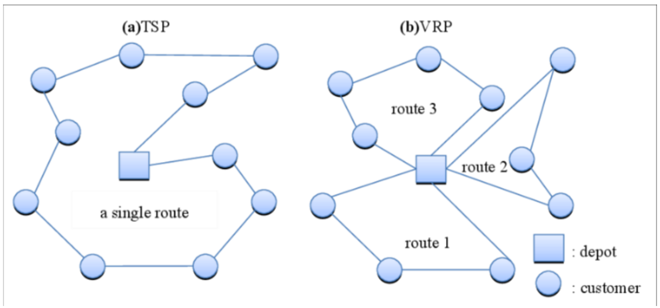

Figure 10: Traveling Salesman Problem (a) and Vehicle Routing Problem (b)--an illustration. Retrieved from (Liu et al., 2014)

Over the years, many different variations of the VRP and the TSP have been introduced. We will cover the models per ADV type, starting with UAVs, then covering S-RADVs and finally the R-ADVs.

2.5.2 S-RADVs Models

A company can decide to only use S-ADVs for their delivery operations. In this case, the problem is simplified to a VRP scenario, where a S-ADV delivers the products to a variable number of customers, depending on the capacity of the S-ADV. This will be referred to as the ADV Direct Delivery Problem.

Companies might also look further and use a hybrid model, the truck-and-robot or mothership van model. This model complements S-ADVs with their limited range, that can be used to drop off and pick up several S-ADVs. However, the mothership itself is not an autonomous vehicle (yet) and requires a driver. The mothership can load and transport up to eight S-ADVs (Figliozzi, 2020).

Two-Tier Model

In the Two-Tier Model, conventional trucks transport several parcels to smaller local hubs (for example S-ADV hubs). From these hubs, the parcels are then distributed to the customers. Bakach et al. (2021) proposes a two-tier model with hubs able to store S-ADVs and parcels, with and without time windows. Since it is a very general model, it could also be applied to UAVs.

Mothership Dispatch-Wait-Collect

In this model, a van with SADVs drives to a drop-off location. Here, the van loads the parcels into the SADVs, and they deliver the parcels to the customers. The truck waits at the same location and continues with the route once every SADV has returned to the truck. Chen et al. (2021) proposed in their paper a new Vehicle Routing Problem with Time Windows and Delivery Vehicles (VRPTWDR). The trucks dispatch Vehicles nearby customers while the driver of the truck is also serving customers.

Mothership Dispatch-Move

This model works almost the same as the Dispatch-Wait-Collect model, except for the fact that the truck does not specifically wait until all the SADVs are back before it continues. This model considers the possibility that some SADVs will drive ahead to the next location, where they will be picked up by the van. Simoni et al. (2020) formulated this problem as the Traveling Salesman Problem with Robot (TSP-R).

Mothership Dispatch-leave

Another version of the mothership van problem is one where the S-ADVs do not have to return to the van. Instead, they return to so called ADV stations where they are stored and charged, waiting to be picked up by another van for a new route. These ADV stations do not store any parcels, only S-ADVs. Ostermeier et al. (2021) present in their paper the Truck-and-Robot Cost-optimal Routing approach (TRX). Instead of the variant presented by Boysen et al. (2018), the paper of Ostermeier focuses on a limited availability of S-ADVs.

2.5.3 UAV Models

Besides S-ADR deliveries, a company can also decide to deliver with UAVs. Given the current capacity limitations with UAV delivery (only able to serve one customer), this resembles a series of individual dispatch problems. However, once the capacity and possibilities of UAV delivery increases, it might become feasible to expand operations to delivery to multiple customers by route. The problem will then become a VRP. This will be referred to as ADV Only Delivery. This ADV Direct Delivery problem is basically a VRP, using the characteristics of the specific vehicle used in delivery.

In addition to the option of UAV-only delivery, a hybrid model is often suggested in the literature. With this approach, a van is used to transfer the UAV(s) and/or deliver the larger parcels. At the same time, the UAVs can charge on the van. This mitigates the constraints of flight duration and radius associated with UAVs. The Flying Sidekick Traveling Salesman Problem (FSTSP) and the Parallel UAV Scheduling Traveling Salesman Person (PDSTSP) are two mathematical models proposed by Murray and Chu (2015) to model the truck and UAV combination (As implied by its name, both are derived from the conventional Traveling Salesman Problem).

In the FSTSP, a delivery truck equipped with a UAV leaves a distribution center to deliver parcels. The UAV is deployed to deliver lightweight parcels to consumers along the truck's route. After delivery, the UAV returns to the truck or distribution center. The objective is to minimize the total delivery time or distance traveled by the UAV pair.

The PDSTSP extends the FSTSP by incorporating multiple UAVs operating in parallel with a delivery truck. Each UAV can independently deliver parcels within its operational range. The PDSTSP is seen as the most viable option if a lot of the customers are located around the depot and inside of the operational range. The FSTSP is seen as the viable option when the depot is located further away from the customers. However, it is not always straightforward, as there can be instances that demonstrate its complexity and show that the opposite is true. Othman et al. (2017) presents a model where a truck, accompanied by a UAV, follows a fixed predetermined route. A UAV is detached at certain locations and returns to another location, where it can charge at the truck. Kitjacharoenchai & Lee (2019) discuss the Vehicle Routing Problem with UAVs (VRPD) with multiple UAVs at a single truck.

Besides these papers, there have been many others researching different aspects of problems regarding UAV delivery such as load-dependent flight speed (Nishira et al., 2023), a delivery model with backhaul option (Jeon et al., 2021), the option of using UAVs for resupply at transshipment points (Moshref-Javadi et al., 2023). Moadab et al. (2022) even discuss the potential of a UAV routing problem using public transportation as moving charging stations.

2.5.4 R-RADVs models
The final models covered in this subsection are the models concerned with the R-ADVs. Because of the advantage that R-ADVs have enough space and thus are suitable for delivering basically all types of parcels, they do not need to be supported by another vehicle. This makes the routing problem of R-ADVs less complex. The R-ADVs also have a simple ADV Direct Delivery problem, which is again basically a VRP.

In contrast to S-ADVs, R-ADVs are larger and heavier and are not easily transferred by vans. Then again, they can achieve higher speeds, carry more parcels and cover more distance. This makes them much more suitable for autonomous transport to multiple customers. However, in the current infrastructure, there will be some zones that are not suitable for autonomous vehicles, so they may need to be moved to an AV-friendly zone. This is where the platoon model comes in, in which the R-ADVs follow one manually operated vehicle to an AV-friendly zone where they can deliver themselves. Scherr et al. (2019) formulated this model as the Service Network Design for Autonomous vehicles in Platoons (SNDAVP)

R-ADVs take up a lot of space, so it's not clear how they'll coexist with other modes of transportation on the road. To address the issue, Wu et al. (2017) explored an ADV Transportation Problem with Lane Reservation (ATP-LR). This approach focusses on optimally reserving lanes in a transportation network to establish transportation routes for each task. The problem is tackled through a two-phase exact algorithm, where the feasible routes are first determined, after which the optimal lane scheme and delivery task path is determined.

2.6 Approaches to solve vehicle routing problems
This subsection of the literature review delves into the approaches used to solve vehicle routing problems (VRPs). As briefly mentioned in earlier sections, VRPs vary widely in complexity, and are impacted by several variables, including fleet size, customer demand patterns, time windows, and geographic limitations. As a result, a wide range of approaches have been developed to deal with these issues, from heuristics and metaheuristic techniques that produce high-quality solutions in a reasonable amount of computational time to exact algorithms that ensure optimal results.

The size and restrictions of the problem frequently influence the strategy used because many VRP versions are computationally challenging. This section first examines exact approaches, which have limitations in terms of scalability but are appropriate for small to moderately large issues. The subsection then shifts to heuristic and metaheuristic techniques, which are appropriate for large-scale, real-world VRPs and compromise optimality for efficiency.

Exact Methods
Exact methods are techniques that can be used to solve a problem to optimality. However, only small or moderately sized problems can be solved to optimality in practice, because the run-time increases considerably with the size of the problem. Integer Linear Programming (ILP), a mathematical program with constraints, is one of the most widely used accurate techniques for modeling VRP problems. It is a type of optimization model where the variables are integer values and the objective function equations are linear. Branch and Bound or Branch and Cut are two methods which are used to solve these models. By breaking the ILP problem down into smaller subproblems (or branches) and then removing some of the branches based on bounds on the optimal solution, the branch and bound methods solves the problem. The branch and cut method uses the branch and bound method to turn an optimal non-integer solution into an optimal integer solution by imposing additional constraints (Genova & Guliashki, 2011).

Heuristics

Since the basic version of each VRP (basic version is TSP) is NP-Hard, each vehicle routing problem which must be solved is essentially NP-hard, which means that solving these (large) problems can take an impractically long time. In these cases, heuristics and metaheuristics are often used. Heuristics are basically 'good enough' solutions, which means that they can be optimal but also non-optimal. This is also dependent on the type of heuristic(s) and the problem size (and its characteristics) itself.

Heuristics can be further divided into three different types: constructive heuristics, improvement heuristics, and metaheuristics (F. Liu a et al., 2023; see Figure 11).

Constructive heuristics are algorithms that use a fixed empirical heuristic procedure to construct routing solutions from zero as explained by F. Liu a et al. (2023). They typically generate a feasible solution fast and are easy to implement in many different scenarios. However, this comes at a cost. The solutions generated often have a gap to the optimal solution. Take for example the Nearest Neighbor Method, which allows for the construction of routes in either parallel or sequential manner. From the starting point, which is typically a depot or hub, a route is constructed by continuously adding additional unallocated customers until no more can be added.

Improvement heuristics, according to F. Liu a et al. (2023), iteratively improve a routing solution by performing a local search in the neighborhood. In optimization, a neighborhood refers to the set of solutions that can be reached from a current solution by making a small change. Intra-route methods are methods focused on improving a single route. For example, exchanging the order of two customers in one route. This can change (improve or decrease!) the total amount of kilometers of a route. Inter-route methods are methods focused on local searches across multiple routes. Many of these techniques are simply intra-route techniques but for multiple routes. So, swapping customers from two different routes. There are also some downsides to these heuristics. It can happen that with only these local searches, the solution can become stuck on a local optimum. This means that the solution can no longer be improved anymore with these local searches (swaps, inserts).

Finally, F. Liu a et al. (2023) explains that metaheuristics try to tackle this problem of being stuck in the local neighborhood. Metaheuristics are high-level frameworks designed to explore the broader solution space. They use methods like randomization, memory or adaptive learning to direct local searches (and sometimes constructive heuristics). This makes them suitable for large problems which tend to fall into local optima. Ant Colony, Tabu Search and Simulated annealing are examples of popular metaheuristics. For example, simulated annealing allows sometimes a change which makes the solution worse, to try and find other better solutions based on the worse solution. By storing information in a short-term memory list (tabu list), Tabu Search prevents looping back to previously explored solution. These are called single-solution-based methods, which focus on the one solution you have (for example a problem with 3 routes). Population-based methods focus on multiple solutions (creating a sort of copy of your problem) and improving all these solutions at the same time. Take for example the ant colony algorithm. Instead of guiding one ant to the end goal (Single-solution-based methods), multiple ants are sent to try many options and by learning from each other (Ant Colony Algorithm), so they can find smarter solutions over time (F. Liu a et al., 2023).

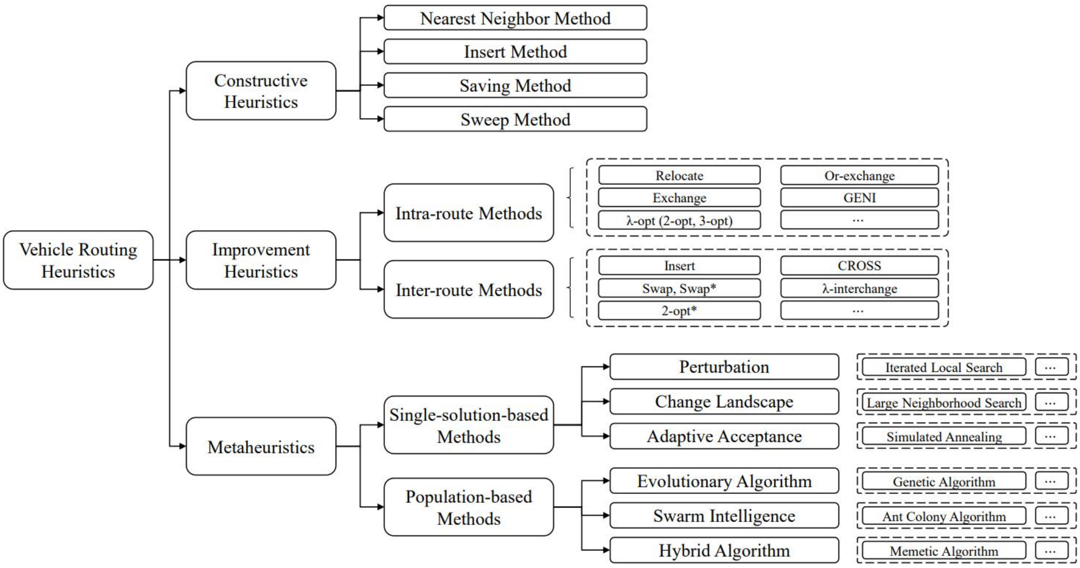

Figure 11: Overview of heuristics, taken from (F. Liu a et al., 2023)

2.7 Evaluating the solution

From network design to vehicle routing, there are several steps involved in creating an effective and sustainable last-mile delivery system. However, the process does not end with proposing a solution. It is essential to evaluate the solution to understand its performance under realistic (as much as possible) conditions.

Key Performance Indicators (KPIs) are the essential indicators used for evaluating the effectiveness of a logistics system. They enable organizations to assess and optimize critical aspects such as cost-efficiency or environmental impact. According to Seuring and Müller (2008), three main components define sustainable development in logistics:

1. **Economic** – Focus on the financial performance and efficiency
2. **Environmental** – Measure the environmental impact of the logistics operations
3. **Social** – Assess social factors such as public perception and workforce impact.

The specific goals and context of the delivery system determine which KPIs are used, and to give a complete picture of system performance, several KPIs are frequently assessed at the same time. A summary of popular KPIs for last-mile delivery systems, divided into the three above-mentioned categories, is shown below (Table 1). These KPIs are compiled from Zis et al. (2023), Aliev et al. (2019) and Morana and Gonzalez-Feliu (2015).

KPIs may specify what should be measured, but it is just as important to determine how to measure it in practical and unpredictable situations. This problem can be addressed by Monte Carlo Simulation (MCS), which offers a probabilistic evaluation framework. The delivery system simulation is repeatedly run by MCS, with key input parameters being randomly varied based on predetermined probability distributions. Traffic-impacted travel times or customer pickup times are examples of these inputs. (Raychaudhuri, 2008)

Table 1: Overview of possible KPIs

Economic	Environmental	Societal
General:	CO2 Emissions (kg/km)	**Safety & Security**
Distance	Nox, Sox Emissions (kg/km)	Accident Rate (%)
Moving time	Total Energy Consumption (kWh/unit)	Emergency Stop Rate (%)
Loading/Unloading Time (Mins)	Charging Efficiency (%)	Failure to Avoid Obstacles (%)
Operational Efficiency	Noise Pollution (dB)	Crime Incidents (#)
Delivery Success Rate (%)		**Workforce impact:**
Load Efficiency (%)		Employment Turnover (%)
Average Delivery Time (Mins)		Employment Change Rate (%)
Delivery throughput (parcels/hour/day)		Training hours per employee
Idle Time / Utilization Rate (%)		**Public Perception**
Reliability & Performance:		Public Perception (Score)
On-Time Delivery Rate (%)		Work-life balance (Score)
Deviation Time (mins)		Public complaint rate (#)
Error Rate		Total Minutes Late
Mean Time Between Failures		Late Customers
Cost Efficiency		Average Time After TimeWindow
Cost per Delivery (€/parcel)		
Energy Consumption per km (Energy/km)		
Maintenance Costs (€)		
Daily operating costs (€)		
Customer Satisfaction		
Customer Satisfaction Score (%)		

2.8 Conclusion

In this chapter, an extensive literature review was conducted to build the theoretical foundation for the remainder of the thesis. The review addressed six core themes: the current state of last-mile delivery, types and characteristics of autonomous delivery vehicles (and their opportunities and limitations), different delivery systems and mathematical models, approaches to solve vehicle routing problems, simulation methods for evaluation and KPIS for performance assessment.

The current last-mile delivery landscape is challenged by five main challenges: Operational Challenges, Infrastructure challenges, delivery challenges, logistical challenges and environmental challenges. ADVs are a promising innovative solution to these challenges. We have covered three main types of ADVs being: Road Autonomous Vehicles (R-ADVs), Sidewalk Autonomous Vehicles (S-ADVs) and Unmanned Aerial Vehicles (UAVs), each with unique operational characteristics and constraints. We have concluded that the current legislation is not suited for ADVS, it does not allow autonomous delivery of parcels. Therefore, some legislative changes would be necessary for the future use of ADVs in the last-mile delivery landscape.

The literature also outlined a wide range of (partly) autonomous delivery systems, including single-tier and two-tier networks, and ADV-aided systems where an autonomous delivery vehicle is aided in delivery by a (often) manually driven van. These systems are all based on variations of the VRP, such as time window and capacity constraints.

The systems have a mathematically based model, which can be solved by multiple approaches. Exact methods such as Integer Linear Programming provide can provide a fast solution for smaller VRPs. Once the size of the problems increases, non-exact methods become necessary to find a (near) optimal solution. Constructive heuristics, like the nearest neighbor approach, create an initial solution, which can then be optimized by local search heuristics like 2-ot, and 3-opt. Metaheuristics are a problem independent technique which can be applied to a broad range of problems.

Finally, the chapter provided a section on the evaluation method of the designed system. A structured overview of relevant KPIs used to evaluate delivery performance across economic, environmental, and societal dimensions is given. The literature provides a wide range of KPIs to evaluate last-mile delivery systems. However, their applicability largely depends on the specific context and objectives of the evaluation. This makes it essential to select KPIs that are not only theoretically relevant but can also be practically measured.

This literature review contributes to the literature in several ways:

- It connects the different domains (technology, logistics, modeling, evaluation) that are all relevant for autonomous last-mile delivery
- It shows that there are many different systems and combinations, which also means there are various decision variables to consider when designing an autonomous delivery system.
- It provides an overview of commonly used KPIs for autonomous delivery evaluation, which will be useful for our own system.

Chapter 3 Framework for Design of an Autonomous delivery System

The previous chapter showed that there is no single way to design an autonomous delivery system. Various types of systems exist, each with their own strengths and limitations, and each better suited to a different delivery environment. For example, systems can be centralized or decentralized, involve a single vehicle type or multiple types, and may use different routing or transfer strategies depending on vehicle capabilities and area characteristics. In this chapter, we aim to answer the following research question:

How can an autonomous B2C last-mile delivery system be designed, using a set of configurable design choices, based on the characteristics of a specific environment?

To answer this question, we introduce a structured decision-making framework that supports system designers in selecting and combining design elements. This framework does not attempt to cover every operational detail. Instead, it focuses on strategic, system level design choices that must be addressed before implementation.

The framework is based on three principles:

- Insights from literature
- Relevance to research scope
- Generality and usability

From these foundations, we identify four design categories that shape the structure of any autonomous delivery system:

- Geographical and Infrastructure considerations
- Demand of the Area
- Depot Characteristics and Placement
- Operational Logistics

We identified the different types of ADVs and learned about the differences between them and the limitations of these ADVs. The autonomous delivery system is based on these ADVs and can be shaped to suit different areas (single depot, two-tier model). This leads us to the first design choice: **geographic and infrastructure considerations**. First, the delivery area should be clearly defined. If we want to apply autonomous delivery to a certain area, we need to know where we need to deliver. Is the service area a closed area? Or there are multiple regions where we need to deliver and is the range of the ADVs sufficient to reach the entire area? In addition, it is also important to know where the delivery points (customers) are and whether they can be reached (directly) by ADVs.

If we know the answer to these questions, we need to know the demand of the location. As we learned from the literature, UAVs and S-ADVs currently have a limited capacity. Therefore, which ADVs (and especially how much) to use may depend on the demand in the delivery area. This leads us to the next design choice: **demand of the area**. The demand of the area includes the volume, frequency and distribution of parcel deliveries that need to be delivered to the customers. Areas with high-demand may require more frequent deliveries or larger capacity vehicles, while low-demand areas might benefit more from fast lower capacity vehicles or one large delivery route. Here it is also necessary to determine which type of parcels will be delivered by ADVS, since (at present) not everything can just be delivered by ADVs (think sofas or refrigerators).

Once we have a good idea of the delivery area and how many packages need to be delivered there, we can think further about placing the depot(s). This leads us to the third design choice: **Depot characteristics and placement**. If the service area is small enough and all customers can be served by the ADVs, a central depot might be the best choice. If not, decentralized depots or even mobile depots should be considered.

Finally, from the literature we obtained information on modelling and solving a vehicle routing problem. For example, a problem can be formulated with homogeneous or heterogeneous fleet, certain capacity and time windows. We also learned that there are different ways to solve a problem (exact, heuristics or metaheuristics). In the final design choice: **operational logistics**, we address these choices. Most of these operational logistics decisions can also be seen as experimental factors, since it might not be known whether for example a homogeneous or heterogeneous fleet is better.

To support the application of the framework in practice, a simulation tool has been created that allows users to configure their autonomous delivery system according to the four previously mentioned categories and evaluate the systems expected performance. This tool serves as a link between the system design and operational assessment.

By inputting parameters such as the service area, demand of the area, depot location, vehicle allocation/configuration and implement operational logistics, we can simulate daily delivery operations. Using the fleet configuration and chosen routing strategies, the tool simulates delivery dynamics and produces performance metrics based on the users' interests.

A menu of choices

3.1 Geographical and Infrastructure considerations

- **3.1.1 Defining the Delivery Area**
  - **A) Basic Boundary** (*Single enclosed shape)*
  - **B) Multi-region model** *(Multiple sub-regions)*
  - **C) Dynamic Regional Adjustments** (*Continuously adjust multiple sub-regions)*
- **3.1.2 Delivery Locations**
  - **A) Predefined General Delivery Locations** *(Fixed predefined locations)*
  - **B) Customized Delivery Locations** *(Based on Accessibility & Reception Needs)*
  - **C) Dynamic Delivery Locations** *(Real-time Adjustments)*
- **3.1.3 Delivery Network Complexity**
  - **A) Basic Uniform Network** *(All vehicles can use same network)*
  - **B) Vehicle-Specific Network Offline Data** *(nodes and edges have predefined attributes)*
  - **C) Vehicle-Specific Network Offline & Online Data** *(Includes real-time updates in addition to static offline data)*

3.2 Demand of the Area

- **A) No demand data available** *(use country/regional averages per capita)*
- **B) Partial data available** *(Extrapolate based on limited sample data)*
- **C) Full data available** *(Use actual demand distribution patterns)*

3.3 Depot Characteristics and Placement

- **A) No Depot** *(Direct Supplier-to-Customer Delivery)*
- **B) Single Centralized Depot** *(Low Complexity, High efficiency in small areas)*
- **C) Multiple Decentralized Depots** *(High Scalability, Faster Deliveries)*
- **D) Mobile/Moving Depots** *(High flexibility, Adaptive Operations)*

3.4 Operational Logistics

- **3.4.1 Vehicle Characteristics**
- **3.4.2 Fleet Composition**
- **3.4.3 Routing Strategy**
  - **A) Exact Optimization** *(Exact mathematical solver for optimal routing)*
  - **B) Constructive Heuristic Only** *(Fast, predefined routing)*
  - **C) Constructive + Improvement Heuristic** (Improves initial routes with heuristic refinement)
  - **D) Constructive + Metaheuristic optimization** (Improves initial routes with metaheuristic refinement)

3.1 Geographical and Infrastructure Considerations

The first elements to consider are the geographical and infrastructure considerations. Designing a logistical delivery system starts with defining the area and the locations that need to be visited. In this context, we focus on the geographical area that will be the focus of the study and the infrastructure that is located there. The aim is to model the area two-dimensionally and visualize the delivery routes in that area.

3.1.1 Defining the area

One of the most important decisions in the design of an autonomous delivery system is defining the geographic area in which you want to deliver. The boundaries of the area affect routing choices, depot placement, and the scope of operations. The more complex the area (due to size, layout, or traffic conditions), the more challenging it becomes to design a flexible and efficient system.

As explained in 2.1 Parcel Delivery Process in the Netherlands (PostNL), parcels are distributed by using different sorting centers/depots. These depots serve as intermediate points to organize and dispatch deliveries more efficiently. To manage this, urban areas are often divided into smaller regions or zones, with each served by its own depot and delivery fleet. Variables such as population density, delivery volume, and physical layout are important when defining these zones.

Several studies have researched how to divide a delivery area into sub-regions to optimize routing and resource allocation. Huang et al. (2018) provide a detailed review of how two-echelon delivery systems can be optimized by considering time windows, depot capacities, energy use, and variable demand.

In this study, the delivery area is represented as a two-dimensional space enclosed by a shape defined by coordinate points (x, y). For most vehicle types, this 2D layout is sufficient. However, the inclusion of UAVs adds a third dimension (altitude) to the model. UAVS are not restricted by road networks but are subject to height regulations and potential no-fly zones. Although we take vertical constraints into account, our model assumes no hard upper altitude limit. This allows us to leave out the locations of high buildings or trees which interfere with the UAV delivery, as we can simply assume that the UAV should fly above them.

The area can take on different forms:

- **Polygonal shapes:** Triangle, rectangle, or other polygons with straight sides.
- **Circular boundaries:** Possible alternative where the boundary follows a curved perimeter.
- **Multi-polygons:** A combination of multiple polygons, which we exclude in this study to ensure the delivery network remains cohesive and connected.

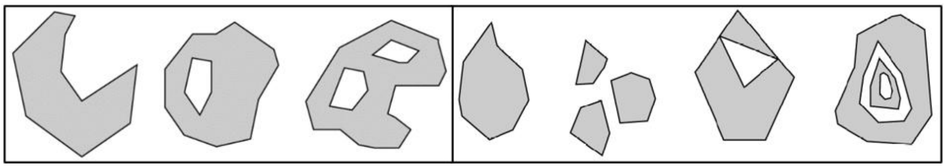

Figure 12: Examples of polygon (left) and multipolygons (right), taken from (Understanding Difference Between Polygon and Multipolygon for Shapefiles in QGIS?, n.d.)

Table 2: Design Choice: Defining the area

Option	Description	Best used when….	Complexity	Required Data	System Adaptability
A) Basic Boundary	Single enclosed zone with no internal subdivisions	Delivery area is relatively small, uniform, or low in demand	Low	Basic Area boundary & demand estimate	Low
B) Multi-Region Model	Area divided into sub-regions	Larger delivery area, non-uniform demand	Medium	Population, infrastructure, demand per region	Medium
C) Dynamic Adjustments	Adaptive sub-regions based on live data	Variable demand or high-uncertainty areas	High	Actual/live data on demand/infrastructure	High

With choice A, the user does not need to perform any further steps, the area is defined and depending on the other design choices, a certain autonomous delivery system is used in this area. In contrast, choices B and C require dividing the region into smaller areas. This can be based on professional expertise or based on methods proposed in literature, such as Huang et al. (2018). In case of C, the continuous adaptation is highly dependent on the availability and frequency of relevant data. It will then be necessary to periodically reassess whether the region's distribution is still correct.

3.1.2 Defining Delivery Locations

After the geographical area is defined, the next step is to determine where parcels will be delivered within that area. The design of delivery locations greatly affects routing efficiency, accessibility for different vehicle types, and ultimately customer satisfaction. Like regional design, delivery location strategies can vary in complexity.

Table 3: Design Choice: Delivery Locations

Option	Description	Best used when….	Complexity	Required Data	System Adaptability
A) General Locations	Fixed drop-off points	High density or centralized locations	Low	Customer address list	Low
B) Customized Locations	Adjusted to accessibility	Mixed urban/residential zones	Medium	Customer address & access constraints per address	Medium
C) Dynamic Locations	Live updates for drop location during the day	Time-sensitive, flexible user preferences	High	List of (multiple) addresses per customer, possibly GPS location	High

With choice A, the user does not need to take any further steps, since the locations are treated with no access constraints. However, this may result in slightly different or inconvenient drop-off locations per vehicle type (an R-ADV might be unable to reach the front door, while a drone may only deliver to the backyard). Choice B specifies the access constraints and allowing the system to link customers with a most suitable delivery method based on their location. Choice C builds on this by dynamically adjusting the delivery location and vehicle.

3.1.3 The delivery network

The delivery network of a logistical system is the foundation of a (autonomous) last-mile delivery system. A network can be defined as:

"A network is simply a collection of connected objects. We refer to the objects as nodes or vertices and usually draw them as points. We refer to the connections between the nodes as edges and usually draw them as lines between points." (*An Introduction to Networks – Math Insight,* n.d.)

Nodes of delivery network

The nodes in the delivery network represent points within the network where decisions are made. They represent a point on earth defined by its latitude and longitude. Each node can be:

- Intersection – where multiple routes or roads converge
- Dead End – terminating point of road
- Delivery Location – final destination of parcels

Edges of delivery network

The edges in the delivery network represent the roads in a defined area. Depending on infrastructure and regulations, edges may vary in suitability for different types of ADVs. In this study, four categories are defined:

- Suitable for S-ADV – accessible only to S-ADVS (such as pedestrian pathways)
- Suitable for R-ADV - accessible only to R-ADVs (regular roads)
- Suitable for both S-ADV & R-ADV – shared infrastructure
- Not suitable – inaccessible or restricted paths for either vehicle type

Table 4: Design choice: Delivery network

Option	Description	Best used when….	Complexity	Required Data	System Adaptability
A) Basic Uniform Network	All vehicles use the same network	Infrastructure is small, highly accessible, and suitable for all vehicle types	Low	Basic road network	Low
B) Vehicle-Specific Network (Offline)	Static access rules per vehicle type	Area contains a lot of infrastructure constraints	Medium	Vehicle access road classification	Medium
C) Vehicle-Specific Network (Online & Offline)	With real-time updates (traffic, weather, etc..)	Dynamic/high-traffic environments	High	Static + real-time network data	High

With choice A, no further steps are needed. All vehicles share the same infrastructure and access rules. In reality, this might result in suboptimal delivery performance, such as inconvenient drop-off locations. With choice B, access constraints are defined per vehicle type. This allows for better alignment between vehicle capabilities and infrastructure characteristics. Finally, choice C introduces the highest complexity by incorporating both offline and real-time updates for the edges.

3.2 Demand of the area

Once the geographical area is defined with its network and delivery locations, the next step is to estimate the demand in this area. Determining depot capacity, optimizing vehicle allocation, and creating effective routing strategies all depend on an understanding of the demand (distribution). However, since private companies often keep parcel delivery data private, demand estimation needs to rely on alternative (or inaccurate) sources.

Table 5: Design Choice: Demand of the area

Option	**Description**	**Best used when….**	**Complexity**	**Required Data**	**System Adaptability**
A) No Demand Data	Use national/regional averages per capita	No historical demand data is available	Low	Demographics & average demand rates	Low
B) Partial Data	Extrapolation from sample demand data	Only partial data is available, or pilot/test areas	Medium	Partial demand data	Medium
C) Full Data	Use actual demand from historical data	All data is available	High	Customer-level or area-wide demand data	High

With choice A, where no historical demand data is available, the system designer should rely on indicators such as population density, demographics, or national e-commerce statistics to estimate parcel volumes. While this approach is less accurate, this estimation method provides a foundation for approximate assessments of the system's scale and feasibility. With choice B, the demand can be scaled from a small period of time to approximate full-area demand. If regional seasonality is available, this would be the preferred choice, or it must come from another source. This option gives a much better idea of the demand of the area. Choice C provides the system designer with real historical data, which allows for the most accurate modeling. The routing and capacity optimization will be much more precise.

3.3 Depot Characteristics and Placement

After estimating the demand in the area, the next step in designing an autonomous delivery system is to establish a depot strategy. Depots function as hubs for vehicle storage, charging, and parcel sorting. An important decision to make is whether to make use of a single depot or multiple depots (or maybe none), which depends on the size and distribution of the demand across the service area.

If a depot is needed, it is convenient to place it in a central location from which both ground ADVs (R-ADVs and S-ADVs) and UAVs can depart from and reach as many customers as possible in the service area.

In this research, we assume that:

- A depot can store and charge multiple vehicles (R-ADVs, S-ADVs and UAVs).
- It functions as a distribution hub where parcels are sorted.
- Employees work at the depot to oversee safety and assist in loading/unloading vehicles.

Table 6: Design Choice: Depot Strategy

Option	Description	Best used when….	Complexity	Required Data	System Adaptability
A) No Depot	Last dispatch point is in (or close to) service area or the supplier (supplier is the last dispatch point) can deliver directly	Supplier with low-volume and small service area. Last dispatch point allows for delivery	Low/Medium	Last Dispatch point or supplier needs the space	Low
B) Single Centralized Depot	One hub for all deliveries	Compact area or centralized operations	Low/Medium	Central Depot Location	Low
C) Multiple Decentralized Depots	Separate depots for regional operations	Larger area or dispersed demand	Medium	Depot Placement Optimization	Medium
D) Mobile/ Moving Depots	Depots that (can) reposition daily	Highly flexible or shifting demand areas	High	Mobile depot routing and scheduling	High

For choice A, the need for a depot depends on the company's structure and the location of the final dispatch point. In cases where the company manufactures or sells its own product and operates in a small service area, it might choose to directly deliver from its own facility (such as retail store or restaurant). Additionally, if the last dispatch point is already located in or near the service area, a separate depot might not be necessary. In such cases, deliveries can be made directly from the existing facility.

In case of choice B, one fixed depot serves as the main hub for all deliveries. The middle mile delivers the parcels to this depot, which in turn handles the deliveries. Normally placed in a central location of the service area, one fixed depot is simple to manage with economies of scale in operations and charging. Suitable for relatively small, well-defined service areas.

In case of choice C, several depots are located around the service area to reduce travel distances and even improve the coverage of the area. It requires careful placement of these depots to balance demand and optimize routing. Suitable for larger, high-demand areas with multiple delivery hotspots. Finally, in case of choice D, the Depots are not fixed but relocated based on demand patterns. It can be implemented using mobile storage units, trucks as temporary depots or modular container hubs. Suitable for dynamic urban environments with constantly changing demands.

3.4 Operational Logistics

In this subsection, we will delve into the operational logistics that are essential for the efficient functioning of an autonomous delivery system. These logistics include decisions such as the allocation of vehicles, routing, scheduling, and fleet management. Since we have now defined and decided what the service area is and its infrastructure, the demand of the area and the depot placement. Now, we can focus on operational logistics. It must be noted that many of these choices/decisions can be inputs for the experiments later on. For instance, a key experimental question may be: *Which fleet configuration best meets the performance of the system?*

3.4.1 Vehicle characteristics

The network described earlier, will be used by several types of vehicles described in the literature review presented in 2.3 Autonomous Delivery Solutions. As we learned from the literature review, we can classify the ADVs into three categories (S-ADV, R-ADV and UAV). Among them there are many different designs and types, of which the system designer is not yet sure which will fit best in the system. Or maybe there is only 1 option for each type, and we need to see if these types meet the expectations of the delivery system. For this purpose, we identified important characteristics for each vehicle category, such as speed, range and capacity (based on data from the literature). These should be quantified and used to integrate the ADVs to the logistics system. Some examples of characteristics are listed in the table (Table 7) below.

Table 7: Example of ADV characteristics

Characteristic	Description	Impact on Delivery System
Average Speed	The vehicle's typical operational speed	Affects travel time and total deliveries per shift
Load time	Time required to load parcels at the depot	Determines depot efficiency and vehicle turnaround
Unload time	Time required to deliver a parcel to the customer	Impacts total stops per trip and overall service time
Maximum Range	The total distance the vehicle can travel	Influences vehicle allocation and depot placement
Storage Capacity	Maximum number of parcels a vehicle can carry	Defines how many stops a vehicle can make per trip
Energy Capacity	Total energy storage available in the vehicle	Limits range and determines recharge needs
Energy Consumption	Energy usage while the vehicle is in transit	Impacts operational efficiency and battery life
Charge speed	Time required to recharge the battery	Determines downtime and depot efficiency

The specific value of these characteristics should be based on the technology that will potentially be used in the autonomous delivery system. In the literature review, we went over some of the vehicles currently under development/in operation. These characteristics will also serve as input parameters in the simulation model developed for this research, allowing the evaluation of various system configurations under different conditions.

3.4.2 Vehicle allocation: Homogeneous vs. Heterogeneous Fleet

Vehicle allocation involves determining how many and which types of vehicles are assigned to each depot within the delivery network. This decision is influenced by factors such as the expected demand, the capacity of the depots, and the geographical characteristics of the area as well as the budget of designing this delivery system.

Table 8: Design Choice: Fleet configuration

Option	Description	Best used when….	Complexity	Required Data	System Adaptability
A) Homo-geneous Fleet	One ADV type	Simplicity and ease of operation	Low	Vehicle specifications	Low
B) Hetero-Geneous Fleet	Multiple ADV types	More flexible delivery capability, different accessibilities per customer	Medium	Vehicle specifications & Customer/Network accessibility	Medium

3.4.3 Routing

In designing an autonomous last-mile delivery system, selecting the appropriate routing strategy is an important step. Routing decisions must consider the characteristics of the ADVs, structure of the delivery network, and the possible constraints of the environment. This section provides a guideline for selecting an appropriate routing strategy by considering these elements.

1 Vehicle-Specific Routing Considerations

The choice of routing strategy depends largely on the type of vehicle being used and their operational constraints. As we mentioned earlier, we categorize ADVs into three groups:

- **Sidewalk Autonomous Delivery Vehicles (S-ADVs)**
- **Road Autonomous Delivery Vehicles (R-ADVs)**
- **Unmanned Aerial Vehicles (UAVs**

2 Network Structures and Routing Implications

Besides vehicle characteristics, the structure of the delivery network plays an important role as well in selecting a routing approach. In general, autonomous delivery networks fall into the following categories (see 2.5 Modelling and Problem Formulations of Autonomous systems):

- **Direct Delivery (ADV-Only Delivery):** ADVs travel directly from a central depot to customers, requiring efficient vehicle routing to minimize travel distances.
- **Two-Tier Delivery:** Vehicles transport first transport parcels to decentralized depots or micro hubs before final delivery to customers.
- **UAV-Aided Delivery:** Combines road-based vehicles with UAVs, where a vehicle carries parcels to an area before deploying UAVs for final delivery. This approach requires vehicle and UAV synchronization algorithms to optimize.
- **S-ADV-Aided Delivery:** Combines road-based vehicles with S-ADVs, where a v vehicle carries parcels to an area before deploying S-ADVs for final delivery. Again, this approach requires vehicle and S-ADV synchronization algorithms to optimize.
- **R-ADV-Aided Delivery:** Uses the platoon model to transfer R-ADVs to an ADV friendly zone.

3 Selecting an Appropriate Routing strategy

The complexity of the routing problem varies based on the number of vehicles, customer demands, and constraints. The appropriate problem-solving approach should be used:

Table 9: Design Choice: Routing strategy

Option	Description	Best used when….	Complexity	Required Data	System Adaptability
A) Exact Optimization	Solve to optimality (e.g. MILP)	Small problem instances with high accuracy needs	High	Full Problem Formulation	Low
B) Constructive Heuristic	Simple rule-based initial routes	Fast, scalable for real-time	Low	Basic delivery inputs	Low
C) Constructive + Improvement Heuristic	Iterative improvement after initial route	Trade-off between speed and quality. Might not achieve optimality	Medium	Heuristic input & tuning	Medium
D) Constructive + Metaheuristic	Advanced metaheuristics after initial route	For larger or more complex systems, with high accuracy needs	High	Metaheuristic configuration	Medium

Each of these approaches is used to generate routes for a problem instance. However, there is also another option which includes continuous real-time route optimization. Which means that throughout the day, each time an ADV returns to the depot to pick up another parcel, the routes are optimized again to see whether the initial routing strategy is still the best. This can also be done with each of the four options.

3.5 Evaluating the model

While the routing optimization discussed above focuses on finding optimized vehicle routes, we also need to evaluate how well these routes perform in dynamic, real-world environments. Simulations incorporate uncertainties such as customer locations, probabilistic service times per customer and other unexpected disruptions, allowing for a nice method of evaluating the route strategy. For this research, we have created a simulation model which simulates a single day of deliveries in the service area chosen by the system designer. Appendix E) Model Documentation outlines more information about the simulation model.

3.5.1 The model

The simulation model evaluates how different ADV configurations perform under varying geographic, demand, and operational logistics scenarios. It helps system designers assess KPIs such as delivery success rate, energy usage, delivery times, and routing efficiency. The basis of the simulation is built upon the *OpenTripModel*, an open-source dictionary for modelling logistics (*About Open Trip Model*, n.d.). The model is structured using a simplified version of the *OpenTripModel (OTM)*. OTM separates logistics into two layers:

- **Entities** are the building blocks of the system. These include **Vehicles** (S-ADVs, R-ADVs, UAVs and Electric vans), **Locations** (Depot or Customer), **Routes** and **Trips**. They store static or planned information, for example a vehicle has an average speed, capacity and battery size. In our model, these static entities sometimes also have properties that can change during the simulation, such as the current battery percentage of a Vehicle.
- **Actions** describe what happens during the simulation. Each action is linked to one or more entities and has an expected duration based on input parameters. For example:
    - A charge action connects a Vehicle to a Location (Depot)
    - A move action links a Vehicle to a Route between two Locations
    - An unload action connects a Vehicle to a Location (Customer)

Each Trip combines a sequence of actions into one delivery round to represent the delivery of the parcels. A Trip is linked to a Vehicle, and every action inside the trip describes what the vehicle is doing step-by-step. Each action has an expected duration which allows the model to estimate the duration of the trip in total. Each action also has a lifecycle and success property, which shows whether the action was (successfully) executed. In this way, actions are the link between vehicles, locations, trips, route, and can be used for evaluation of the performance of the system.

Monte Carlo Simulation

To make the simulation more realistic, we incorporate a Monte Carlo simulation component. While each action has an expected duration at the start of the day, in reality, these durations will vary due to uncertainty in traffic and customer behavior. The Monte Carlo Simulation helps capture this randomness by sampling from probability distributions instead of using fixed values.

- **Unload Action** are based on the vehicle's unload time and the service time at the customers location. The service time represents the time a customer takes to come outside and interact with the ADV. This varies each time, so we model it using a probability distribution

- **Move Action** expected durations are calculated from the distance between two points and the vehicle's average speed. However, real-world scenarios such as congestion introduce variability, so we apply randomness to the actual travel time using a distribution

The exact probability distributions we used are discussed in 4.7 Evaluating the model, where we apply this framework to the campus case.

3.5.2 Conceptual Model

This subsection explains how the simulation operates from start to finish. It outlines the logical flow of a single simulation run, showing how inputs are used to create trips, how parcels are delivered, and how outputs are collected. We refer to the entities and actions as classes from this point on.

Input Configuration

At the start of the simulation, the user provides all required inputs:

- Number and type of vehicles (UAVs, S-ADVs, R-ADVs, or a mix)
- Vehicle characteristics (e.g., speed, capacity, battery, energy consumption
- Service area (e.g. University of Twente Campus)
- Depot location (coordinates)
- Number of deliveries to make (e.g. 106 customers, each with one parcel)
- Time window settings

These inputs are flexible and allow the user to experiment with different fleet setups, demand levels, and delivery constraints. Not forgetting, the service area of the delivery system. The service area as an input means indeed that the model is suitable for different areas if the designated service area is defined in OpenStreetMap. If this is not the case, since OpenStreetMap is open source, the system designer can add a new service area in OpenStreetMap themselves.

Initial class creation

Each class (Vehicle, Location, Trip, Route, Action) has a formula to create this class where the user can enter the input values in to create it. Using the input data, the model creates the following classes:

- **Vehicle class:** Each vehicle is created with its type (e.g., UAV), and all its operational characteristics like speed, range, and capacity
- **Depot (Location class):** The depot is added first, based on its coordinates and operating hours.
- **Customer Locations (Location class):** The customer locations are created based on their coordinates, demand, time windows and service time.

The number of customers are determined based on the defined service area and the types of buildings selected by the user within the area. The service area is extracted from OpenStreetMap data, which includes various building categories such as 'dormitory', apartments' or 'academic'. It is up to the system designer to decide which buildings are eligible for parcel delivery.

At this point, the simulation has all the necessary vehicles, one or more depot(s), and the customer locations. This means that the customer can be allocated to vehicles.

Allocating the customers

Once the initial classes are created, it is time for allocating the customers to vehicles. This is based on the delivery logic implemented in the simulation.

The vehicle’s capacity and the routing algorithm implemented by the user determine how many customer locations will be visited in a single trip. Each trip consists of multiple actions, which describe what the vehicle will do. The trips are always sequenced in the following order:

- **Charge Action:** The vehicle charges (at the start of a trip) to make sure it has sufficient amount of battery
- **Load Action:** The parcels are loaded into the vehicle
- **Move Action:** The vehicle travels from one location to another
- **Wait Action:** In case the vehicle arrives before customer’s time window opens, it must wait
- **Unload Action:** The parcel is delivered to the customer
- **(Repeat Move/Wait/Unload):** These actions are repeated for each stop
- **Return Move Action:** Vehicle moves back to the depot

Running the simulation day

Once all customers are allocated to trips, the simulation day can begin. Deliveries officially start at the opening time of the depot. A SimulationClock tracks the time, and actions/trips are triggered based on their planned start times.

The first scheduled trips for each vehicle have a status “requested”. As soon as the SimulationClock reaches the planned start time of the “requested” trips, they are prepared by updating the expected durations to the real durations (Monte Carlo), and the trip is executed. During execution, the simulation updates the relevant states, for example the vehicle’s battery level, current location, and the progress of the action.

After the vehicle completes its trip, the trips status is set as “completed” and the next scheduled trip for the vehicle is labeled as “requested”. This loop continues until the SimulationClock reaches the depot’s closing time. By the end of the day, the model outputs key performance indicators based on the actions.

This concludes the general explanation of the simulation model. In the next chapter, we apply the framework (including the simulation model) to the University of Twente campus case. There, we go into more detail about how the simulation model was set up for this case, how the inputs were handled, what the outputs are, and how reliable and realistic the model is for answering our research questions.

3.5.3 Assumptions of the model

This simulation model is built on several core assumptions derived from earlier chapters, historical context, and specific coding choices made during implementation. some assumptions can be changed by adjustments in the underlying code, such as the method of customer sampling. The following list covers the main and most important assumptions of the model:

- **Independent days**
  Each day is independent, meaning it is not affected by the day before and does not consider the next day.
- **Independence of time of day**
  The model does not account for variations in campus activity throughout the day. Pedestrian and vehicular traffic tends to increase during class changes. The model assumes that the day is evenly distributed.
- **Static Geographical input:**
  The campus delivery area with its nodes and edges is assumed to remain static throughout the simulation. So, all infrastructure elements are available all the time, with no temporary shutdowns or maintenance activities.
- **Customer generation and sampling**
  The University of Twente campus consists of multiple buildings that vary in size and number of residents (potential customers). However, since detailed information on the distribution of residents per building is not publicly available, this distribution is not modeled. Also, the number of residents per building still does not give us the necessary information (Number of residents does not equal number of potential customers). Instead, it is assumed that all buildings have an equal probability of being selected as customer location. Customer generation is performed by sampling building locations with replacement, meaning that each sampled building represents a single customer with a (standard) demand of one parcel. Sampling with replacement allows for the possibility that a single building is selected multiple times, representing multiple customers per building.
- **Generic Distribution of Waiting Times:**
  Each delivery vehicle (UAV, R-ADV or S-ADV) drop off has the same generic distribution of waiting times at the customer, no matter the location. So, we assume customers must be present at parcel drop off for every delivery vehicle.
- **Average Vehicle Speed:**
  The model assumes a constant average speed for the autonomous delivery vehicles.
- **Fixed Battery and Recharging Parameters:**
  Each vehicle is assumed to have a consistent battery capacity with a predetermined recharging time and battery consumption while moving. This means we neglect the (small) amount of battery which the vehicles consume on standby.
- **No Vehicle Interference**
  The simulation does not model interactions between (delivery) vehicles. Each vehicle operates independently based on the schedule and routing algorithm.
- **Static Routing Strategy:**
  The routing strategy is based on the expected durations of each action; the schedule does not adapt dynamically during a simulation run if the real durations are updated.
- **Delivery Success:**
  The model assumes that all deliveries have a certain delivery success. Failed deliveries stay in the vehicle and are not delivered at another time. The idea is that customers can pick up the parcel themselves at the warehouse in case of a failed delivery.

3.6 Conclusion

Chapter 3 focused on transforming the theoretical foundation from literature into a practical framework for designing autonomous last-mile delivery systems. At the start of the chapter, we wanted to answer the following research question:

How can an autonomous B2C last-mile delivery system be designed, using a set of configurable design choices, based on the characteristics of a specific environment?

Recognizing that each urban environment has unique logistical characteristics, this chapter introduced a "menu of design choices" that can be used for different scenarios/environments. The idea is that a system designer can use this framework to make/consider design decisions for their own area. This framework does not attempt to cover every operational detail. Instead, it focuses on strategic, system level design choices that must be addressed before implementation. These design decisions can then be evaluated by the simulation model.

The framework consisted of four major design decisions:

1) Geographical and Infrastructure considerations: define the delivery area and delivery network
2) Demand Estimation quantifies the delivery load based on data
3) Depot Configuration: addresses the placement of depots
4) Operational logistics: covering fleet composition, vehicle characteristics, and routing strategies

Each design component was described, highlighting how the decisions affect the performance of the delivery system. The design framework serves two purposes:

1) It structures the process of designing autonomous delivery systems
2) It enables evaluation across different scenarios/areas using the simulation model

In the next Chapter, we apply this framework to the case of the University of Twente Campus.

Chapter 4 Application of the framework to the UT campus

In this chapter, we'll apply the framework of the previous chapter to design the autonomous delivery system for the University of Twente Campus. Just like in chapter 3, we will go over each element of the framework and describe the actions in a separate subsection.

4.1 Geographical and Infrastructure Considerations

From section 3.1 Geographical and Infrastructure Considerations, we know that geographic and infrastructural context are very important. In the case of the University of Twente campus, these considerations are particularly important due to its unique layout. This section aims to define the delivery area, identify the specific delivery locations and define the delivery network.

4.1.1 Defining the area

The first step is to define the service area in which the deliveries take place. Since the campus is closed and well-structured environment, this study considers the entire campus as a single area.

Selected approach: Choice A) Basic Boundary

As said before, we treat the campus as a single service area instead of differentiating between residential, academic or commercial areas. The service radius of each available vehicle (S-ADV, R-ADV and UAV) is also large enough to cover the entire campus (Radius of approximately 1500m) without the need for subdivision into different operational zones. The choice to treat it as a single area simplifies logistics, routing, and vehicle allocation.

To accurately define the boundaries of the service area, we use the entire University of Twente campus as the geographical basis, as illustrated in Figure 13. We retrieved additional geospatial data using OpenStreetMap (OSM) to outline the campus surroundings and coordinates. A set of 33 boundary nodes has been identified to define the exact perimeter of the study area. The full list of nodes, along with their coordinates, is provided in Appendix B.

4.1.2 The delivery locations

Now that we have defined the delivery service area, we can identify the delivery locations.

Selected approach: Choice A) Predefined General Delivery Locations

A total of 192 potential delivery locations were selected using OpenStreetMap data. These locations were chosen based on their building types (Residential and Institutional). Residential delivery locations were identified using tags such as “apartments”, “house”, “residential” and “dormitory”. University buildings were identified using tags such as “university”, “school”, and “public”.

Of the 192 locations, 34 are institutional buildings (Red buildings in Figure 13). However, these locations were excluded from the delivery strategy for two reasons:

1. Deliveries to university buildings are handled internally by the university itself through two daily consolidated delivery rounds.
2. A disproportionate share (approximately 80%) of these deliveries is directed to just three buildings, making them unrepresentative of typical last-mile delivery patterns.

By excluding these 34 buildings, the research focuses solely on the remaining 158 residential delivery locations (Blue buildings in Figure 13), providing a more accurate basis for evaluation of autonomous B2C last-mile delivery performance.

The campus consists of various types of buildings, each with different population densities. However, due to privacy concerns, detailed data on individual building occupancy or usage is not accessible. As a result, all delivery locations are treated uniformly in this research.

How Are the Delivery Locations Assigned?

- The University of Twente Campus was mapped using OpenStreetMap (OSM), identifying a set of nodes representing intersections, pathways, and accessible locations
- Each customer's delivery point is linked to the nearest node, ensuring that deliveries are made to the closest locations
- These nodes serve as the delivery points for that customer, regardless of the vehicle type. It might occur that certain locations use the same node as delivery point.

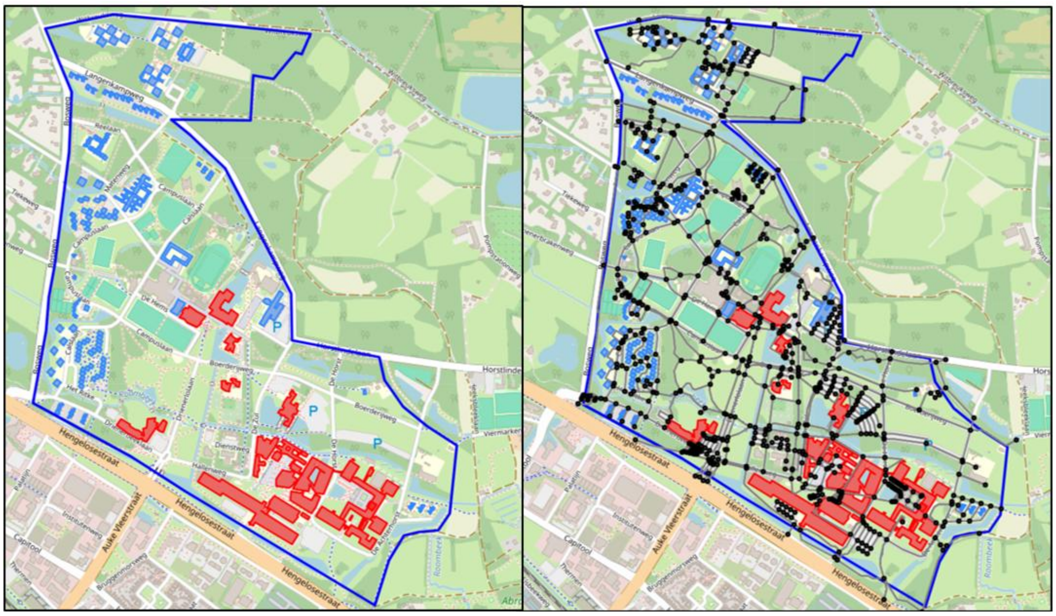

Figure 13: Residential buildings (blue) and campus facility buildings (red), with delivery network plotted on the right

4.1.3 The delivery network

The delivery network is retrieved using OpenStreetMap as well (See Figure 13). This provides a representation of the roads, sidewalks, and pathways. This information, however, is not complete enough to use for our research, so we chose to assume that each road of the network is accessible to R-ADVs and S-ADVs.

Selected Approach: Choice A) Basic Uniform Network with a separate UAV Distance Matrix
The road network is used by the S-ADV and R-ADV, while we created a separate UAV Distance matrix for UAV delivery.

How was the delivery network defined?

1. Extracting Data from OpenStreetMap

- The University of Twente Campus was used as input for extracting data on the nodes and edges. These can be used by all ground vehicles (S-ADVs & R-ADVs).
- The nodes are defined as intersections, dead ends, and delivery points. The edges as roads and pathways that any ground vehicle can use.
- The UAV network is created using python library *Networkx (NetworkX Documentation, n.d.)* which calculates the Euclidean distance between each node and the depot.

4.2 Customer demand

Obtaining detailed information on parcel delivery volumes from commercial logistics providers in the Netherlands is remarkably difficult. Publicly available data is typically limited to annual parcel volumes at national level, with no specification by region or delivery day. As a result, there is no direct information on campus-specific delivery volumes or seasonality patterns. For this research, were able to obtain data (from PostNL) on parcel deliveries at the campus between January 1 until March 14th, 2024. A total of 2500 parcels were delivered to the university facilities and 5000 parcels to the campus residents. We also used an external dataset from the *Final Report: Single-carrier Consolidation – Central London Trial (Clarke et al. 2017)*, which includes parcel delivery of a full year. While this data is from a different country/environment, it serves as an estimation of the general delivery trends and seasonality patterns.

Selected Approach: Choice B) Partial Data Available

Please read Appendix C) Extra demand calculations for more details on the dataset and calculations. To estimate the annual parcel demand on the University of Twente Campus, we begin by analyzing the data from the London Trial. This dataset includes daily delivery data of one or more delivery vans. By summing up the total number of parcels we get to a total of 2,005,728 parcels delivered in a year. Since we have daily delivery data, we can also calculate the weekly package distribution. March 14th falls in the middle of week 11, so we sum the parcel distribution percentages from week 1 to 10 and add half of week 11. This results in 0.194709 (19.47%), so 19.47% of parcels were delivered from January 1 until March 14th in London.

We use this proportion to estimate the total annual parcel volume for campus deliveries. We know that around 5000 parcels were delivered during the first 10.5 weeks of the year. Dividing this by the 19.47% share gives an estimated annual volume:

To estimate full-year parcel demand, we use:

$$\frac{5{,}000}{0{,}194709} = 25.680\, packages\ per\ year\ (PostNL)$$

Since PostNL holds 50% of the market share, the total number of parcels delivered annually to the campus is approximately:

$$25{,}680 * 2 = 51.360\, parcels\ per\ year\ (Total\ amount)$$

To determine the number of parcels suitable for autonomous delivery, it is important to consider size and weight constraints. Back in 2019, Amazon's CEO Worldwide consumer stated that: ***"between 75 and 90% of Amazon deliveries could technically be handled by the UAV"*** (D'Onfro, 2019). Since UAVs have the most restrictive limitations in terms of payload and volume, any parcel that is eligible for UAV delivery is also suitable for S-ADVs and R-ADVs. Therefore, the share of parcels eligible for delivery via autonomous methods can be estimated at 75-90%. In this study we take the most conservative number 75%.

$$51.360 * 0{,}75 = 38{,}520\, parcels\ per\ year\ (Total\ amount)$$

In addition to the weekly distribution of parcels, we also calculated the daily and monthly distributions to gain insight into seasonal variations across different time periods. By combining weekly and daily distributions, we roughly estimate the daily number of parcels delivered.

This enables us to compute key demand statistics, including:

- Median and Average Daily Parcel Volume
- Top 10% busiest days (highest demand periods)
- Lowest 10% least busy days (low-demand periods)

These insights (See Table 10) provide a data-driven understanding of delivery trends, which can help in modelling demand fluctuations and optimizing logistics planning.

Table 10: Median, Average and Range of parcels delivered per day

METRIC	ALL DAYS	TOP 10% DAYS	LOWEST 10% DAYS
MEDIAN	109	172	48
AVERAGE	106	193	46
RANGE	29-306	153-306	29-52

4.3 Depot Characteristics and Placement

At the University of Twente, the depot serves as the central hub for handling autonomous deliveries. This depot will handle vehicle storage, charging and loading activities.

Selected Approach: Choice A – Single Centralized Depot
The centralized approach is chosen based on the following considerations:

- **Campus Suitability:** The University of Twente is a compact, well-defined area with relatively stable parcel demand of ~38,520 parcels per year and ~106 parcels per day.
- **Cost Efficiency:** Centralizing operations reduces infrastructure needs, enables shared use of charging and maintenance facilities. It requires less staff and coordination is easier.
- **Operational Simplicity:** All vehicles start and return to one location
- **Existing location present:** The University of Twente already has a small parcel distribution center named "Garage", located near the Ravelijn (See Figure 14). The building is assumed to be large enough to support the storage, charging and monitoring of the ADVs.

Another option would be to build a new depot and place it at another location on campus. In our case, we assume an existing location, but there are also papers that calculate what the best locations might be such as *Sartika & Gamal*, (2019).

Figure 14: Location of new Depot/"old Garage" (Yellow)

This does also mean that we use a two-tier delivery system (See 2.4 Autonomous Delivery Systems). In this model, a delivery van (for example from PostNL) transports parcels from a regional sorting center (in this examples case, PostNL's depot in Hengelo) to the central campus depot. From there, the parcels are delivered to the customers using the ADVs.

4.4 Operational Logistics

In this section, we focus on the operational logistics of the autonomous delivery system. These are decisions about vehicles, routing, scheduling, and fleet management. These logistics include decisions related to ADV types, routing, scheduling, and fleet configuration. Unlike the previously discussed system design choices (delivery area, demand of the area, and the depot location) these operational logistics choices are not all fixed.

While certain operational logistics might be fixed in another scenario, in our scenario this is not the case. This really depends on each situation and the aim of the research (wishes of the system designer). We will elaborate on each aspect (vehicle characteristics, fleet allocation, and routing strategy) and clarify whether the decisions are fixed for this research.

4.4.1 Vehicle Characteristics (Fixed)

The University of Twente delivery network will be utilized by the three types of vehicles mentioned before: S-ADVs, R-ADVs and UAVs. For modelling purposes, the operational characteristics are treated as fixed inputs. These include their average speed, load time, unload time, service range, storage capacity, noise, battery capacity, energy usage and charge speed. These characteristics influence how each ADV performs in the simulation and determines factors such as delivery time and energy consumption.

Please note that in other research projects, these vehicle parameters could be varied as experimental inputs. For instance, a system designer might want to compare multiple UAV models to identify the most suitable one for the delivery system. The simulation model developed for this research allows for changes in all of the vehicle characteristics.

The autonomous delivery system modeled for the campus case includes the following three ADV types:

- **S-ADV**, based on the Starship Robot
- **R-ADV**, based on the Macrostep parcel locker
- **UAV**, based on the Zipline Drone

These ADVs are chosen because they have (relatively) a lot of publicly available information on their characteristics and have already been tested or used in real-life delivery scenarios. This makes them suitable and realistic candidates for the simulation model. Also, to compare the autonomous system to the current system. We introduce the CargoLEV TC delivery vehicle (Vleugel, 2023). A Light Electric Vehicle (LEV) that is currently being employed in the last-mile logistics in Amsterdam, The Hague, Arnhem and Groningen. This makes it a good candidate for comparisons with the autonomous system.

Table 11: Light Electric Vehicle Characteristics

Characteristic	***Value***	**Source or Explanation**
Speed	*20 km/h*	Estimated guess based on infrastructure
Load time	*60 seconds/parcel*	Estimated guess
Unload time	*180 seconds/parcel*	Estimated guess
Maximum range	*115 km*	Max range (132km) (*DPG Media Privacy Gate*, n.d.). With load probably less
Storage capacity	*300*	6 m^3 storage room, fits approximately 600 shoeboxes (0.34m*0.28m*0.14m). Assuming inefficient use of storage room say 300 units.
Noise	*56 (dB)*	Legal limit (*Waarom Een Elektrische Auto Geluid Maakt | ANWB*, n.d.)
Tank	*11960 (Wh)*	104*115=11960
Energy cons. moving	*104 (Wh/km)*	(*DPG Media Privacy Gate*, n.d.)
Charge speed	*1440 (Wh)*	Assume same charge speed as R-ADV

4.5.2 Vehicle allocation (Experimental)

The fleet allocation is a key experimental factor in this thesis. Since we are not limited by a budget or specific ADV delivery goal, we decided to use the fleet allocation as the main experimental factor. Different experiments will test how varying fleet combinations (homogeneous vs heterogeneous) impact the delivery performance.

In some scenarios, the fleet might be restricted to one vehicle type (such as only UAV delivery) and the system designer might just be interested in the performance of a certain number of UAVs under different conditions. In these cases, the vehicle allocation becomes a fixed variable.

4.5.3 Routing

Finally, we will make the decision on the routing strategy.

Selected Approach: B) Constructive Heuristic Only

For this research, we have decided to base the routing strategy on constructive heuristics only. The focus of this thesis is on how to design an autonomous delivery system, not optimizing the autonomous delivery system at the campus. A heterogeneous electric fleet where each vehicle can make multiple trips is already a complex model to optimize, and we have decided that this would fall out of the scope of this thesis. We use the Two-Tier model, with an ADV depot placed in the center of the campus.

To achieve this, we use Solomon's Nearest Neighbor constructive heuristic, a well-established method for solving the Vehicle Routing Problem (Solomon, 1987). This method efficiently assigns Sidewalk Autonomous Delivery Vehicles (S-ADVs), Road Autonomous Delivery Vehicles (R-ADVs), and Unmanned Aerial Vehicles (UAVs) to different deliveries while potentially considering time windows. Solomon's nearest neighbor's constructive heuristic, however, does not deal with a potential penalty of missing the time window (hard constraint). For this thesis, we have transformed Solomon's cost function to a cost function that factors in a penalty of missing the time windows.

Solomon's Cost Function with Parameterized Weights

Solomon's Cost Function is used for the routing decisions, by evaluating the desirability of each potential next customer during the routing process. This heuristic calculates the cost of visiting a candidate customer *j* after serving customer i, considering multiple factors:

1. Travel distance between the two locations

2. Travel time required to get to this location

3. The urgency of delivery based on the remaining time within customer *j'*s allowable delivery window. For example: Possible delivery to a customer with TimeWindow (09:00-17:00) at precisely 11:00 means a remaining time of 6 hours, making it a less time-critical choice than another customer whose TimeWindow closes at 12:00.

We define a tunable cost function that integrates weighted contributions from each of the core factors. The general form of the cost function is given by:

$$C_{i,j} = w_1 * d_{i,j} + w_2 * T_{i,j} + w_3 * \max\left(0, v_{i,j}\right) + PenaltyCosts * \max\left(0, -v_{i,j}\right)$$

Where:

- Time Window Customer *i* $[e_i, l_i]$
- Service time at customer *i* s_i
- Service beginning at customer *j* = b_j = $\max(e_j,\ \ b_i + s_i +\ t_{i,j})$
- $d_{i,j}$ = Distance between the current location *i* and customer *j*.
- $T_{i,j}$ = Travel time required to reach customer *j* computed as: $T_{i,j} = b_j - (b_i + s_i)$
- $v_{i,j}$ = Slack time computed as: $v_{i,j} =\ l_j - (b_i + s_i +\ t_{i,j})$
- $PenaltyCost$ = 10000 (Standard)
- $C_{i,j}$ = Composite cost of selecting (next) customer *j* after visiting location *i*.
- w1 = **Weight for travel distance** (default: 0.5, TimeWindowDelivery: 0.05)
- w2 = **Weight for travel time** (default: 0.4, TimeWindowDelivery: 0.05)
- w3 = **Weight for slack time** (default: 0.1, TimeWindowDelivery: 0.90)
- $w_1 + w_2 +\ w_3 = 1$

The decision on the weights is as follows: For delivery on days where each customer has the same window of delivery (09:00-17:00), the slack time is not relevant because it's the same for each customer. During the regular days without TimeWindows, the weights are thus set on distance/time. In experiments where meeting TimeWindows is important, this slack time weight is increased a lot.

This approach allows the routing algorithm to be adapted dynamically based on:

- Delivery urgency (increasing w_2 prioritizes travel time)
- Battery efficiency (reducing w_1 for UAVs can prevent over-prioritization of short distances).
- Time window sensitivity (adjusting w_3 ensures compliance with delivery deadlines)

```
FUNCTION SolomonsNearestNeighbor(customers, vehicles, depot, w1, w2, w3, w4):
    UNALLOCATED_CUSTOMERS = customers  # List of unallocated customers
    ALLOCATED_ROUTES = []  # List of completed routes

    FOR EACH vehicle IN vehicles:
        INITIALIZE a new route with the depot as the starting point
        current_location = depot
        remaining_capacity = vehicle.capacity
        current_time = depot.start_time
        route = [depot]

        WHILE UNALLOCATED_CUSTOMERS IS NOT EMPTY:
            best_customer = None
            best_cost = INFINITY

            FOR EACH customer IN UNALLOCATED_CUSTOMERS:
                IF customer.demand > remaining_capacity:
                    CONTINUE

                travel_distance = Haversine(current_location, customer)
                arrival_time = current_time + travel_distance / vehicle.speed

                # Compute slack time
                slack_time = (customer.time_window_close - arrival_time) / 60
                penalty = w4 * MAX(0, -slack_time)

                # Cost function with parameterized weights
                cost = (w1 * travel_distance) + (w2 * arrival_time) + (w3 * MAX(0, slack_time)) + penalty

                IF cost < best_cost AND arrival_time <= customer.time_window_close:
                    best_customer = customer
                    best_cost = cost

            IF best_customer IS NONE:
                BREAK  # No feasible customers, end the route

            ADD best_customer to route
            UNALLOCATED_CUSTOMERS.REMOVE(best_customer)
            remaining_capacity -= best_customer.demand
            current_location = best_customer
            current_time = MAX(arrival_time, best_customer.time_window_open) + best_customer.service_time

        ADD route to ALLOCATED_ROUTES

    RETURN ALLOCATED_ROUTES
```

Figure 15: Algorithm Solomons Nearest Neighbor

4.6 KPI Selection

In this section, we will select the Key Performance Indicators that will be useful for our simulation study. The KPIs identified from the literature review form the total set, but not all of them can be directly measured or used within the simulation framework. Therefore, we must narrow them down to KPIs that can be logically and quantitatively evaluated in our model. The remaining KPIs are selected based on their measurability within our simulation framework and scope of this study. Some of these KPIs are derived from the KPIs in the literature section to better fit our model (Instead of *Delivery Success rate*, we went with *total successful deliveries*).

Table 12: Final KPIs

Final KPIs	
Economic	
Distance (km)	Total amount of kilometers driven
Total Successful Deliveries (#)	Total Orders successfully delivered to customers
Average Delivery Time (mins)	Average time between deliveries
Utilization Rate (%)	Percentage of time the ADV performs a 'useful' action: Loading, Unloading, Moving, Charging
Unplanned Deliveries (#)	The average number of deliveries not being planned. If the ADV does not expect to be able to deliver the package on time, this customer will not be included in the route.
Environmental	
Energy Consumption (kWh)	Measures the total kilowatt-hours per day
Social	
Average Time After TimeWindow (mins)	Measures the average time the customers need to wait/be at their location before the ADV arrives with their package
Late Customers (#)	Average Number of Customers receiving their order too late
Total Minutes Late (mins)	Total minutes of deliveries being late

We want to note that many KPIs can be determined in different ways. Energy Consumption can also be looked at from an economic perspective, because if a certain vehicle consumes a lot of energy this will increase operating costs. The same goes for distance, we also think that distance is somewhat a social KPI. A delivery system that drives a total of 300 kilometers around the campus or one that drives a total of 20 kilometers will certainly be looked at different from the public viewpoint.

4.7 Evaluating the model

To test the performance of the autonomous delivery system on the University of Twente campus, the delivery system is evaluated using the simulation model. From 3.5 Evaluating the model we learned that we need a couple of inputs for our model:

Input Configuration

- Number and type of vehicles (UAVs, S-ADVs, R-ADVs, or a mix)
- Vehicle characteristics (e.g., speed, capacity, battery, energy consumption
- Service area (e.g. University of Twente Campus)
- Depot location (coordinates)
- Number of deliveries to make (e.g. 106 customers, each with one parcel)
- Time window settings

In the next subsection, we will cover all of the inputs we use in our simulation model for the campus case

4.7.1 Inputs of the simulation model

The simulation model is driven by a variety of inputs derived from previous chapters and supplemented with historical data. Essentially, the model can be rewritten and changed on many levels, but the model which is used during this thesis has the following inputs:

- **Geographical Data:** Input of the model is the area “Universiteit Twente”, which is a defined area found on OpenStreetMap (*OSM Universiteit Twente Service Area*, n.d.)
- **Housing tags:** Buildings with the following tags are used as input: House, Dormitory, Apartments.
- **Depot Location:** The depot location is manually chosen in the model with latitude: 52.239623555802815, longitude: 6.853921115398408 and OSMID 7801064041.
- **Road Network:** For the R-ADV and S-ADV, the entire road network is used in the *Universiteit Twente* service area (Instead of using only sidewalks/bicycle lanes)
- **Vehicle Specifications:** Performance parameters for the ADVs such as average speed, storage capacity, battery life, and loading/unloading times as given in Section 4.4.1 Vehicle Characteristics (Fixed) and Appendix D) Vehicle Characteristics
- **Customer Demand:** 43 Low, 106 Average and 196 High (4.2 Customer demand)
- **Driving Time Distribution:** To account for the uncertainty in driving from one location to another, the expected duration (computed from route length and vehicle speed) is modeled using a normal distribution. The distribution has a mean equal to the expected duration and a standard deviation equal to 5% of that expected duration, representing some variability due to factors such as traffic conditions.
- **Customer Pickup Time Distribution (ADVs) – Gamma (α=4, θ=1):** To simulate how long it takes for a customer to come outside and collect a parcel from an ADV, we model this time using a Gamma distribution with shape parameter α=4 and scale θ=1. This results in a mean pickup time of 4 minutes. Starship Technologies (2024b) explains on their website that the maximum time a starship robot waits for customer pickup is 12 minutes. We assume that most of the customers pick up their parcel much faster than 12 minutes, the pickup time follows most-likely a right-skewed distribution. There is no real data on the pickup time of autonomous delivery on campus, so we decided to use this gamma distribution with a mean of four minutes. It is important to know that this pickup time is **on top** of the unload time per vehicle. See Figure 16 on the next page for an overview of all the service times per customer (Unload time + Pickup time)
- **Customer Pickup Time Distribution (Electric van) – Gamma (α=2, θ=1):** We use the same type of distribution for the service time (pickup time) for the Electric Van (human deliverer). Since a human operator will deliver these packages to the houses of the customers, the service time at a customer is much lower. The human deliverer does not need to person to hand the package too. Also, the human deliverer can often just place the package on a place around the house, which saves lots of time. See Figure 16 for an overview of the service times of the vehicles used.

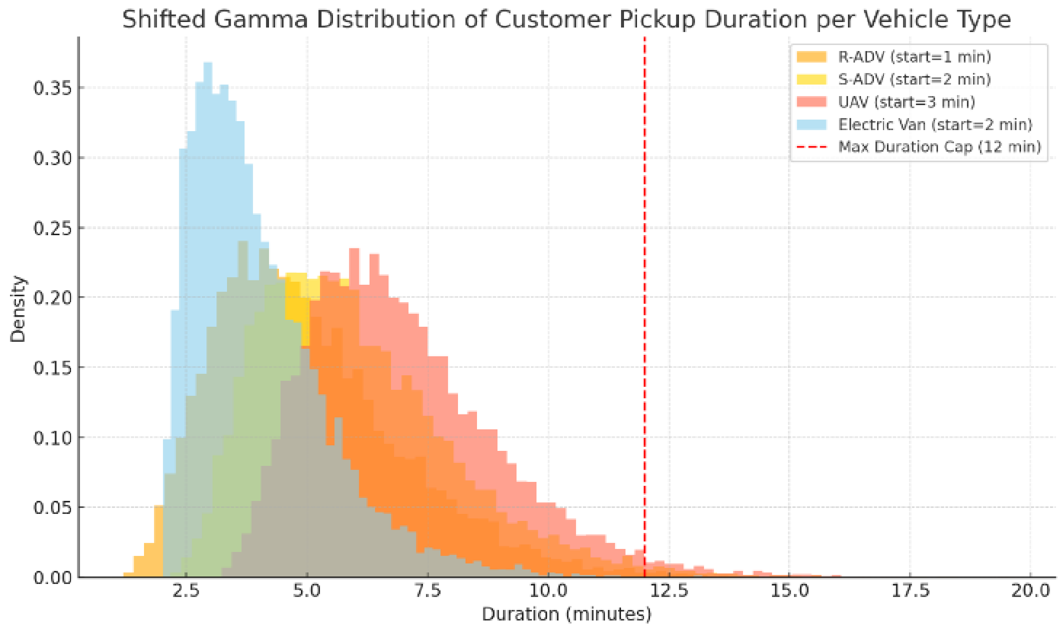

Figure 16: Distribution of waiting times

Besides these inputs, the model also has some additional inputs:

- **Delivery Success:** Each delivery has a predetermined success percentage, for experimentation this is placed at 100%.
- **SimulationClock SpeedUp factor:** The simulation model lets the user choose between Manual Mode and Experimentation Mode. In Manual Mode, the simulation runs in real-time but with a speed-up factor (default is 20x), so the simulation progresses 20 times faster than actual time.

4.7.1 Validity of the simulation model

The model is validated using techniques discussed in Robinson (1997). The problem of our simulation model and routing strategy is that there is no real world for comparison. Just like Robinson (1997) explained in validating models that have no real world for comparison, a model that is valid to one person may not be valid to another. We based our simulation on the Campus case as discussed earlier, but this campus does not have an autonomous delivery system. Despite the lack of real-world data, we employ validation techniques such as black-box validation.

Data validation

The main input from our simulation model is the data from OpenStreetMap, which is a public website accessible to anyone. Using python packages such as OSMnx (Boeing, 2025) and NetworkX (*NetworkX Documentation*, n.d.), the road network of the campus is transformed to a usable delivery network which can calculate the shortest path between two different nodes (and its distance). The packages Networkx and OSMnx are tested by a lot of people worldwide, which gives us the confidence that these packages provide valid data for our delivery network. There is only one possible flaw for this thesis (which can be solved by anyone if noticed!). OpenStreetMap is a public website where basically everyone can change the characteristics of the map (just like Wikipedia can be changed by everyone). In this research, we assume that there is no misinformation (only a lack of information sometimes) in OpenStreetMap, and that if there was anything inaccurate, that the community would take care of it (just like Wikipedia).

White box validation

White-box validation is used to assess the behaviour of small parts of the simulation model. Our simulation focusses on delivering parcels and thus driving from node to node. First, we want to check whether each node and delivery location is correctly depicted on the world map. The customer generation tab in the simulation model offers the possibility to check each location (with a selection of tags). This results in a list of locations with their latitude and longitude.

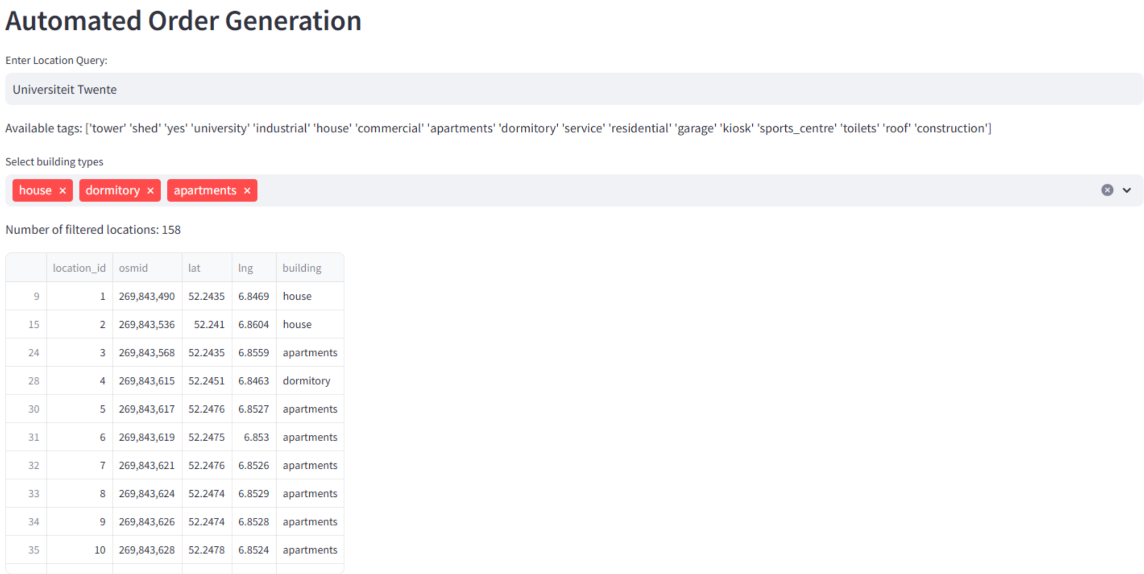
Automated Order Generation

Enter Location Query:

Universiteit Twente

Available tags: ['tower' 'shed' 'yes' 'university' 'industrial' 'house' 'commercial' 'apartments' 'dormitory' 'service' 'residential' 'garage' 'kiosk' 'sports_centre' 'toilets' 'roof' 'construction']

Select building types

house × dormitory × apartments ×

Number of filtered locations: 158

	location_id	osmid	lat	lng	building
9	1	269,843,490	52.2435	6.8469	house
15	2	269,843,536	52.241	6.8604	house
24	3	269,843,568	52.2435	6.8559	apartments
28	4	269,843,615	52.2451	6.8463	dormitory
30	5	269,843,617	52.2476	6.8527	apartments
31	6	269,843,619	52.2475	6.853	apartments
32	7	269,843,621	52.2476	6.8526	apartments
33	8	269,843,624	52.2474	6.8529	apartments
34	9	269,843,626	52.2474	6.8528	apartments
35	10	269,843,628	52.2478	6.8524	apartments

Figure 17: Automated Order Screenshot pt1

This provides the user with a list of possible locations. Each of these locations can be hovered over and check if the latitude longitude and whether the latitude/longitude is correct. These values can be inserted in an online website and checked if they correspond to the location in on the map as well. These packages that were used to display this map are used by so many people that they give the good locations/results. Another thing is whether the distances between nodes are correct. We can take a random set of nodes and connect them to each other with the shortest path. This could be checked in the real world, but because OpenStreetMap is based on latitude and longitude, this already covers this check. The shortest path algorithm can be checked multiple times but again, these packages are very accurate (at the level we of which we use it) so there is basically no need. With the manual mode, other classes such as customers with their demand and TimeWindows can be checked as well.

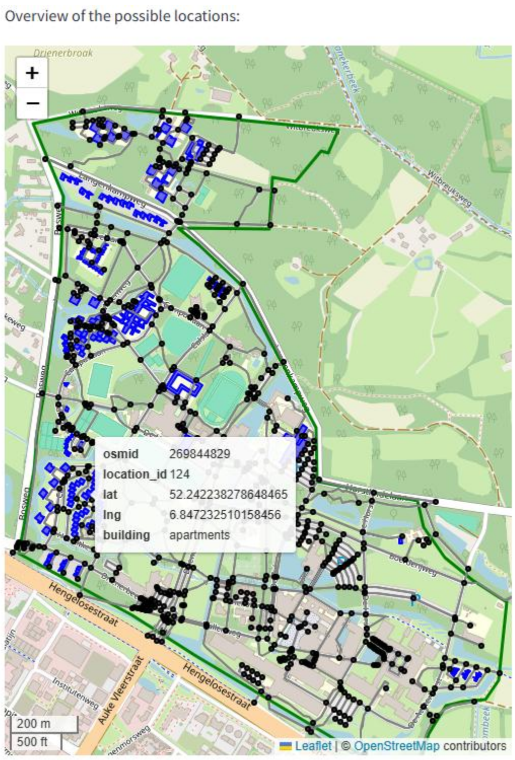

Figure 18: Automated Order Screenshot pt2.

Another thing that should be looked at is whether the service time distribution and the deviation in the driving time. The model can generate multiple random service times, and these can be plotted to see if they fit the distribution. The same goes for the deviation in driving time, the differences can be plotted to see if they only deviate 5%.

Black-box validation

Black-box validation is used to assess the overall behaviour of our model. We check whether given the input parameter, realistic outputs are obtained. It is not possible to validate the ADV deliveries in the simulation model, however, it is possible to validate the deliveries with the electric van. For this validation, we simulated the deliveries with the electric van for 20 and 100 customers. We used the input parameters discussed in this chapter (Electric van characteristics), so also Solomons Nearest Neighbor Algorithm. Figure 19, shows the delivery routes of 20 and 100 customers.

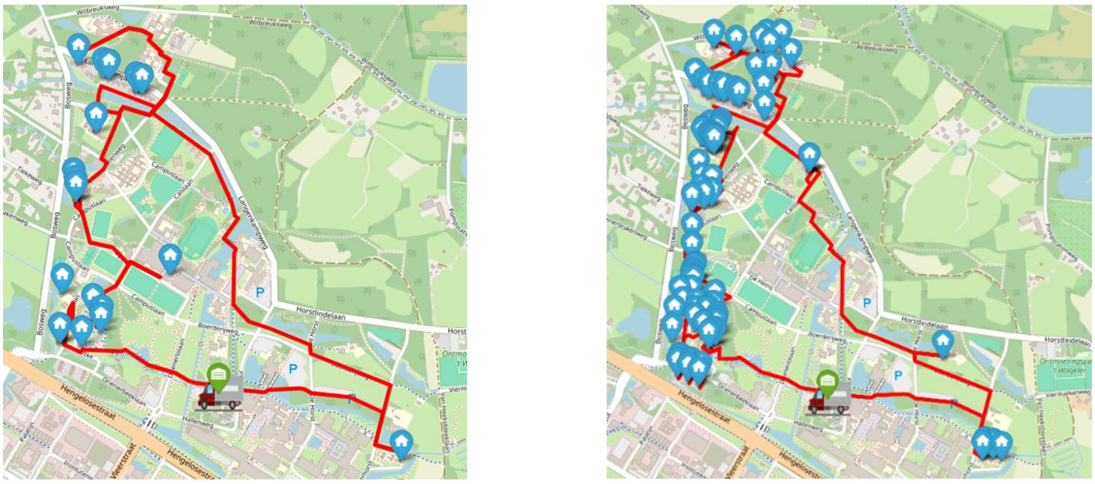

Figure 19: Delivery Route to 20 random customers (left) vs Delivery route to 100 random customers (right)

Some interesting results from the two routes are given in Table 13: Results 20 vs 100 customers.

Table 13: Results 20 vs 100 customers

	Distance	Expected Duration	Actual Duration
20 customers	7.46 km	1:47:22	1:42:51
100 customers	9.56 km	7:13:40	7:09:11

With an unloading time of 2 minutes, and an average pickup time of 2 minutes, it is expected that the electric van will take around 4 minutes for each customer. Assuming on average, the time it takes a van to drive from one customer to another is in the case of 20 customers 60 seconds, and in case of 100 customers 30 seconds, we can estimate the duration of the trips:

- For the 20-customer case, it is expected to take: 5 minutes x 20 + 5 minutes loading at depot = 105 minutes (1:45:00)
- For the 100-customer case, it is expected to take: 4,5 minutes x 100 + 5 minutes loading at depot = 450 minutes (7:30:00)

This roughly corresponds to the real time it takes the delivery van to deliver every package. The route for 100 customers is 2000 meters longer, which is approximately 25 meters per extra customer (2000/80). Looking at the picture, this distance between customer is very likely.

Now, it might seem like the delivery routes of the van are always in the same direction, starting in the bottom left corner and continuing upwards until moving down to the bottom right. Figure 20 shows an example of when the closest customer is not at the bottom left, but north of the depot. Which indicates that the Nearest Neighbor Algorithm does indeed choose the nearest neighbors.

It also highlights the limitation of the algorithm: the delivery route is not always optimal. In this case, a shorter total route could have been achieved if the van had started deliveries at the bottom-left corner and visited location 37 last.

Figure 20: Different Delivery Route example

Concluding, from what we have noticed, is that information that is available from OpenStreetMap is valid. This would indicate that the activities on the delivery network are also correct. Via the Manual Mode, it is also possible to follow the deliveries in real time (with a specific actions table which show the data under the model such as expected arrival time vs actual arrival time). From the black-box validation, we conclude that our model performs as expected and we can only say with great confidence that we do not know how the model can be incorrect (except that there are many assumptions such as average speed which do not give an entirely accurate picture of what would happen in real life).

4.8 Conclusion

Chapter 4 applied the autonomous delivery design framework to the specific case of the University of Twente campus. The campus was chosen for its semi-controlled environment, diverse infrastructure, and because it can be compared to a small village. The potential for real-life testing is also much higher on a university campus compared to other environments.

The main research question of this chapter was:
How can the autonomous last-mile delivery system be configured for the University of Twente campus using the developed framework and simulation model?

The decisions of the "menu of choices" were described and explained in this chapter. We have explained how the area of the University of Twente is generated. We have shown what the total set of delivery locations will be and why they were chosen. The daily demands were calculated using a test case in London, since we had almost no data on the demand at the campus. The "garage" of the University of Twente will be used as the depot location, because it is centrally located, well connected to the road network and this location is already being used for deliveries to university buildings. Solomon's nearest neighbor was used for the constructive heuristic which formed the basis of our routing strategy. We then covered the relevant KPIs to be used for performance evaluation in the simulation phase.

For this simulation model, we covered the inputs of the simulation model based on the design decisions. We also validated the model. Using the simulation model created for this research, the autonomous last-mile delivery system configured for the University of Twente campus can be evaluated. In the next chapter, we will discuss and present the experiments.

5 Experiments

This chapter describes the experimental setup used to evaluate the performance of the proposed autonomous last-mile delivery system. Building on the assumptions, inputs and design decisions discussed in the previous chapters, we present the simulation configuration, define the experimental factors, and detail the structure of the experimental design. The goal is to systematically assess how different configurations impact delivery performance on the campus. We aim to answer the following research question:

RQ. How does the autonomous delivery system perform under different design configurations, and which setup offers the best operational performance?

5.1 Experimental Settings

The simulation models deliveries for a single day, meaning one simulation run represents one full day of parcel deliveries.

5.1.1 Experimental factors

The experiments focus on varying the following input parameters:

Table 14: Experimental Factors

Exp. Factor	Description	Factor settings
Fleet Composition	Homogeneous vs Heterogeneous	UAVs only, S-ADVs only, R-ADVs only or combinations
Fleet size	Number of available vehicles per type	1-20 vehicles
Daily Demand	Total number of orders per day	Low (43), Average (106) and High (193)
Time Windows	Whether deliveries must occur within time windows (wide or narrow)	Without Time Window, with narrow time windows or with wide time windows
Opening Time	Night/Evening delivery or not	09:00-17:00 or 09:00-21:00

5.1.2 Ranges and Combinations

The experiments will be conducted in four different phases:

Phase 1

The goal of phase 1 is to determine the number of ADVs in a homogeneous fleet necessary to meet the demand (low, average, high) of a regular day without TimeWindows (09:00-17:00).

Phase 2

Once we determined the homogeneous fleet composition for a regular day without TimeWindows, we are interested in the heterogeneous fleet composition. Using the homogeneous number of R-ADVs, S-ADVs and UAVs as an upper bound, we experiment with different combinations of ADVs per demand type.

Phase 3

Once we know the homogeneous and heterogeneous fleet composition for a regular day without TimeWindows, we are interested in the impact of night delivery. So, we start experimenting on days 09:00 to 21:00. Using the homogeneous number of ADVs as an upper bound, we try to find the new homogeneous number of ADVs for night delivery. Then, we use this number as upper bound for the heterogeneous fleet.

Phase 4

Finally, we experiment with time window delivery (09:00-13:00 & 13:00-17:00) on a regular 09:00-17:00 day, and experiment to find the minimum number of ADVs we need to deliver to each customer without being late.

5.1.3 Warmup & Replications

The Key Performance Indicators described are the output of the simulation. Our simulation focuses on a single day delivery, independent of the preceding or following day. The day starts at the opening time of the depot after which the parcels can directly be delivered to customers and ends at the closing time of the depot. Since each day is independent and random, certain days might give us very extreme results. Which is why we need to run multiple days and take the average of those days. In this section we determine the minimum number of days needed for a single experiment.

We are dealing with a terminating simulation because of our closing time (natural event specifies the end of a simulation run), and the system has a transient behaviour (performance depends on initial conditions, in our case customer demand). This means that we do not need to use a warmup period.

To account for stochastic variation in simulation outcomes, each experiment is replicated multiple times. The number of replications is determined based on the precision of the confidence interval around the KPIs. Specifically, the half-width of the confidence interval, relative to the mean, is compared to a threshold γ'. The following criterion should be used:

$$\frac{t_{n-1,1-\alpha/2} * \sqrt{S^2/n}}{\bar{X}} < \gamma'$$

A pilot study with for example 10 replications is normally used to estimate the sample mean $\bar{X}$ and variance S^2, after which the number of replications *n* is increase until the criterion is satisfied. A frequently used threshold is a 5% error margin, which ensures that the 95% confidence interval half width relative to the mean is below 5%.

We chose to apply a fixed number of 10 replications per experiment. In each phase, we only used simulation configurations that successfully delivered all packages (on time) to be considered valid. Any replication in which one or more packages remained undelivered were not considered sufficient. For the other KPIs, no relative error is calculated, which represents a limitation of the study.

5.2 Experimental Results

In this section, we will go over the experimental results. The first experimental phase focuses on gaining an understanding of the number of vehicles necessary to deliver parcels on a low, average and high demand day.

5.2.1 Phase 1

The main goal of phase 1 is to gain a rough understanding of the number of vehicles we need to cover the demand on a regular day (09:00-17:00), without considering time windows constraints. This phase gives us the basis we need for further experimentation. Due to the extensive number of experiments conducted, most detailed results and tables are provided in Appendix F) Detailed Experiments Section. In phase 1, we focus exclusively on the key performance indicator *Successful Deliveries*, aiming to identify the minimum number of vehicles needed to fulfill all deliveries. Phase 1 is structured as follows: we first determine the number of required R-ADVs to cope with low, average and high demand, followed by the same analysis for S-ADVs and UAVs.

We start this phase by first analyzing the scenario Average demand. An initial experiment was conducted using a fleet size of 5, 10, 15, and 20 R-ADVs Vehicles. The results, shown in Table 15, reveal that the key performance indicators (KPIs) were nearly identical across all configurations. All setups successfully completed all deliveries with no unplanned stops and minimal variation in delivery time or energy consumption. Notably, the configuration with only 5 R-ADVs already achieved full delivery coverage, with a utilization rate of 26.64%.

Table 15: Initial experiment homogeneous fleet R-ADVs (average demand)

Exp	# of R-ADVs	Total Distance	Successful Deliveries	Trips	Util. (%)	Avg Delivery Time
1	5	19.11	106.0	5.0	26.64	00:06:01
2	10	19.13	106.0	5.0	13.46	00:06:05
3	15	18.86	106.0	5.0	8.72	00:05:55
4	20	19.42	106.0	5.0	6.51	00:05:53

Since the configuration with 5 R-ADVs showed only ~26% utilization, we investigated whether even fewer street Vehicles could achieve full delivery capacity. A second experiment tested configurations with 1 to 5 Vehicles. The results below (see Table 16) demonstrate that two R-ADVs were sufficient to handle all regular-day deliveries.

Table 16: Number of R-ADVs required for average demand

Exp	# of R-ADVs	Total Distance	Successful Deliveries	Trips	Util. (%)	Avg Delivery Time
1	1	14.86	81.33	4.0	98.74	00:05:50
2	2	20.4	106.0	5.0	66.96	00:06:03
3	3	20.06	106.0	5.0	44.6	00:06:02
4	4	20.61	106.0	5.0	33.92	00:06:09
5	5	19.84	106.0	5.0	26.52	00:05:59

This indicates that two R-ADVs can comfortably meet daily delivery demands, and a single robot has a maximum delivery capacity of around 80 parcels (regular day, no time windows). Next, the experiments for low demand (Table 17) and high demand (Table 18) are shown in tables below.

Table 17: Low demand R-ADVs (regular day, no time windows)

Exp	# of R-ADVs	Total Distance	Successful Deliveries	Trips	Util. (%)	Avg Delivery Time
1	1	11.66	43	2	53.19	00:05:56
2	2	11.8	43	2	28.42	00:06:23

A single R-ADVs was able to handle all 43 deliveries comfortably, with a utilization of 53.19%. Which, interestingly, does match what we saw earlier, that the street robot is able to deliver around 80 parcels a day (43/80 is around 53%)

Table 18: High demand R-ADVs (regular day, no time windows)

Exp	# of R-ADVs	Total Distance	Successful Deliveries	Trips	Util. (%)	Avg Delivery Time
1	2	25.42	167	8	101.83	00:05:52
2	3	26.7	192	8	77.59	00:05:50
3	4	28.1	192	8	58.53	00:05:52

With 3 R-ADVs, all 192 deliveries were completed. Again, notice the fact that this is conform to the 80 packages a day.

Conclusion Phase 1

The same type of experiments was conducted with the S-ADVs and the UAVs. All these experiments and their outcomes can be found in Appendix F) Detailed Experiments Section. A summary of the key findings is presented below in Table 19.

Table 19: Conclusion Phase 1

	Number of Vehicles necessary to fulfill entire demand		
Vehicle	**Low (43)**	**Average (106)**	**High (192)**
S-ADVs	5	11	20
R-ADVs	1	2	3
UAVs	2	4	6

From these results, we observe a linear relationship between the number of vehicles and the demand level. Based on our own observations under the conditions of a regular working day (09:00-17:00, no time windows), we can formulate the following equation:

$$Demand_{Regular\ Day} \approx 80 * \alpha + 10 * \beta + 35 * \gamma$$

Where:

- $\alpha = Number\ of\ R - ADVs$
- $\beta = Number\ of\ S - ADVs$
- $\gamma = Number\ of\ UAVs$

This equation will be further tested and refined in Phase 2, where we will explore the use of heterogeneous fleets. It must be noted that this equation is a rough estimation, since the total number of customers a vehicle can serve heavily depends on the location of the customers (a lot of generated customers near the hub allows the vehicle to deliver faster and thus more parcels).

5.2.2 Phase 2

Now, we are interested in the heterogeneous fleet scenario, where a different combination of vehicles is allowed to deliver parcels. From this point on, we begin using the KPIs discussed earlier, since a lot of combinations will satisfy the demand of the parcels. We do not consider time windows in this phase, so social KPIs will not be measured. Again, more detailed information on the experiments is given in the Appendix. To estimate sufficient heterogeneous combinations, we will apply two guiding principles:

- Avoid Redundant Capacity: Combinations in which a single vehicle type already meets the entire demand on its own are excluded. This helps focus on true mixed-fleet strategies.
- Use of estimation formula: The equation derived in Phase 1 is used as a reference to guide the search for effective fleet configurations:
$$Demand_{Regular\ Day} \approx 80 * \alpha + 10 * \beta + 35 * \gamma$$
- In general, we use the following table for creating our heterogeneous combinations:

S-ADVs	R-ADVs	UAV
Only S-ADVs	-	-
-	Only R-ADVs	-
-	-	Only UAVs
High S-ADVs	Low R-ADVs	-
Low S-ADVs	High R-ADVS	-
-	Low R-ADVs	High UAVs
-	High R-ADVs	Low UAVs
High S-ADVs	-	Low UAVs
Low S-ADVs	-	High UAVs

Low Demand Performance (Homogeneous & Heterogeneous, Regular Day, no TW)

Table 20: Heterogeneous Fleet Low Demand (Regular Day, No TimeWindows)

Combo ID	Demand	S-ADV	R-ADV	UAV	Reps	Del	Tot. Energy	Util.	Distance	Del. time
1	Low (43)	5	0	0	10	43	2.98	83.45	85.26	46:34
2	Low (43)	0	1	0	10	43	1.06	54.92	11.74	06:07
3	Low (43)	0	0	2	10	43	2.66	56.1	63.25	12:32
4	Low (43)	2	0	1	10	43	2.84	78.53	73.01	26:21
5	Low (43)	1	0	1	10	42.8	2.77	97.7	69.34	22:01

In the low demand scenario, the heterogeneous combinations are limited by the small number of vehicles available, especially R-ADVs (only 1) and UAV (only 2). As such, there is only one useful heterogeneous combination (Combo 4) involving both S-ADVs and UAVs. Combination 5 did not always succeed in delivering all the 43 parcels (42.8 deliveries, meaning 1 in 5 scenarios this combination can only deliver 42 parcels). This configuration matches the demand capacity closely (2*10 + 1*35 ≈ 55, 43/55 ≈ 78%) and performs relatively well across the KPIs.

Average Demand Performance (Homogeneous & Heterogeneous, Regular Day, no TW)

The results for heterogeneous fleet with average demand are shown in Table 21:

Table 21: Heterogeneous Fleet Average Demand (Regular Day, no TimeWindows)

Combo ID	Demand	S-ADV	R-ADV	UAV	Del.	Tot. Energy	Util.	Distance	Del. time
1	Average (106)	11	0	0	106	7.54	95.37	215.37	47:30
2	Average (106)	0	2	0	106	1.86	67.14	20.66	06:04
3	Average (106)	0	0	4	106	6.87	75.82	163.74	13:46
4	Average (106)	3	1	0	106	3.53	96.44	75.12	17:31
5	Average (106)	0	1	1	106	3.3	91.22	60.99	08:18
6	Average (106)	8	0	1	106	6.99	89.89	190.96	36:44
7	Average (106)	1	0	3	106	6.86	97.71	166.83	17:44

With the average demand, we observe that a lot more options seem to be viable due to the increase in available ADVs. The utilization rates indicate that these options are the most 'efficient' ones. Since adding an S-ADV or UAV to a certain combination would mean that they essentially have overcapacity.

Example Calculation using demand estimation formula

Let's verify combo 6 using the demand formula:

$$Demand_{Regular\ Day} \approx 80 * 1 + 10 * 0 + 35 * 1 \approx 120$$

Thus, Combo 6 has a total capacity of 120, giving it +-8% overcapacity (106/120) compared to actual demand.

High Demand Performance (Homogeneous & Heterogeneous, Regular day, no TW)

With the high demand, we observe that there are a lot of combinations possible to meet the demand on a single day. Again, these utilization rates indicate that these options are the most 'efficient' ones. What does become very clear at this point is the increase in vehicles in perspective of the entire day of demand. For example, the number of S-ADVs increases a lot based on demand while the number of R-ADVs and UAVs remain quite stable. This makes sense, based on the average delivery time per package, which is very high for S-ADVs (because of their slow speed).

Table 22: Heterogeneous Fleet High Demand (Regular Day, no TimeWindows)

Combo ID	Demand	S-ADV	R-ADV	UAV	Del.	Tot. Energy	Util.	Distance	Del. time
1	High (192)	20	0	0	192	13.7	95.41	391.47	47:49
2	High (192)	0	3	0	192	2.56	79.04	28.39	05:56
3	High (192)	0	0	7	192	12.31	79.27	293.17	13:53
4	High (192)	12	1	0	192	9.23	93.45	242.49	30:25
5	High (192)	3	2	0	192	4.4	96.04	82.86	12:00
6	High (192)	0	1	4	192	8.74	87.68	194.95	10:57
7	High (192)	0	2	1	192	4.31	96.44	71.11	07:15
8	High (192)	17	0	1	192	13.31	93.48	371.34	42:10
9	High (192)	4	0	5	192	12.32	91.14	306.02	20:33
10	High (192)	10	0	3	192	12.4	92.91	327.1	30:17
11	High (192)	2	1	3	192	8.09	91.37	184.66	13:44
12	High (192)	4	1	2	192	8.3	93.26	197.85	18:42

Example Calculation using demand estimation formula

Let's verify combo 11 using the demand formula:

$$Demand_{Regular\ Day} \approx 80 * 1 + 10 * 2 + 35 * 3 \approx 205$$

Thus, Combo 11 has a total capacity of 205, giving it +-6.5% overcapacity (192/205) compared to actual demand.

Comparison with electric van

From these results, we observe that the average demand scenario represents the practical upper limit for operating a single electric van, given the utilization rate of 98.01%. This indicates that the vehicle is nearly fully occupied during the available delivery window (09:00-17:00). So, there is basically no buffer time for delays. The 0.2 late customers in the average scenario implies that on some days, the final customer gets their package delivered a little bit after 17:00. The late customers in the high scenario are explained because one delivery van is fully packed until 17:00, creating the same problem as the average demand scenario.

Table 23: Performance electric van regular days

Demand	# of vans	Del.	Tot. Energy	Utilization	Distance	Del. time	Late customers
Low	1	43	1.01	43.16	9.69	04:48	0
Average	1	106	1.31	98.01	12.62	04:25	0.2
High	2	192	1.7	85.97	16.36	04:17	0.7

Conclusion Phase 2

From Phase 2, we learn that the estimated demand formula developed is moderately effective for predicting demand in a heterogeneous fleet composition. We also observe that the number of S-ADVs does not scale proportionally well with demand when compared to R-ADVs and UAVs. This may be improved by including night delivery in the next phase, which could provide the S-ADVs with longer delivery windows. Additionally, when combining different types of ADVs in a fleet, we see changes in the average delivery times. Specifically, R-ADVs and UAVs demonstrate lower average delivery times, suggesting they might be more suitable for deliveries with time windows.

5.2.3 Phase 3

Now, we are interested in the performance of the system with night delivery (09:00-21:00) instead of normal hours (09:00-17:00). Just like the previous phase, we start with the low demand case, continue with the average demand case and finally the high demand case. Since we have a longer delivery time, the expected performance is that we need fewer vehicles for each demand scenario.

Low Demand Performance (Homogeneous, Night Delivery, No TW)

The night delivery scenario (See Table 24) creates the possibility to deliver every parcel by either one UAV, one R-ADV or four S-ADVs. These values also show that a heterogeneous fleet is not possible/necessary for the low demand scenario (combinations will always lead to overcapacity).

Table 24: Performance low Demand Heterogeneous & Homogeneous fleet with night delivery

Combo ID	Demand	S-ADV	R-ADV	UAV	Del	Late cust	Tot. Energy	Util.	Distance	Del. time
1	Low (43)	4	0	0	43	0	3.09	72.09	88.18	48:16
2	Low (43)	0	1	0	43	0	1.08	37.47	11.95	06:18
3	Low (43)	0	0	1	43	0	2.76	90.44	65.75	15:20

Average Demand Performance (Homogeneous and heterogeneous, Night Delivery, No TW)

The night delivery scenario with average demand (See Table 25) creates opportunities for some heterogeneous fleets. In this night delivery scenario, one R-ADV is still able to deliver all the packages on its own. Because of this, there is no efficient heterogeneous fleet combination with R-ADVs. This leaves only combinations with S-ADVs and UAVs.

Table 25: Performance Average Demand Heterogeneous and Homogeneous fleet (Night Delivery, No TW)

Combo ID	Demand	S-ADV	R-ADV	UAV	Del	Late cust	Tot. Energy	Util.	Distance	Del. time
1	Avg (106)	8	0	0	106	0	7.54	87.5	215.34	47:41
2	Avg (106)	0	1	0	106	0	1.8	88.8	20	06:01
3	Avg (106)	0	0	3	106	0	6.81	72.5	162.21	14:46
4	Avg (106)	2	0	2	106	0	6.86	83.87	172.01	22:49
5	Avg (106)	5	0	1	106	0	7.08	84.76	190.37	34:35

High Demand Performance (Homogeneous and heterogeneous, Night Delivery, No TW)

In the high demand scenario, it is no longer possible for the single R-ADV to deliver each parcel. Two R-ADVs are now necessary to meet the demand of 192 parcels, with a utilization rate of 80%. The S-ADVs do not scale well, even in the night delivery scenario. Instead of the 20 S-ADVs we needed, we now still need 15 of them. The R-ADVs can now also be considered in the heterogeneous fleet combinations. Besides the combinations between two ADVs, this high demand scenario also has one heterogeneous fleet combination with all three ADVs.

Table 26: Performance High Demand Homogeneous and Heterogeneous Fleet (Night Delivery, No TW)

Combo ID	Demand	S-ADV	R-ADV	UAV	Del.	Tot. Energy	Util.	Distance	Del. time
1	High (192)	15	0	0	192	13.5	84.03	385.76	47:22
2	High (192)	0	2	0	192	2.45	77.61	27.27	05:49
3	High (192)	0	0	5	192	12.26	79.52	291.98	14:58
4	High (192)	7	1	0	192	7.65	82.35	191.84	24:48
5	High (192)	0	1	2	192	6.11	82.91	126.52	09:20
6	High (192)	8	0	2	192	12.6	87.44	335.96	32:49
7	High (192)	2	0	4	192	12.5	87.18	306.61	19:41
8	High (192)	2	1	1	192	5.65	88.33	125.49	13:17

Comparison with electric van

On regular days, the electric could just deliver the parcels to the customers. On some days, the last parcel was delivered a bit late (after 17:00). This changes in the night delivery scenario (Table 27). The performance of the electric van does not really change that much. Only that in the high demand scenario, the utilization is much lower. This is also expected because we have a longer delivery time window.

Table 27: Performance Electric Van Night Delivery

Demand	# of vans	Del.	Tot. Energy	Utilization	Distance	Del. time	Late customers
Low	1	43	1.02	41.83	9.78	04:48	0
Average	1	106	1.5	67.11	14.38	04:32	0
High	2	192	1.7	58.1	16.35	04:21	0

Conclusion Phase 3

From Phase 3, we can conclude that the following number of ADVs is necessary to meet the demand of parcels:

Table 28: Conclusion Phase 3 (Number of ADVs)

	Number of Vehicles necessary to fulfill entire demand		
Vehicle	**Low (43)**	**Average (106)**	**High (192)**
S-ADVs	4	8	15
R-ADVs	1	1	2
UAVs	1	2	5
Electric Van	1	1	2

From the experiments, we can still observe some linear relationships between the number of vehicles and the demand level. Where the R-ADV and UAV really improve in the potential number of parcel deliveries on a day-to-day basis, the S-ADV does not really show the improvement you would want in the night scenario.

Just like in Phase 1, we formulate another number of vehicles demand equation.

$$Demand_{Night\ Delivery} \approx 120 * \alpha + 12 * \beta + 50 * \gamma$$

Where:

- $\alpha = Number\ of\ R - ADVs$
- $\beta = Number\ of\ S - ADVs$
- $\gamma = Number\ of\ UAVs$

Example Calculation using demand estimation formula

Let's verify *Average demand (2 S-ADVS, 0 R-ADV, 2 UAV)* using the demand formula:

$$Demand_{Night\ Delivery} \approx 120 * 0 + 14 * 2 + 50 * 2 \approx 128$$

This combination had a utilization of 83.5% according to the experiments, 106/124 results in 82.8%.

Let's verify *High demand (2 S-ADVS, 1 R-ADV, 1 UAV)* using the demand formula:

$$Demand_{Night\ Delivery} \approx 120 * 1 + 14 * 2 + 50 * 1 \approx 198$$

This combination had a utilization of 88.3% according to the experiments, 192/198 results in 96.9%. So, the equation works quite well in the average case and underestimates the performance in the high demand case.

5.2.4 Phase 4

For the final phase, we will check the compatibility of the autonomous delivery system with customer specific time windows. Which means that the customers can choose their own preferred time window. For example, on a 09:00-17:00 day, customers can select morning delivery 09:00-13:00 or afternoon delivery 13:00-17:00. One of the most interesting KPIs right now are the number of late customers (and the total minutes being late). Also, we are interested in the KPI *Average time after the start of TimeWindow* which, as the name suggests, is the average amount of time customers must wait after the start of the TimeWindow. A lower *Average Time After Start of TimeWindow* is preferred, since customers won't have to wait as long. Before we start experimenting, we will have to take another look at our Solomons nearest neighbor's cost function. This cost function, as we have used (and talked about) before, has the following weights:

$$C_{i,j} = w_1(0.025) * d_{i,j} + w_2(0.025) * T_{i,j} + (0.95)w_3 * \max(0, v_{i,j}) + PenaltyCosts * \max(0, -v_{i,j})$$

The reason that we chose for this distribution of weights now becomes useful. Earlier, we only had to deal with distance and time. Where now, we must deal with TimeWindows. In this scenario, we prioritize the slack time of the time window of the customer, thus giving it a higher weight compared to distance/travel time.

Low Demand Performance (Homogeneous and heterogeneous, Regular Day, TW)

From this point on, the important information to save on the tables becomes a bit much, so we use some abbreviations in the tables (LM = Total Late Minutes, LC = Average Late Customers, ATAT = Average Time After TimeWindows). In the low demand scenario (Table 29: Performance Low Demand Homogeneous (Regular Day, TW)Table 29), it takes 8 S-ADVs and 4 Drones to deliver to each customer on time in their TimeWindow. As one can see, the R-ADV scenario does not get to the point of achieving zero late customers. This has to do with our Solomons Nearest Neighbor algorithm which is not really suitable for TimeWindow deliveries with high-capacity vehicles. Further explanation in conclusion.

Table 29: Performance Low Demand Homogeneous (Regular Day, TW)

ID	S-ADV	R-ADV	UAV	Del.	Dist.	LM	LC	Energy	Util.	ATAT	Del. time
1	10	0	0	43	109.37	0	0	3.83	73.42	69.15	01:21:57
2	0	2	0	43	15.7	186.3	3	1.41	64.81	79.49	14:30
3	0	0	4	43	65.05	0	0	2.73	62.68	37.36	28:03
4	7	1	0	43	96.26	0	0	3.51	69.26	55.76	01:01:59
5	0	1	3	43	34.13	138.32	1.1	2.05	50.19	58	22:24
6	7	0	1	43	93.04	0	0	3.42	69.86	51.12	01:02:22
7	3	0	3	43	75.74	0	0	2.99	64.17	42.22	42:58

Average Demand Performance (Homogeneous and heterogeneous, Regular Day, TW)

In the average demand scenario, the number of S-ADVs necessary to deliver the parcels to get zero late customers is 18, and the number of UAVs is 9. The R-ADVs do not improve after 5, since they drive a total of 5 trips which is enough to deliver each parcel (5*24 = 120, 120>106).

Table 30: Performance Average Demand Homogeneous (Regular Day, TW)

ID	D	S-ADV	R-ADV	UAV	Dist.	LM	LC	Energy	Util.	ATAT	Del. time
1	106	18	0	0	218.82	0	0	7.66	77.48	71.25	01:03:11
2	106	0	5	0	26.84	114.43	2.5	2.42	37.91	72.29	12:53
3	106	0	0	9	163.14	0	0	6.85	63.79	41.21	26:02
4	106	15	1	0	198.37	0	0	7.52	72.9	77.43	52:56
7	106	6	2	0	98.23	278.88	2.6	4.42	71.62	90.26	25:59
8	106	0	2	3	78.92	0	0	4.22	61.61	70.02	13:56
10	106	0	1	8	99.46	540.57	5.2	4.88	61.97	74.24	25:26
11	106	15	0	3	224.85	0	0	8.29	71.55	59.45	57:20
12	106	10	0	5	198.74	0	0	7.72	68.78	49.79	44:30

High Demand Performance (Homogeneous and heterogeneous, Regular Day, TW)

In the high demand scenario, just like in the low and average scenario, the R-ADVs do not improve after 8, since they drive a total of 8 trips which is enough to deliver each parcel (8*24 = 192). We observe that the 30 S-ADVS are sufficient for time window delivery without late customers, as well as 11 drones.

Table 31: Performance High Demand Homogeneous (Regular day, TW)

ID	D	S-ADV	R-ADV	UAV	Dist.	LM	LC	Energy	Util.	ATAT	Del. time
1	192	30	0	0	382.94	0	0	13.44	53.57	74.52	01:00:11
2	192	0	8	0	35.79	41.49	0	3.22	36.88	66.18	11:05
3	192	0	0	11	288.68	0	0	12.12	48.92	55.63	20:11
4	192	15	3	0	224.68	2.46	0.2	9.32	70.88	80.42	31:53
5	192	0	3	8	83.37	26.9	1.4	5.07	64.4	80.78	17:42
7	192	26	0	3	430.44	79.61	1.33	15.64	78.66	75.4	57:07
8	192	8	0	8	328.83	0	0	13.05	77.25	64.75	30:53

Conclusion Phase 4
From Phase 4, we can conclude that the following number of ADVs is necessary to meet the demand for parcels:

Table 32: Conclusion Phase 4 (Number of ADVs)

	Number of Vehicles necessary to achieve 0 late customers		
Vehicle	**Low (43)**	**Average (106)**	**High (192)**
S-ADVs	10	18	30
R-ADVs	**2**	**5**	**8**
UAVs	4	9	11

The main conclusion drawn from phase 4 is the inability of the Solomons Nearest Neighbor to deal with high-capacity vehicles, such as R-ADVS, in case of strict TimeWindow constraints. The problem lies in how the algorithm assigns customers to trips.

For example, in the low-demand scenario (43 parcels), a single R-ADV could technically deliver all parcels in two trips (2x24 = 48 capacity). However, because the algorithm now prioritizes meeting the strict time windows rather than minimizing distance, each trip takes a bit longer. As a result, one R-ADV is no longer enough, and a second R-ADV is needed to meet the demand and time windows.

What happens is this:
- The first R-ADV is filled with the “best” customers according to the cost function (fitting within the time windows and travel time)
- The second R-ADV is then assigned to the remaining customers, who are often more spread out.

In some cases, if a very large share of morning deliveries ends up in the second R-ADV (for example the case of 26 total morning deliveries vs 17 afternoon deliveries), it becomes nearly impossible for that second vehicle to meet all the morning time windows (especially when they are further apart).

Not to forget that the service times per customer are not fixed (we use a distribution), so even small delays can cause the last few deliveries on a trip to be late. This issue affects the Late Customers KPI, which makes it almost impossible to get a good idea of the effectiveness of R-ADVs at strict time windows. The sequential filling of the vehicles using Solomons nearest neighbor prevents effective use of R-ADVs with strict TimeWindows. This is a limitation of our research.

The UAVs and S-ADVs performed as expected and increased very much in necessary amounts. It remains the question, however, whether this increase in amount is worth the strict time windows.

5.3 Conclusion

In this chapter, we experimented with different configurations for the campus case. At the start of the chapter, we aimed to answer the following research question:

How does the autonomous delivery system perform across different configurations, and how does it compare to the current system?

To answer this question, we answer it in different parts.

ADV parcel delivery capacities

First of all, we obtained a rough estimation of the number of parcels an ADV can deliver in a day (regular day 09:00-17:00 and night delivery 09:00-21:00).

$$Demand_{Regular\ Day} \approx 80 * \alpha + 10 * \beta + 35 * \gamma$$

$$Demand_{Night\ Delivery} \approx 120 * \alpha + 12 * \beta + 50 * \gamma$$

This gives us an estimated guess on how many parcels an ADV can deliver in a day, which is useful information in value comparison. Deliveries with strict TimeWindows did not really gave us an accurate demand estimation formula, since it also really depends on the distribution of Morning/Afternoon deliveries.

Vehicle-Type Performance analysis

Across all the experimental phases, the ADV performance remained somewhat the same. There was not a vehicle that performed very differently in night delivery scenarios or smaller TimeWindow scenario. It must be noted that the sequential filling of vehicles with Solomons Nearest Neighbor algorithm is not suitable for higher capacity ADVs/vehicles in cases of strict time windows.

R-ADVs (Road Autonomous Delivery Vehicles) demonstrated the best all-round performance in terms of energy efficiency, delivery time, and scalability. With a capacity of 24 and the ability to deliver around 80 packages a day, they are ideal for campus-wide deployment or at least as the backbone of a heterogeneous fleet.

UAVs (Unmanned Aerial Vehicles) showed surprisingly good results for a single capacity vehicle. Due to their speed and ability to bypass the traffic by skipping traditional road infrastructure, the delivery times are very short compared to the S-ADVs. This is also apparent in the total kilometers they must fly in a single delivery day. It is always significantly shorter than the S-ADVs, but not even close to the low kilometers the R-ADVs drive during a day. With a capacity of 1 and the ability to deliver around 35 packages a day makes them scale very well. However, it makes them perhaps more suitable for supplementing deliveries in a heterogeneous fleet, but they are also certainly suitable for a homogeneous fleet.

S-ADVs (Sidewalk Autonomous Delivery Vehicles) turned out to be the worst of the three. This might have been expected from the start, but the experiments showed that they really do not come close to the R-ADVs and UAVs in whatever area. With a capacity of 1 and the ability to deliver around 10 packages a day, they are certainly not suitable for homogeneous package delivery. So, they can only be suitable for supplementary deliveries. However, since the UAV beats the S-ADVs in each aspect, this might not be the first supplementary choice.

5.3.1 Economic, Social and Environmental Assessment

Assessing the ADVs across economic, social and environmental factors remains a challenge, especially given the limited data availability and the subjective nature of some criteria. This section still tries to assess the three ADVs based on available cost estimates, energy consumption, and potential social impacts.

Economic Assessment

Economic comparisons are difficult to make without having the prices of each vehicle. The vehicles used throughout this thesis-such as the Express Robot (R-ADV), Starship Robot (S-ADV), and Zipline Drone (UAV) do not share detailed cost information. In fact, all of the characteristics of these ADVs were already very difficult to find.

There is one source online (Condliffe, 2022) stating that Kristjan Korjus said in 2018 that the price of a S-ADV was around 5000 dollars. Although the price is probably affected by inflation, the increasing adoption and technology maturity of such systems typically drive costs down. Therefore, we assume 5000 dollars (roughly €4300) as a reasonable estimate for the unit price in this research. Since we do not have the official prices of the other ADVs, we propose a relative valuation based on their average delivery capacities per day:

- S-ADV: ≈ 10 deliveries/day
- R-ADV: ≈ 80 deliveries/day
- UAV: ≈ 35 deliveries/day

Assuming cost scales linearly with the delivery capacity:

- An R-ADV could be valued at 8x the S-ADV price, € 34,400
- A UAV could be valued at 3.5x the S-ADV price, € 15,000

This over simplified approach allows a rough economic ranking. A UAV is likely to cost significantly less than € 15,000, making it a more cost-efficient option. Similarly, while an R-ADV may indeed cost more than an S-ADV, it is unlikely to be eight times as expensive. Based on these estimations, the UAV appears to offer the best cost-efficiency, followed by the R-ADV, with the S-ADV being the least economically favorable. We should also consider the operating costs, which in this case are electricity costs. Assuming a price of 0.32 euros per kWh (Overstappen.nl, 2025), the electricity price per day for each type of vehicle is:

- Homogeneous S-ADV fleet (11): Total 7.54 kWh (€ 2,41)
- Homogeneous R-ADV fleet (2): Total 1.86 kWh (€ 0,60)
- Homogeneous UAV fleet (4): Total 6.87 kWh (€ 2,20)

Even though the S-ADVs are the most cost efficient per vehicle per operating day, in total they are the least energy efficient, followed closely by UAVs and the most energy efficient are R-ADVs.

Environmental Assessment

Environmental comparisons are also difficult to make. Because each option is electrical, there are no emissions per vehicle (assuming electricity comes from clean sources). From the cost assessment, we know that R-ADVs are the most energy efficient, followed by UAVs and finally the S-ADVs.

Social Assessment

The social assessment is arguably the most difficult to assess, mainly because it lacks the objective, simulation-based data available for economic and environmental performance. One of the most important social KPIs, the social perception of ADVs by campus students and residents, is pretty subjective and hard to quantify without survey-based research.

Nevertheless, we can again make some estimations based on vehicle activity. On an average day, the delivery system requires:

- 11 S-ADVs, covering a total of 215 kilometers
- 2 R-ADVs, covering a total of 20 kilometers
- 4 UAVs, covering a total of 163 kilometers

Given these numbers, students and residents are much more likely to encounter S-ADVs in daily campus life. With their slow speed and regular occurrence on the network, S-ADVs are more visible and potentially more disruptive. The high number of vehicles and the total distance traveled imply a constant, noticeable presence, which could negatively influence public perception. R-ADVs are much larger and potentially noisier (likely near the legal limit of 56 dB because of their participation on the road with other cars). However, they have much lower presence on the roads. This limited presence, both in vehicle numbers and distance driven, likely results in lower visibility and lower perceived disturbance.

UAVs are a completely different story, because they do not participate in traffic. They are, however, flying a total of 163 kilometers per day. They cover this distance through direct flight paths across campus, making them frequently visible overhead. They are also quit fast and maybe a bit loud (around 60 dB), which could raise some concerns about noise pollution and safety, especially in usually quiet places.

This analysis suggests that social acceptability is likely the highest for R-ADVs, moderate for UAVs and the lowest for S-ADVs.

5.3.1 Recommendations for the University of Twente Campus

Finally, we end with a recommendation to the University of Twente Campus. We recommend not to implement deliveries with time windows. Since it would be way too costly to achieve this compared to delivery without time windows.

If you want an autonomous delivery system with full day (09:00-17:00) delivery to customers at their home, we recommend using a heterogeneous fleet of **one R-ADV and two UAVs**. This heterogeneous fleet can deliver up to around 150 parcels a day between 09:00-17:00 (covering days with low demand, average demand and even slightly higher than average demand).

On a lower demand day, the single R-ADV or the two UAVs can deliver the parcels homogeneously (each with around 55% utilization). On an average demand day, a single R-ADV and UAV can deliver these packages with 91% utilization. On days of high demand, we recommend using an option of evening delivery, which allows this fleet to deliver 192 parcels with a utilization of 83%. If evening delivery is not an option, the heterogeneous fleet needs an additional **two drones** to cover the high demand days.

As we read before, findings, basing the decision on economic, environmental and social is quite difficult.

Comparison with current scenario (electric van)

From the experiments, we know that only one electric van (current situation) can deliver the parcels on low demand and average demand days. On high demand days, this number doubles. So, we can draw up some rough comparisons. On low demand and average demand days, the autonomous delivery system requires one worker (09:00-17:00) which monitors the autonomous vehicles and loads them with packages. The same is true for the worker on the electric van, which delivers the packages himself. This means that in personnel costs, this is the same. In higher demand days, the autonomous delivery system requires late night delivery, so the worker needs to work for four more hours (a total of 12 hours of work). In the case of electric van, however, there are now 2 vans required. This means that a total of 2 workers need to work for 8 hours, which totals 16 hours (4 hours more). This autonomous delivery system is therefore more efficient in terms of staff costs (on high demand days), which is usually a high cost in business.

Another advantage of the autonomous delivery system is that personnel are only required at a single location to monitor operations and load parcels. In contrast, the current delivery process typically requires two delivery workers, especially on high-demand days, to carry out the deliveries using vans. The autonomous system's ability to operate with just one staff member, as opposed to the necessity of two in the current setup, significantly increases the operational flexibility.

The argument could be made that we exclude the larger packages which are not considered because they do not fit in with the ADVs, which would indicate that in the autonomous delivery scenario, there should be someone delivering them as well. However, we compared the same type and number of packages for both situations, so this is also the case for the current scenario. In the current situation, the electric bus can only handle the demand (98%). So, this will no longer be possible in the situation with the large packages. This means that with average demand, an additional driver must also be used in the current scenario (which evens it out).

Chapter 6 Conclusion and Recommendations

The most important findings and recommendations are summarized in this chapter by systematically revisiting each stage of the research and answering the associated research questions. This structured approach highlights how the study addressed the central problem and provides a comprehensive overview of key findings.

Additionally, we outline the limitations of the study and explain how the research could have been improved. Finally, we give recommendations for future research regarding this topic.

6.1 Most important Findings and Recommendations

At the start of this research, we encountered the problem that the current last-mile delivery system was approaching a point that it would no longer be sustainable due to various challenges. Autonomous delivery presented a potential solution, but it was unclear what approach should be used to implement this. To address this problem, we formulated the following main research question:

How to design an autonomous B2C last mile delivery system, using the Campus of the University of Twente as a case study?

This question was answered using four different stages, which we will cover right now.

Stage 1. Literature review

The thesis started with a literature review with the goal of answering the following question: *What does existing literature reveal about the technologies, models and methods available for designing and evaluating autonomous last-mile delivery systems?*

The review identified five key challenges facing the current last-mile delivery system: Operational Challenges, Infrastructure challenges, delivery challenges, logistical challenges and environmental challenges. Autonomous Delivery Vehicles emerged as a promising innovative solution to these issues. Three types of ADVs were identified: Road Autonomous Vehicles (R-ADVs), Sidewalk Autonomous Vehicles (S-ADVs) and Unmanned Aerial Vehicles (UAVs), each with unique operational characteristics and constraints. The literature also detailed various possible autonomous delivery system architectures, including single-tier networks, two-tier networks and ADV-aided systems. Mathematical modeling approaches are widely used to model and optimize these delivery systems. Exact methods such as Integer Linear Programming provide can provide a fast solution for smaller VRPs. Once the size of the problems increases, non-exact methods become necessary to find a (near) optimal solution. Constructive heuristics, like the nearest neighbor approach, create an initial solution, which can then be optimized by local search heuristics like 2-ot, and 3-opt. Metaheuristics are a problem independent technique which can be applied to a broad range of problems. Finally, the literature revealed a wide range of KPIs to evaluate last-mile delivery systems. However, their applicability largely depends on the specific context and objectives of the evaluation.

Stage 2. Design Framework with Simulation Model

The second stage of the study focused on creating an autonomous last-mile delivery system and presenting the simulation model developed for evaluation. In this stage, the aim was to answer the following question: *How can an autonomous B2C last-mile delivery system be designed, using a set of configurable design choices, based on the characteristics of a specific environment?*

We have created a general framework structured as a "menu of choices", helping system designers make informed decisions about service area boundaries, demand patterns, depot placement, and operational logistics. These design choices are not only theoretically grounded but also serve as direct inputs to the simulation model used for evaluation.

Stage 3. Case study: Applying the Framework to the University of Twente
In this stage, the system design framework and simulation model was applied to the specific use case of the University of Twente campus. We aimed to answer the following question: *How can the autonomous last-mile delivery system be configured for the University of Twente campus using the developed framework and simulation model?*

Based on campus characteristics and available data, a feasible autonomous delivery system was configured. This included decision-making in:

- Service area: University of Twente campus
- Demand estimations: Parcel demand was estimated using partial delivery data from PostNL and extrapolated with seasonality patterns derived from an external dataset
- Depot configuration: The garage on campus will be used as a depot, meaning that we use a two-tier delivery system
- Operational logistics: Three ADV types were used for the characteristics (Starship Robot for S-ADV, Macrostep for R-ADV and Zipline for UAV), fleet configuration are used as simulation input, and we use a constructive heuristic (Solomons Nearest Neighbor) for our routing strategy

This information was entered into the simulation model created for this research and used for evaluation.

Stage 4. Experiments and Evaluation
This stage focused on conducting simulation experiments using the previously selected configurations. The goal was to answer the following question: *How does the autonomous delivery system perform under different design configurations, and which setup offers the best operational performance?*

A series of experiments were conducted to test the system under different configurations, including variations in fleet composition (homogeneous vs heterogeneous), customer demand levels (43, 106, 192), and delivery constraints (09:00-17:00 delivery, 09:00-21:00 and strict time windows of 09:00-13:00 & 13:00-17:00).

For campus deliveries, we recommend not to implement deliveries with strict time windows. Since it would be way too costly to achieve this compared to delivery without time windows. For an autonomous delivery system with full day (09:00-17:00) delivery to customers at their home, we recommend using a heterogeneous fleet of **one R-ADV and two UAVs**. This heterogeneous fleet can deliver up to around 150 parcels a day between 09:00-17:00 (covering days with low demand, average demand and even slightly higher than average demand).

On a lower demand day, the single R-ADV or the two UAVs can deliver the parcels homogeneously (each with around 55% utilization). On an average demand day, a single R-ADV and UAV can deliver these packages with 91% utilization. On days of high demand, we recommend using an option of evening delivery, which allows this fleet to deliver 192 parcels with a utilization of 83%. If evening delivery is not an option, the heterogeneous fleet needs an additional **two drones** to cover the high demand days.

Concluding
So, to design an autonomous B2C last mile delivery system, using the University of Twente campus as a case study, it is essential to follow a structured approach that combines literature insights, a configurable design framework, and simulation-based evaluation.

The recommended design, a two-tier delivery system using a heterogeneous fleet of one R-ADV and two UAVs (operating within a full-day delivery window), proved to be the most efficient across varying demand levels. This configuration balances operational efficiency, scalability, and practicability, while also avoiding the high costs associated with strict time windows or overcapacity.

This research does not only provide a concrete system design for the campus but also offers a general framework and methodology that can guide system designers for similar implementations in other environments. The framework is highly adaptable and flexible, making sure that it can be tailored to the specific needs of many different environments.

In the following sections, we reflect on the limitations of this study and propose ideas for future research.

6.2 Limitations

Although the findings and recommendations of this research are very interesting, the research also had some limitations. In this section we acknowledge the limitations of this research and give suggestions on how this research could have been improved.

Model/Framework limitations
First of all, the framework only serves as a general guide in designing autonomous delivery systems. While it provides a structured and adaptable approach, it does not claim to capture every possible design factor or operational nuance.

As we have discussed during this research, the Solomons Nearest Neighbor is not really effective for strict window customer allocation because of its sequential filling algorithm.

Creating a simulation model which represents the real word completely is very challenging, which is why a lot of simulation models are based on certain assumptions. Assumptions such as: Independent days, no traffic interference and average vehicle speed (see **Error! Reference source not found.**) can potentially be removed by more extensive programming. This would make the model more accurate and thus also give more accurate results.

Although the literature review covered various economic, environmental, and social Key Performance Indicators (KPIs), this research found a lack of quantifiable KPIs that accurately assess and compare autonomous delivery systems on these dimensions using the simulation model. For example, environmental impact can be evaluated on CO_2 emissions and energy usage, but this is minimal for all of the ADV options studied. Similarly, social impacts such as effects on employments and user acceptance were not really quantitatively measured.

Campus specific limitations
Another limitation of the research is the almost non-existent information on package/customer distribution on the campus. Due to privacy reasons this information cannot be made public, but it does make for a limitation of the research. More accurate information about the housing of students (numbers per house) and the package distribution (distribution of parcels per house, seasonality) will improve the model significantly and give more accurate results. The same principle holds for more accurate information on the ADV characteristics.

We chose to apply a fixed number of 10 replications per experiment. In each phase, we only used simulation configurations that successfully delivered all packages (on time) to be considered valid. Any replication in which one or more packages remained undelivered were not considered sufficient. For the other KPIs, no relative error is calculated, which represents a limitation of the study.

The implementation of autonomous delivery systems is subject to legal and regulatory frameworks. This research assumed feasibility of UAV and ground ADV deployment on campus, while there is a high probability that UAV delivery is not possible with the Twente Airport near the campus.

6.3 Future Research

In this section, we outline opportunities for future research building upon the foundations of this study. We start by discussing the possibilities for the specific case of the University of Twente campus, followed by the potential of this research in a broader sense.

A natural first step of this research would be to deal with the limitations of the model and perhaps create a more accurate simulation model.

For this scenario of the University of Twente case, an extension of parcel delivery to the university buildings could be considered. So, also consider the deliveries to the university buildings. As well as experimenting with what else is possible with autonomous transport. For example, one very interesting option is the Generalized Traveling Salesman Problem, in which each node has a set of different TimeWindows. For example, morning delivery at home and afternoon delivery at work. With the help of an autonomous system that can dynamically adapt, this could possibly work well in the future.

This research was conducted with not only the University of Twente campus in mind, but also the broader picture of autonomous last-mile delivery. So, we encourage future research by using the framework that we have created and applying it to different use cases. For similar use cases to test whether this framework is useful for cases like a campus, but also completely different cases to (stress) test the possibilities of the framework.

Every component of the framework can also be more extensively researched. For instance, future research could investigate alternative demand forecasting methods or evaluate routing algorithms that better handle complex constraints such as dynamic traffic or customer preferences.

Finally, this research could be used as the foundation of a digital twin model, which means that the routing and findings of a day can be used as input for actual deliveries.

Bibliography

About Open Trip model. (n.d.). OpenTripModel. https://opentripmodel.org/

Acm. (2023, July 3). *Post- en pakketmonitor ACM: 43 miljoen minder pakketten in 2022*. ACM.nl. https://www.acm.nl/nl/publicaties/post-en-pakketmonitor-acm-43-miljoen-minder-pakketten-2022

Akeb, H., Moncef, B., & Durand, B. (2018a). Building a collaborative solution in dense urban city settings to enhance parcel delivery: An effective crowd model in Paris. *Transportation Research Part E: Logistics and Transportation Review*, *119*, 223–233. https://doi.org/10.1016/j.tre.2018.04.007

Akeb, H., Moncef, B., & Durand, B. (2018b). Building a collaborative solution in dense urban city settings to enhance parcel delivery: an effective crowd model in Paris. *Transportation Research Part E: Logistics and Transportation Review*, *119*, 223–233. https://doi.org/10.1016/j.tre.2018.04.007

Aliev, K., Antonelli, D., Awouda, A., & Chiabert, P. (2019). Key performance indicators Integrating collaborative and mobile robots in the factory networks. In *IFIP advances in information and communication technology* (pp. 635–642). https://doi.org/10.1007/978-3-030-28464-0_56

Allen, J., Piecyk, M., Piotrowska, M., McLeod, F., Cherrett, T., Ghali, K., Nguyen, T., Bektaş, T., Bates, O., Friday, A., Wise, S., & Austwick, M. Z. (2018). Understanding the impact of e-commerce on last-mile light goods vehicle activity in urban areas: The case of London. *Transportation Research Part D: Transport and Environment*, *61*, 325–338. https://doi.org/10.1016/j.trd.2017.07.020

Alumni website University of Twente. (n.d.). Universiteit Twente. https://www.utwente.nl/en/alumni/inspiring-alumni/ut-canon/stories/opening-tht/

Alverhed, E., Hellgren, S., Isaksson, H., Olsson, L., Palmqvist, H., & Flodén, J. (2024). Autonomous last-mile delivery robots: a literature review. *European Transport Research Review*, *16*(1). https://doi.org/10.1186/s12544-023-00629-7

An introduction to networks - Math Insight. (n.d.). https://mathinsight.org/network_introduction#:~:text=A%20network%20is%20simply%20a,them%20as%20lines%20between%20points.

Autonomous Delivery Robots Market Insights. (n.d.). https://www.mordorintelligence.com/industry-reports/autonomous-delivery-robots-market

Autoriteit Consument & Markt Dashboard. (2024). https://public.tableau.com/views/Post-enpakketmonitor/MARKTAANDELEN?%3Alanguage=en-US&%3Asid=&%3Adisplay_count=n&%3Aorigin=viz_share_link&publish=yes&%3AshowVizHome=no#1

AVENUE21. Planning and policy considerations for an age of Automated Mobility. (2023). In *Springer eBooks*. https://doi.org/10.1007/978-3-662-67004-0

Bahabry, A., Wan, X., Ghazzai, H., Menouar, H., Vesonder, G., & Massoud, Y. (2019). Low-Altitude navigation for Multi-Rotor drones in urban areas. *IEEE Access*, *7*, 87716–87731. https://doi.org/10.1109/access.2019.2925531

Bakach, I., Campbell, A. M., & Ehmke, J. F. (2021). A two-tier urban delivery network with robot-based deliveries. *Networks*, *78*(4), 461–483. https://doi.org/10.1002/net.22024

Bakach, I., Campbell, A. M., & Ehmke, J. F. (2022). Robot-Based Last-Mile deliveries with pedestrian zones. *Frontiers in Future Transportation*, *2*. https://doi.org/10.3389/ffutr.2021.773240

Baldacci, R., Mingozzi, A., & Roberti, R. (2012). Recent exact algorithms for solving the vehicle routing problem under capacity and time window constraints. *European Journal of Operational Research*, *218*(1), 1–6. https://doi.org/10.1016/j.ejor.2011.07.037

Bellan, R. (2024, November 19). Nuro expands driverless autonomous vehicle testing in push to attract customers. *TechCrunch*. https://techcrunch.com/2024/11/19/nuro-expands-driverless-autonomous-vehicle-testing-in-push-to-attract-customers/

Boeing, G. (2025). Modeling and analyzing urban networks and amenities with OSMNX. *Geographical Analysis*. https://doi.org/10.1111/gean.70009

Boysen, N., Fedtke, S., & Schwerdfeger, S. (2020). 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, N., Schwerdfeger, S., & Weidinger, F. (2018). Scheduling last-mile deliveries with truck-based autonomous robots. *European Journal of Operational Research*, *271*(3), 1085–1099. https://doi.org/10.1016/j.ejor.2018.05.058

Bruneteau, F., & Bruneteau, F. (2022, August 25). *What are the six levels of vehicle automation? - PTOLEMUS Consulting Group*. PTOLEMUS Consulting Group. https://www.ptolemus.com/what-are-the-six-levels-of-vehicle-automation/

Buldeo, H., Touami, S., & Dablanc, L. (2022a). Autonomous e-commerce delivery in ordinary and exceptional circumstances. The French case. *Research in Transportation Business & Management*, *45*, 100774. https://doi.org/10.1016/j.rtbm.2021.100774

Buldeo, H., Touami, S., & Dablanc, L. (2022b). Autonomous e-commerce delivery in ordinary and exceptional circumstances. The French case. *Research in Transportation Business & Management*, *45*, 100774. https://doi.org/10.1016/j.rtbm.2021.100774

Campus | Universiteit Twente. (n.d.). Universiteit Twente. https://www.utwente.nl/campus/#van-nanolab-tot-starbucks

Capgemini Research Institute. (2019). The last-mile delivery challenge. In *Sogeti*. Capgemini. Retrieved January 23, 2024, from https://www.sogeti.com/explore/reports/the-last-mile-delivery-challenge/thank-you/

Chen, C. (2023, October 18). *Amazon reveals first photos of the new Prime Air delivery drone*. US About Amazon. https://www.aboutamazon.com/news/transportation/amazon-prime-air-drone-delivery-mk30-photos#:~:text=These%20drones%20use%20vertical%20take,in%20one%20hour%20or%20less.

Chen, C., Demir, E., Huang, Y., & Qiu, R. (2021a). The adoption of self-driving delivery robots in last mile logistics. *Transportation Research Part E Logistics and Transportation Review*, *146*, 102214. https://doi.org/10.1016/j.tre.2020.102214

Chen, C., Demir, E., Huang, Y., & Qiu, R. (2021b). The adoption of self-driving delivery robots in last mile logistics. *Transportation Research Part E: Logistics and Transportation Review*, *146*, 102214. https://doi.org/10.1016/j.tre.2020.102214

Clarke, S., Leonardi, J., & Greater London Authority. (2017). *Consolidating deliveries of retail and non-retail clients into single electric van deliveries*. Greater London Authority. https://www.london.gov.uk/sites/default/files/single_carrier_consolidation_trial_using_electric_vehicles_finalreport-.pdf

Condliffe, J. (2022, June 29). Why sidewalk delivery robots still need safety drivers. *MIT Technology Review*. https://www.technologyreview.com/2018/01/30/145935/why-sidewalk-delivery-robots-still-need-safety-drivers-too/

Depot Definition | GIS Dictionary. (n.d.). https://support.esri.com/en-us/gis-dictionary/depot

DJI DB2000 Intelligent Flight Battery. (n.d.). DroneLand. https://droneland.nl/dji-db2000-intelligent-flight-battery

D'Onfro, J. (2019, June 5). *Amazon's new delivery drone will start shipping packages "In a matter of months."* Forbes. https://www.forbes.com/sites/jilliandonfro/2019/06/05/amazon-new-delivery-drone-remars-warehouse-robots-alexa-prediction/#41ce1326145f

Dopson, E. (2023, October 6). *What is Logistics? Definition and Components*. Shopify. https://www.shopify.com/blog/what-is-logistics

DPG Media Privacy Gate. (n.d.). https://www.autoweek.nl/kentekencheck/VVZ-41-G/

Drone Delivery FAQs - Amazon Customer Service. (n.d.). https://www.amazon.com/gp/help/customer/display.html?nodeId=T3jxhuvPfQ629BOIL4

Ecommerce News Europe. (2020, October 1). *The Future Mobility Network tests delivery robot*. Ecommerce News. https://ecommercenews.eu/the-future-mobility-network-tests-delivery-robot/

Een Postgeschiedenis. (n.d.). *Het Pakkettennetwerk*. https://een-postgeschiedenis.jouwweb.nl/reis-door-de-tijd/pakketten/het-pakkettennetwerk?

Electric vs. Gas Cars: Is It Cheaper to Drive an EV? (2023, November 17). https://www.nrdc.org/stories/electric-vs-gas-cars-it-cheaper-drive-ev

Engadget is part of the Yahoo family of brands. (n.d.). https://www.engadget.com/nuro-lays-off-30-percent-of-staff-shifts-focus-to-rd-204646216.html

Engesser, V., Rombaut, E., Vanhaverbeke, L., & Lebeau, P. (2023). Autonomous Delivery Solutions for Last-Mile Logistics Operations: A literature review and research agenda. *Sustainability*, *15*(3), 2774. https://doi.org/10.3390/su15032774

Express Robot. (2025, January 17). Macrostep BV. https://www.macrostep.eu/producten/sectoren/express-robot.html

Extending the supply chain: how cutting-edge companies bridge the critical last mile into customers' homes. (2005). *Choice Reviews Online*, *42*(05), 42–2901. https://doi.org/10.5860/choice.42-2901

Feiten en cijfers postbezorging. (n.d.). https://www.postnl.nl/api/assets/blt43aa441bfc1e29f2/bltb088c95d9bd21d09/infographic-postbezorging-bij-postnl.pdf

Figliozzi, M. (2017). Lifecycle modeling and assessment of unmanned aerial vehicles (Drones) CO 2 e emissions. *Transportation Research Part D: Transport and Environment*, *57*, 251–261. https://doi.org/10.1016/j.trd.2017.09.011

Figliozzi, M. (2020). Carbon emissions reductions in last mile and grocery deliveries utilizing air and ground autonomous vehicles. *Transportation Research Part D: Transport and Environment*, *85*, 102443. https://doi.org/10.1016/j.trd.2020.102443

Frequently asked questions. (n.d.). https://lmudining.sodexomyway.com/kiwibot/kiwibot_FAQ#:~:text=What%20is%20Kiwibot%3F,built%20more%20than%20400%20robots.

Gao, M., Hugenholtz, C. H., Fox, T. A., Kucharczyk, M., Barchyn, T. E., & Nesbit, P. R. (2021). Weather constraints on global drone flyability. *Scientific Reports*, *11*(1). https://doi.org/10.1038/s41598-021-91325-w

Gehrke, S. R., Phair, C. D., Russo, B. J., & Smaglik, E. J. (2023). Observed sidewalk autonomous delivery robot interactions with pedestrians and bicyclists. *Transportation Research Interdisciplinary Perspectives*, *18*, 100789. https://doi.org/10.1016/j.trip.2023.100789

Genova, K., & Guliashki, V. (2011). LINEAR INTEGER PROGRAMMING METHODS AND APPROACHES-A SURVEY. *CYBERNETICS AND INFORMATION TECHNOLOGIES*, *11*(1), 0. https://www.sid.ir/En/Journal/ViewPaper.aspx?ID=395331

Gevaers, R., Van De Voorde, E., & Vanelslander, T. (2009). Characteristics of Innovations in Last-Mile Logistics - Using Best Practices, Case Studies and Making the Link with Green and Sustainable Logistics. *European Transport Conference, 2009Association for European Transport (AET)*. https://trid.trb.org/view/1107297

Gevaers, R., Van De Voorde, E., & Vanelslander, T. (2011). Characteristics and Typology of Last-mile Logistics from an Innovation Perspective in an Urban Context. In *Edward Elgar Publishing eBooks*. https://doi.org/10.4337/9780857932754.00009

Giret, A., Julián, V., Corchado, J. M., Fernández, A., Salido, M. A., & Tang, D. (2018). How to choose the Greenest Delivery Plan: A framework to measure key performance indicators for sustainable urban logistics. In *IFIP advances in information and communication technology* (pp. 181–189). https://doi.org/10.1007/978-3-319-99707-0_23

Gosling, S. (2025, February 13). *What do US consumers want from e-commerce deliveries?* McKinsey & Company. https://www.mckinsey.com/industries/logistics/our-insights/what-do-us-consumers-want-from-e-commerce-deliveries

Ha, N., Akbari, M., & Au, B. (2022). Last mile delivery in logistics and supply chain management: a bibliometric analysis and future directions. *Benchmarking: An International Journal*, *30*(4), 1137–1170. https://doi.org/10.1108/bij-07-2021-0409

Hoffman, K., & Padberg, M. (2001). Traveling salesman problem. In *Springer eBooks* (pp. 849–853). https://doi.org/10.1007/1-4020-0611-x_1068

Huang, Y., Savelsbergh, M., & Zhao, L. (2018). Designing logistics systems for home delivery in densely populated urban areas. *Transportation Research. Part B: Methodological/Transportation Research. Part B, Methodological*, *115*, 95–125. https://doi.org/10.1016/j.trb.2018.07.006

Iddenden, G. (2022, November 14). Amazon unveils new MK30 delivery drone - Latest Retail Technology News From Across The Globe - Charged. *Charged Retail*. https://www.chargedretail.co.uk/2022/11/14/amazon-unveils-new-delivery-drone-mk30/

Industry - Starship deliveries. (2024, August 2). Starship Deliveries. https://starshipdeliveries.com/industry/#:~:text=Green%20technology,model%20working%20towards%20CO2%20neutrality

Ioniță, S. (2017). Autonomous vehicles: From paradigms to technology. *IOP Conference Series*, *252*, 012098. https://doi.org/10.1088/1757-899x/252/1/012098

Jennings, D. (2020). *A study of road autonomous delivery robots and their potential impacts on freight efficiency and travel*. PDXScholar. https://pdxscholar.library.pdx.edu/cengin_fac/549

Jennings, D., & Figliozzi, M. (2019a). Study of sidewalk autonomous delivery robots and their potential impacts on freight efficiency and travel. *Transportation Research Record*, *2673*(6), 317–326. https://doi.org/10.1177/0361198119849398

Jennings, D., & Figliozzi, M. (2019b). Study of sidewalk autonomous delivery robots and their potential impacts on freight efficiency and travel. *Transportation Research Record*, *2673*(6), 317–326. https://doi.org/10.1177/0361198119849398

Jeon, A., Kang, J., Choi, B., Kim, N., Eun, J., & Cheong, T. (2021). Unmanned aerial vehicle Last-Mile delivery considering backhauls. *IEEE Access*, *9*, 85017–85033. https://doi.org/10.1109/access.2021.3087751

Kahl, M. (2020, June 9). Rethinking last-mile logistics, post-COVID-19: Facing the 'next normal.' *Automotive World*. https://www.automotiveworld.com/articles/rethinking-last-mile-logistics-post-covid-19-facing-the-next-normal/

Kitjacharoenchai, P., & Lee, S. (2019). Vehicle Routing Problem with Drones for Last Mile Delivery. *Procedia Manufacturing*, *39*, 314–324. https://doi.org/10.1016/j.promfg.2020.01.338

Laporte, G., Toth, P., & Vigo, D. (2013). Vehicle routing: historical perspective and recent contributions. *EURO Journal on Transportation and Logistics*, *2*(1–2), 1–4. https://doi.org/10.1007/s13676-013-0020-6

Li, X., Tupayachi, J., Sharmin, A., & Ferguson, M. M. (2023). Drone-Aided Delivery Methods, Challenge, and the Future: A Methodological review. *Drones*, *7*(3), 191. https://doi.org/10.3390/drones7030191

Liu, F., a, Lu, C., a, Gui, L., a,b, Zhang, Q., a, Tong, X., c, Yuan, M., c, Computer Science, City Univeristy of Hong Kong, Hong Kong, State Key Laboratory of Digital Manufacturing Equipment and Technology, Huazhong University of Science and Technology, Wuhan, China, & Huawei Noah's Ark Lab, Hong Kong. (2023). Heuristics for Vehicle Routing Problem: A survey and Recent advances [Journal-article]. *arXiv*. https://arxiv.org/abs/2303.04147v1

Liu, W., Lin, C., Chiu, C. R., Tsao, Y. S., & Wang, Q. (2014). Minimizing the Carbon Footprint for the Time-Dependent Heterogeneous-Fleet Vehicle Routing Problem with Alternative Paths. *Sustainability*, *6*(7), 4658–4684. https://doi.org/10.3390/su6074658

Lu, Y., Yang, J., & Yang, C. (2023). A humanitarian vehicle routing problem synchronized with drones in time-varying weather conditions. *Computers & Industrial Engineering*, *184*, 109563. https://doi.org/10.1016/j.cie.2023.109563

MacroStep BV. (2025, January 22). Macrostep BV. https://www.macrostep.eu/

Maersk. (2025, March 20). Closing the logistics loop with last-mile delivery. *Maersk*. https://www.maersk.com/logistics-explained/transportation-and-freight/2025/02/07/last-mile-delivery

Map features. (n.d.). OpenStreetMap Wiki. https://wiki.openstreetmap.org/wiki/Map_features

Marques Brownlee. (2025, March 25). *The Truth about Drone Deliveries!* [Video]. YouTube. https://www.youtube.com/watch?v=88yQTzlmsiA

McKinsey & Company. (2016). Parcel Delivery: The future of last mile. In *McKinsey & Company*. Retrieved January 11, 2024, from https://bdkep.de/files/bdkep-dateien/pdf/2016_the_future_of_last_mile.pdf

McKinsey&Company. (2018). Fast forwarding last-mile delivery - implications for the ecosystem. In *McKinsey & Company*. Retrieved January 11, 2024, from https://www.mckinsey.com/~/media/mckinsey/industries/travel%20logistics%20and%20infrastructure/our%20insights/technology%20delivered%20implications%20for%20cost%20customers%20and%20competition%20in%20the%20last%20mile%20ecosystem/fast-forwarding-last-mile-delivery-implications-for-the-ecosystem.pdf

Mercedes-Benz Media. (n.d.). Mercedes-Benz Media. https://media.mercedes-benz.com/article/594669ab-1a6f-45f3-859d-42b3f8912435

Ministerie van Infrastructuur en Waterstaat. (2023, February 13). *Self-driving vehicles*. Mobility, Public Transport and Road Safety | Government.nl. https://www.government.nl/topics/mobility-public-transport-and-road-safety/self-driving-vehicles

Ministerie van Infrastructuur en Waterstaat. (2024, January 4). *Nederland als toepassingsland voor slimme mobiliteit*. Mobiliteit Nu En in De Toekomst | Rijksoverheid.nl. https://www.rijksoverheid.nl/onderwerpen/mobiliteit-nu-en-in-de-toekomst/nederland-als-testland-voor-mobiliteit

Moadab, A., Farajzadeh, F., & Valilai, O. F. (2022). Drone routing problem model for last-mile delivery using the public transportation capacity as moving charging stations. *Scientific Reports*, *12*(1). https://doi.org/10.1038/s41598-022-10408-4

Morana, J., & Gonzalez-Feliu, J. (2015). A sustainable urban logistics dashboard from the perspective of a group of operational managers. *Management Research Review*, *38*(10), 1068–1085. https://doi.org/10.1108/mrr-11-2014-0260

Moshref-Javadi, M., Van Cauwenberghe, K. P., McCunney, B. A., & Hemmati, A. (2023). Enabling same-day delivery using a drone resupply model with transshipment points. *Computational Management Science*, *20*(1). https://doi.org/10.1007/s10287-023-00453-3

Murray, C., & Chu, A. (2015). The flying sidekick traveling salesman problem: Optimization of drone-assisted parcel delivery. *Transportation Research Part C: Emerging Technologies*, *54*, 86–109. https://doi.org/10.1016/j.trc.2015.03.005

Nederlandse consument geeft online €33,3 miljard uit in 2022. (n.d.). Thuiswinkel.org. https://www.thuiswinkel.org/webshops/nieuws/nederlandse-consument-geeft-online-33-3-miljard-uit-in-2022/#:~:text=Nederlanders%20hebben%20in%202022%20voor,met%205%25%20naar%20347%20miljoen.

Neolix. (n.d.). https://www.neolix.ai/

Netherlands Enterprise Agency, RVO. (2023, November 6). *Rules for flying drones*. business.gov.nl. https://business.gov.nl/regulation/drones/

NetworkX Documentation. (n.d.). https://networkx.org/

Nishira, M., Ito, S., Nishikawa, H., Kong, X., & Tomiyama, H. (2023). An Integer Programming Based Approach to Delivery Drone Routing under Load-Dependent Flight Speed. *Drones*, *7*(5), 320. https://doi.org/10.3390/drones7050320

Nuro. (2020). *Nuro R2 specifications* [Press release]. https://nuro.sfo3.digitaloceanspaces.com/Nuro_R2-Spec-Sheet-Letter.pdf

Nuro. (2022). *Nuro R3 specifications* [Press release]. https://nuro.sfo3.digitaloceanspaces.com/Nuro_Nuro-Spec-Sheet-Letter.pdf

Onfleet, Inc. (2023, November 22). *What is Last Mile Delivery? Everything You Need to Know*. Onfleet Last Mile Delivery Blog | Onfleet. https://onfleet.com/blog/what-is-last-mile-delivery/

Op welke dagen bezorgt PostNL? | PostNL. (n.d.). https://www.postnl.nl/klantenservice/algemene-bezorgdagen-tijden/welke-dagen/

OSM Universiteit Twente Service Area. (n.d.). OpenStreetMap. https://www.openstreetmap.org/way/144601742#map=15/52.24346/6.85458

Ostermeier, M., Heimfarth, A., & Hübner, A. (2021). Cost-optimal truck-and-robot routing for last-mile delivery. *Networks*, *79*(3), 364–389. https://doi.org/10.1002/net.22030

Othman, M. S., Shurbevski, A., Karuno, Y., & Nagamochi, H. (2017). Routing of Carrier-vehicle Systems with Dedicated Last-stretch Delivery Vehicle and Fixed Carrier Route. *Journal of Information Processing*, *25*(0), 655–666. https://doi.org/10.2197/ipsjjip.25.655

Overstappen.nl. (2025, June 4). *Stroomprijs 2025 | Bekijk de kWh prijs van 20 leveranciers*. https://www.overstappen.nl/energie/stroomprijs/#:~:text=Momenteel%20betreft%20de%20stroomprijs%20gemiddeld%20%E2%82%AC%200%2C32%20per%20kWh.&text=Wij%20berekenen%20de%20gemiddelde%20stroomprijs,voor%20bestaande%20klanten%20kunnen%20afwijken.

Poeting, M., Schaudt, S., & Clausen, U. (2019). A Comprehensive Case Study in Last-Mile Delivery Concepts for Parcel Robots. *Winter Simulation 2019*. https://doi.org/10.1109/wsc40007.2019.9004811

Pop, P. C., Cosma, O., Sabo, C., & Sitar, C. P. (2023). A comprehensive survey on the generalized traveling salesman problem. *European Journal of Operational Research*, *314*(3), 819–835. https://doi.org/10.1016/j.ejor.2023.07.022

POSTNL. (2023). Rates for mail, packets and parcels. In *PostNL*. Retrieved November 29, 2023, from https://www.postnl.nl/api/assets/blt43aa441bfc1e29f2/blt5797459cb3672f36/64a281e45323fe6ad8c8397f/def-tarievenfolder-per-1-juli-2023-enn.pdf

Railway-News. (2021, November 26). *DPD Netherlands trials Cleveron's autonomous delivery vehicle*. Future Transport-News. https://futuretransport-news.com/dpd-netherlands-trials-cleverons-autonomous-delivery-vehicle/#:~:text=Autonomous%20Vehicles,-by%20Future%20Transport&text=Cleveron's%20unmanned%20autonomous%20vehicle%20took,public%20streets%20of%20the%20Netherlands.

Rashidi, H., Matinfar, F., & Parand, F. (2021). Automated Guided Vehicles - A review on applications, problem modeling and solutions. *International Journal of Transportation Engineering*, *8*(3), 261–278. https://doi.org/10.22119/ijte.2021.246669.1531

Raychaudhuri, S. (2008). Introduction to Monte Carlo simulation. *2008 Winter Simulation Conference*. https://doi.org/10.1109/wsc.2008.4736059

Reed, S. M., Campbell, A. M., & Thomas, B. W. (2022). The value of autonomous vehicles for Last-Mile deliveries in urban environments. *Management Science*, *68*(1), 280–299. https://doi.org/10.1287/mnsc.2020.3917

Robinson, S. (1997). Simulation model verification and validation. *1997 Winter Simulation Conference*, 53–59. https://doi.org/10.1145/268437.268448

Sartika, H., & Gamal, M. D. H. (2019). Depot location problems by considering its distribution problems. *Journal of Research in Mathematics Trends and Technology*, *1*(1), 18–26. https://doi.org/10.32734/jormtt.v1i1.753

Scherr, Y. O., Neumann-Saavedra, B. A., Hewitt, M., & Mattfeld, D. C. (2019). Service network design with mixed autonomous fleets. *Transportation Research Part E: Logistics and Transportation Review*, *124*, 40–55. https://doi.org/10.1016/j.tre.2019.02.001

Seuring, S., & Müller, M. (2008). From a literature review to a conceptual framework for sustainable supply chain management. *Journal of Cleaner Production*, *16*(15), 1699–1710. https://doi.org/10.1016/j.jclepro.2008.04.020

Simoni, M. D., Kutanoğlu, E., & Claudel, C. (2020). Optimization and analysis of a robot-assisted last mile delivery system. *Transportation Research Part E: Logistics and Transportation Review*, *142*, 102049. https://doi.org/10.1016/j.tre.2020.102049

Singhal, G., Bansod, B. K. S., & Mathew, L. (2018). Unmanned Aerial Vehicle Classification, Applications and Challenges: A Review. *Preprint*. https://doi.org/10.20944/preprints201811.0601.v1

Sivak, M., & Schoettle, B. (2015). Road safety with self-driving vehicles: general limitations and road sharing with conventional vehicles. *UMTRI*. https://deepblue.lib.umich.edu/handle/2027.42/111735

Solomon, M. M. (1987). Algorithms for the Vehicle Routing and Scheduling Problems with Time Window Constraints. *Operations Research*, *35*(2), 254–265. https://doi.org/10.1287/opre.35.2.254

Sonneberg, M., Leyerer, M., Kleinschmidt, A., Knigge, F., & Breitner, M. H. (2019). Autonomous unmanned ground vehicles for urban logistics: Optimization of last mile delivery operations. *Proceedings of the . . . Annual Hawaii International Conference on System Sciences*. https://doi.org/10.24251/hicss.2019.186

Soper, S., & Day, M. (2022, October 6). Amazon abandons autonomous Home-Delivery robot scout in latest cut (AMZN). *Bloomberg.com*. https://www.bloomberg.com/news/articles/2022-10-06/amazon-abandons-autonomous-home-delivery-robot-in-latest-cut

Srinivas, S., Ramachandiran, S., & Rajendran, S. (2022). Autonomous Robot-driven Deliveries: A review of recent developments and future directions. *Transportation Research Part E: Logistics and Transportation Review*, *165*, 102834. https://doi.org/10.1016/j.tre.2022.102834

STARSHIP DELIVERY ROBOT. (n.d.). https://www.post.ch/-/media/post/ueber-uns/medienmitteilungen/2017/factsheet-lieferroboter.pdf?la=en

Starship Robot Dimensions & Drawings | Dimensions.com. (n.d.). Dimensions.com. https://www.dimensions.com/element/starship-robot

Starship Technologies. (2023, October 19). *About Starship: The global leader in autonomous delivery - Starship Technologies: Autonomous robot delivery*. Starship Technologies: Autonomous Robot Delivery. https://www.starship.xyz/company/

Starship Technologies. (2024a, January 5). *How far can the robots deliver? - Starship Technologies: Autonomous robot delivery*. Starship Technologies: Autonomous Robot Delivery. https://www.starship.xyz/faqs/how-far-can-the-robots-deliver/

Starship Technologies. (2024b, May 1). *The Starship Robot FAQs - Starship Technologies: Autonomous robot delivery*. Starship Technologies: Autonomous Robot Delivery. https://www.starship.xyz/faqs/category/the-starship-robot/

Statista. (2023, August 29). *Retail e-commerce sales growth worldwide 2017-2027*. https://www.statista.com/statistics/288487/forecast-of-global-b2c-e-commerce-growth/

Tarasov, K. (2022, November 12). A first look at Amazon's new delivery drone, slated to start deliveries this year. *CNBC*. https://www.cnbc.com/2022/11/11/a-first-look-at-amazons-new-delivery-drone.html

Team, W. (2023, December 13). *Fast-Moving consumer goods*. WallStreetMojo. https://www.wallstreetmojo.com/fast-moving-consumer-goods/

Thomas, R., & Tokar, T. (n.d.). *Reducing the cost of last mile delivery: a channel shifting behavioral approach*. Walton College. https://walton.uark.edu/initiatives/supply-chain-research/working_paper-reducing_the_cost_of_last_mile_delivery.php

Transporter | udelv. (n.d.). https://udelv.com/transporter/

Understanding difference between Polygon and Multipolygon for shapefiles in QGIS? (n.d.). Geographic Information Systems Stack Exchange. https://gis.stackexchange.com/questions/225368/understanding-difference-between-polygon-and-multipolygon-for-shapefiles-in-qgis

Understanding last mile delivery for business success. (2025, March 30). https://packagex.io/blog/what-is-last-mile-delivery

University map - Materials Science and Technology of Polymers. (n.d.). https://archive.mtpgroup.nl/university-map.aspx

University of Twente. (n.d.). *Map of the University of Twente*. Retrieved February 29, 2024, from https://www.utwente.nl/download/campusmap.pdf

Vakulenko, Y., Shams, P., Hellström, D., & Hjort, K. (2019). Online retail experience and customer satisfaction: the mediating role of last mile delivery. *The International Review of Retail, Distribution and Consumer Research*, *29*(3), 306–320. https://doi.org/10.1080/09593969.2019.1598466

Veelgestelde vragen | DHL eCommerce. (n.d.). https://www.dhlecommerce.nl/nl/consument/faq#:~:text=Wij%20bezorgen%20je%20pakket%20van,met%20zaterdag%2C%20in%20verschillende%20tijdvensters.&text=Waar%20zie%20ik%20mijn%20pakketstatus,van%20je%20pakket%20te%20checken.

Vivatechnology. (2025). *Autonomous Delivery: Robots delivering the future of logistics*. https://vivatechnology.com/news/autonomous-delivery-robots-delivering-the-future-of-logistics

Vleugel, D. (2023, April 6). *PostNL kiest voor de ID.Buzz elektrische busjes van Volkwagen* -. https://logistiekmet.nl/postnl-id-buzz-elektrische-busjes-volkswagen/

Waarom een elektrische auto geluid maakt | ANWB. (n.d.). ANWB. https://www.anwb.nl/auto/elektrisch-rijden/techniek-en-onderhoud/geluid-elektrische-auto

Watkins, S., Burry, J., Abdel-Rohman, M., Marino, M., Prudden, S., Fisher, A., Kloet, N., Jakobi, T., & Clothier, R. (2020). Ten questions concerning the use of drones in urban environments. *Building and Environment*, *167*, 106458. https://doi.org/10.1016/j.buildenv.2019.106458

Weinberg, D. A., Dwyer, H. A., Fox, S., & Martelaro, N. (2023). Sharing the Sidewalk: Observing Delivery Robot Interactions with Pedestrians during a Pilot in Pittsburgh, PA. *Multimodal Technologies and Interaction*, *7*(5), 53. https://doi.org/10.3390/mti7050053

Wessling, B. (2022, January 12). *Nuro unveils third-generation autonomous delivery vehicle*. The Robot Report. https://www.therobotreport.com/nuro-unveils-third-generation-autonomous-delivery-vehicle/

What is an Autonomous Car? (n.d.). Synopsys. https://www.synopsys.com/automotive/what-is-autonomous-car.html#:~:text=Autonomous%20cars%20rely%20on%20sensors,different%20parts%20of%20the%20vehicle

Winkle, T. (2016). Safety Benefits of Automated Vehicles: Extended Findings from Accident Research for Development, Validation and Testing. In *Springer eBooks* (pp. 335–364). https://doi.org/10.1007/978-3-662-48847-8_17

World Bank Open Data. (n.d.). World Bank Open Data. https://data.worldbank.org/indicator/SP.URB.TOTL.IN.ZS?amp%3Blocations=DE-GB-FR-NL-BE&amp%3Bstart=1960&amp%3Bview=chart&end=2022&most_recent_year_desc=false&start=2022&view=map

World Economic Forum. (2020). The Future of the Last-Mile Ecosystem. In *World Economic Forum*. Retrieved December 5, 2023, from https://www.weforum.org/publications/the-future-of-the-last-mile-ecosystem/

Wu, P., Chu, F., Che, A., & Fang, Y. (2017). An efficient two-phase exact algorithm for the automated truck freight transportation problem. *Computers & Industrial Engineering*, *110*, 59–66. https://doi.org/10.1016/j.cie.2017.04.030

Yahoo is part of the Yahoo family of brands. (n.d.). https://au.news.yahoo.com/new-zealand-woman-finds-drone-hovering-outside-her-room-39299704.html?guccounter=1&guce_referrer=aHR0cHM6Ly93d3cuZ29vZ2xlLmNvbS8&guce_referrer_sig=AQAAADFAeH5GJIH394R--XyjndW6q0plCN8zekM0Uk1SSUnp8QzKLUHg7RdSo_pcdVnudFAZx51wHR4yOysOuSSXtFd1GDEvpvM8YlD0RKcF2NmQH2FSgMhca_mFhd1K3g3-W4swD9QdKmfxKd9m4I9qQpex8O_fge-QoR_DIiNsguSS

Young, J. (2023, July 29). *How Much Weight Can A Drone Carry? Real Cases (lbs & Kg) - DroneGuru*. DroneGuru. https://www.droneguru.net/weight-a-drone-carry/#:~:text=A%3A%20The%20amount%20of%20weight,(44%20pounds)%20or%20more.

Zipline fact sheet. (n.d.). Zipline. https://www.flyzipline.com/about/zipline-fact-sheet

Zis, T. P., Psaraftis, H. N., & Reche-Vilanova, M. (2023). Design and application of a key performance indicator (KPI) framework for autonomous shipping in Europe. *Maritime Transport Research*, *5*, 100095. https://doi.org/10.1016/j.martra.2023.100095

Appendix A) Detailed KPI selection

Economic dimension

The economic dimension of delivery systems evaluates efficiency, reliability, cost-effectiveness, and customer satisfaction. These KPIs provide insights into the financial sustainability and viability of logistics operations, particularly in autonomous parcel delivery.

1. General

- **Distance (km)** – Total Distance
- **Moving time (hrs/mins)** – Total moving time
- **Loading/Unloading time (hrs/mins)** – Total unloading/loading time
- **Average Delivery Time (mins) -**Time taken from pickup to delivery
- **Delivery throughput (parcels/hour/day)** – number of parcels delivered per timeframe
- **Idle Rate / Utilization rate (%) -** Percentage of time a vehicle is doing something 'useful' such as loading, unloading, moving, charging.

2. Reliability and Performance

- **On-Time Delivery Rate –** Percentage of deliveries completed within the expected timeframe
- **Deviation Time (mins) –** Total Time deviating from customers time windows
- **Error rate (%) –** Number of incorrect or failed deliveries
- **Mean time between failures** – Average time between system failures

3. Cost efficiency

- **Cost per delivery (€/parcel) -** Total cost per successful parcel
- **Energy Consumption per km (kWh/km)** – Evaluates energy efficiency
- **Maintenance Costs (€) –** Expenses for vehicle and system upkeep
- **Daily Operating Costs (€) –** Total Daily Operational Expense

4. Customer Satisfaction

- **Customer Satisfaction Score (%)**

Environmental dimension

Environmental KPIs focus on reducing the ecological footprint of logistics operations. These metrics are typically quantifiable and standardized, making them easier to track.

- **Total Energy Consumption (kWh/day) –** measures the total energy consumption of the system
- **CO2 emissions (kg/km) –** Total carbon footprint per ton-km
- **Nox, Sox emissions (kg/km) –** Measures harmful pollutants affecting air quality
- **Charging efficiency (%) –** Measures energy transfer efficiency
- **Noise Pollution (dB)** – Measures the sound emissions generated by logistics operations.

Social/Societal Dimensions

The social dimension is much more difficult to quantify, as it depends on survey-based data such as public perception or data like workforce impact and operational safety.

Safety & Security

- **Accident Rate (%) –** Measures the frequency of Accidents
- **Emergency Stop Rate (%)** – Number of times a vehicle executes an emergency stop
- **Failure to avoid obstacles (%) –** Rate of unsuccessful obstacle avoidance attempts
- **Crime incidents (#) –** Number of crime incidents

2. Workforce Impact

- **Employment Turnover (%) –** How much staff can keep their job
- **Training Hours per Employee –** How much training hours are required for employees.

3. Public Perception & Complaints

- **Public Perception (Score) –** Survey based score on public perception
- **Work-life balance score (Score)** – Survey based score on work-life balance
- **Public complaint rate (#) –** number of complaints per day/week
- **Total minutes late (mins)** – Total minutes of the ADVs being too late on a single day
- **Late customers –** Average number of customers being late on a single day
- **Average Time After TimeWindows –** Average time the customer must wait in their time window until they receive their package.

Appendix B) Latitude and Longitude polyline campus

This appendix shows the latitude and longitude coordinates of the campus area.

Node_Number	Latitude	Longitude	Node_Code
1	52.2346702	6.8609214	1581427650
2	52.2360103	6.8615901	1581427654
3	52.2367369	6.8623460	1581427658
4	52.2368486	6.8630453	1581427662
5	52.2368848	6.8640121	1581427671
6	52.2373811	6.8645515	1581427670
7	52.2394155	6.8648960	1581427678
8	52.2400050	6.8634598	1581427682
9	52.2405613	6.8622939	1581427688
10	52.2413688	6.8613765	1581427690
11	52.2420523	6.8611371	1581427692
12	52.2425458	6.8568649	1581427696
13	52.2439463	6.8568437	1581427702
14	52.2476731	6.8536048	1581427708
15	52.2486595	6.8519172	1581427712
16	52.2491214	6.8510241	1581427716
17	52.2490803	6.8548766	8390256680
18	52.2504145	6.8550101	8390256679
19	52.2503258	6.8559694	8390256678
20	52.2510970	6.8570688	8390256677
21	52.2510467	6.8573486	8390256676
22	52.2516628	6.8577781	8390256675
23	52.2522441	6.8481536	670918703
24	52.2518093	6.8448758	1581427731
25	52.2499176	6.8461647	1581427719
26	52.2481438	6.8460123	1581427709
27	52.2448187	6.8453539	1581427706
28	52.2438252	6.8450541	8224509753
29	52.2428283	6.8449027	7531709842
30	52.2406722	6.8442565	1581427689
31	52.2390723	6.8490913	1581427673
32	52.2380581	6.8516822	1581427672
33	52.2346702	6.8609214	1581427650

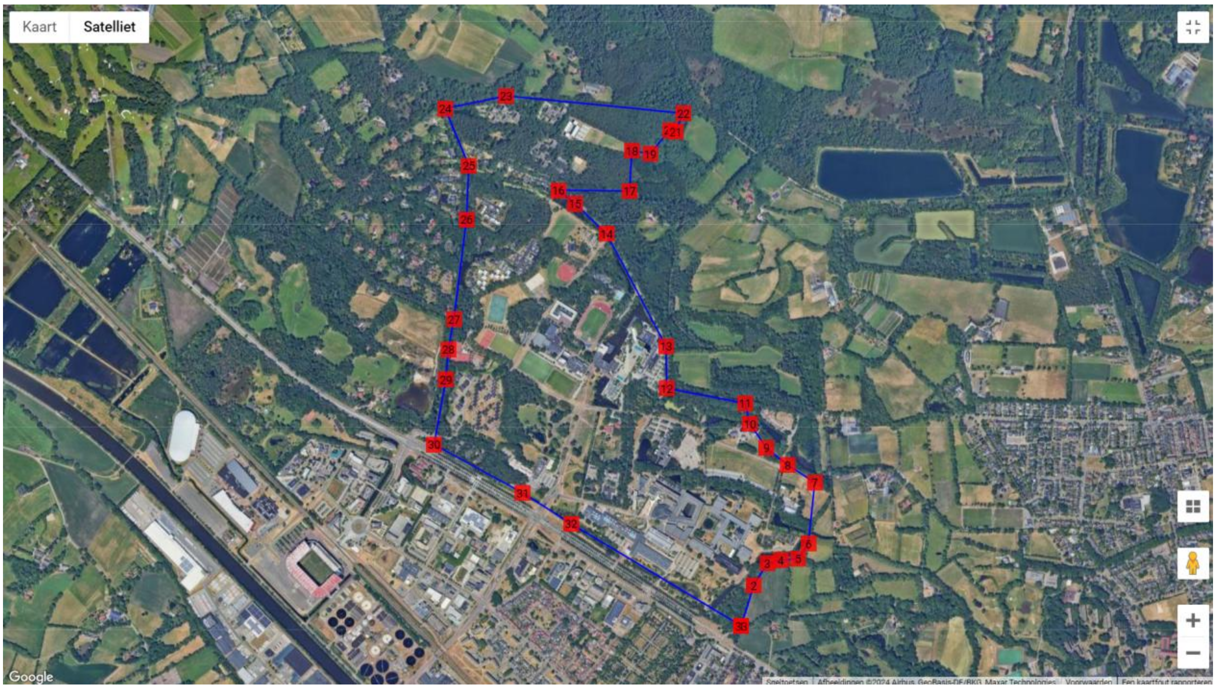

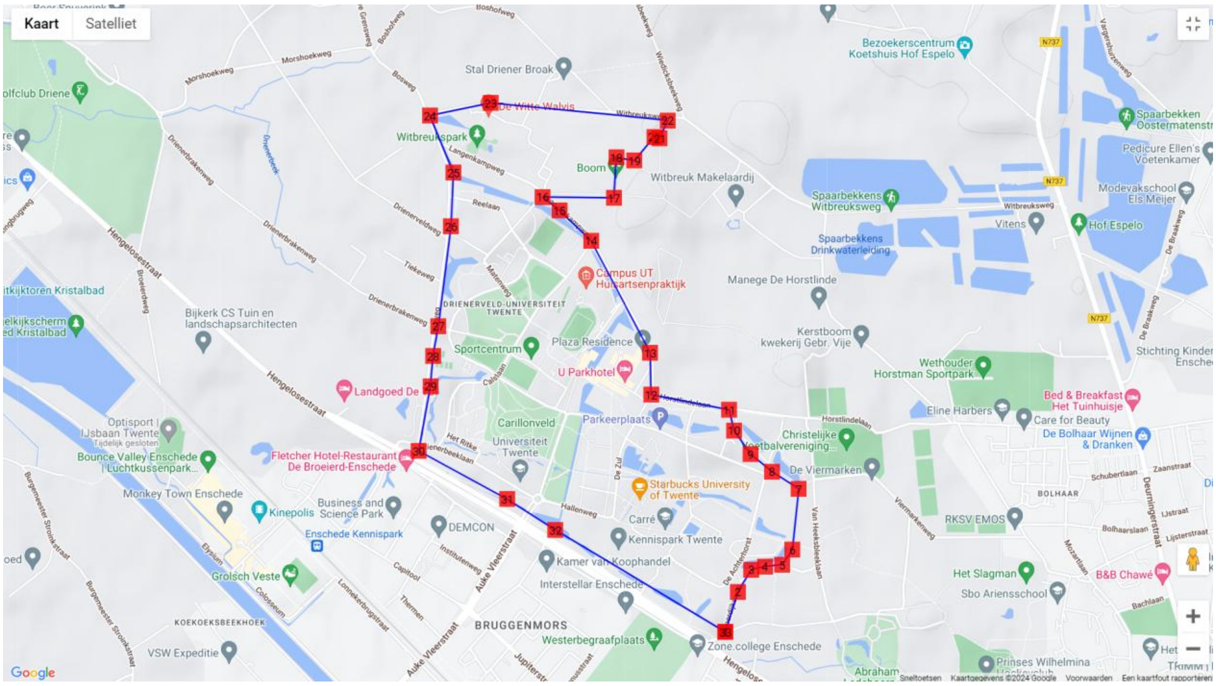

Appendix C) Extra demand calculations

This appendix describes the demand calculations in more detail. The first picture shows how the data is presented in the dataset, with daily total number of parcels per van round.

MORNING	INPUTS								
Date	ID of van round	Courier Number	Successful deliveries	Attempted & unsuccessful	Successful pickups	Successful collections	Carried fwd	Bulk Catalogue	Total number of parcels
01/Jul/15	1	789174	160	26	4	0	31	0	164
01/Jul/15	2	183648	197	11	10	5	8	0	212
01/Jul/15	3	202114	73	21	0	0	21	0	73
01/Jul/15	4	169505	262	15	4	2	11	0	268
01/Jul/15	5	762270	145	14	0	0	26	0	145
01/Jul/15	6	762233	81	8	2	0	11	0	83
01/Jul/15	7	183543	208	2	5	2	2	0	215
01/Jul/15	8	183580	98	8	5	1	10	0	104
01/Jul/15	9	169289	232	4	3	0	4	0	235
01/Jul/15	10	169362	256	1	9	1	1	0	266
01/Jul/15	11	169435	144	1	2	0	0	0	146
01/Jul/15	12	209698	139	18	10	1	18	0	150
01/Jul/15	13	169340	31	51	0	0	57	0	31
01/Jul/15	14	762249	157	19	2	8	17	0	167
01/Jul/15	15	169475	185	7	14	1	7	0	200
01/Jul/15	16	169718	138	17	8	3	18	0	149
01/Jul/15	17	202138	231	3	3	2	3	0	236
01/Jul/15	18	191082	13	107	0	0	112	0	13
01/Jul/15	19	762284	136	22	2	1	20	0	139
01/Jul/15	20	183516	64	13	0	2	14	0	66
01/Jul/15	21	762225	87	11	4	6	18	0	97
01/Jul/15	22	169297	130	4	5	1	2	0	136
01/Jul/15	23	169424	163	22	4	1	8	0	168
01/Jul/15	24	169734	245	4	2	0	3	0	247
01/Jul/15	25	169645	196	18	9	0	12	2	207
01/Jul/15	26	183594	107	24	2	2	22	0	111
01/Jul/15	27	169613	150	11	2	0	9	0	152
01/Jul/15	28	189204	187	12	10	2	9	0	199
01/Jul/15	29	183632	76	67	0	4	75	0	80
01/Jul/15	30	169184	122	18	1	0	12	0	123
01/Jul/15	31	183575	153	3	14	0	3	0	167
01/Jul/15	32	789212	174	6	12	2	9	0	188
01/Jul/15	33	202165	139	9	6	1	6	0	146
01/Jul/15	34	169510	100	40	1	1	38	0	102

As we can see, these are only datapoints from July 1st. The dataset has more than 13000 rows of delivery data, and by summing each row "Total number of parcels" we get to a total of 2,005,728 parcels delivered. Since we have daily volume data, we can calculate the weekly package distribution.

Week number	Percentile	Week number	Percentile	Week number	Percentile
1	**0,010605631**	19	0,016981872	37	0,016466847
2	**0,024004762**	20	0,02141069	38	0,015995696
3	**0,021054211**	21	0,020495312	39	0,01569755
4	**0,019482213**	22	0,018773741	40	0,016798398
5	**0,019314194**	23	0,018419755	41	0,018263203
6	**0,019793322**	24	0,02025151	42	0,018308075
7	**0,018604725**	25	0,018911846	43	0,018540908
8	**0,017184791**	26	0,019197528	44	0,018580295
9	**0,016850249**	27	0,026170062	45	0,01878122
10	**0,018712417**	28	0,018827089	46	0,020048092
11	**0,018204372**	29	0,017891767	47	0,02135485
12	0,017988989	30	0,013538233	48	0,022506552
13	0,014673981	31	0,017579661	49	0,033668091
14	0,015712507	32	0,019513124	50	0,038014146
15	0,018110142	33	0,01619363	51	0,033512537
16	0,01905693	34	0,015873048	52	0,019648237
17	0,017417126	35	0,017563706		
18	0,017962066	36	0,011490098		

Since March 14th falls in the middle of week 11, we sum the parcel distribution percentages from week 1 to 10 and add half of week 11. This results in 0.194709 (19.47%).

Next, we use this proportion to estimate the total annual parcel volume for campus deliveries. We known that around 5000 parcels were delivered during the first 10.5 weeks of the year. Dividing this by the 19.47% share gives an estimated annual volume:

$$\frac{5{,}000}{0{,}194709} = 25.680\ packages\ per\ year\ (PostNL)$$

Since PostNL holds 50% of the market share, the total number of parcels delivered annually is:

$$25{,}680 * 2 = 51.360\ packages\ per\ year\ (Total\ amount)$$

To determine the number of parcels suitable for autonomous delivery, it is important to consider size and weight constraints. Back in 2019, Amazon's CEO Worldwide consumer stated that: *"between 75 and 90% of Amazon deliveries could technically be handled by the UAV"* (D'Onfro, 2019). Since the UAVs have the most restrictive limitations in terms of payload and volume, any parcel that is eligible for UAV delivery is also suitable for S-ADVs and R-ADVs. Therefore, the share of parcels eligible for delivery via autonomous methods can be estimated at 75-90%. In this study we take the most conservative number 75%.

$$51.360 * 0{,}75 = 38{,}520\ packages\ per\ year\ (Total\ amount)$$

The graph below shows the expected number of packages delivered each week to the campus based on the total of 38,520 packages per year.

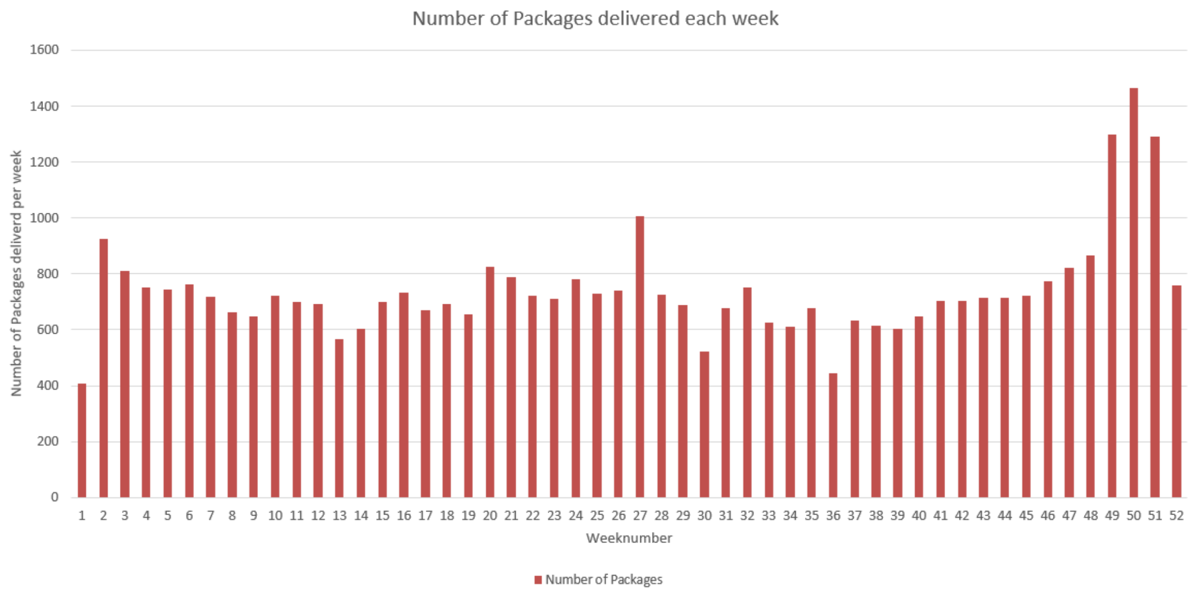

In addition to the weekly distribution of parcels we also calculated the daily and monthly distributions to gain insight into seasonal variations across different time periods. By combining weekly and daily distributions, we roughly estimate the daily number of parcels delivered. This enables us to compute key demand statistics, including:

- Median and Average Daily Parcel Volume
- Top 10% busiest days (highest demand periods)
- Lowest 10% least busy days (low-demand periods)

Parcel distribution of the daily demand (daily percentages combined with weekly) with high demand days in red, low demand days in green.

		Percentage	0,0713	0,1561	0,1085	0,2088	0,1824	0,1777	0,0952466
Weeknumber	Percentage	Number of P.	Sun	Mon	Tue	Wed	Thu	Fri	Sat
1	0,0106056	409	$\boxed{29}$	64	44	85	75	73	39
2	0,0240048	925	66	144	100	193	169	164	88
3	0,0210542	811	58	127	88	169	148	144	77
4	0,0194822	750	53	117	81	157	137	133	71
5	0,0193142	744	53	116	81	155	136	132	71
6	0,0197933	762	54	119	83	159	139	135	73
7	0,0186047	717	51	112	78	150	131	127	68
8	0,0171848	662	47	103	72	138	121	118	63
9	0,0168502	649	46	101	70	135	118	115	62
10	0,0187124	721	51	113	78	151	132	128	69
11	0,0182044	701	50	109	76	146	128	125	67
12	0,017989	693	49	108	75	145	126	123	66
13	0,014674	565	40	88	61	118	103	100	54
14	0,0157125	605	43	94	66	126	110	108	58
15	0,0181101	698	50	109	76	146	127	124	66
16	0,0190569	734	52	115	80	$\boxed{153}$	134	130	70
17	0,0174171	671	48	105	73	140	122	119	64
18	0,0179621	692	49	108	75	144	126	123	66
19	0,0169819	654	47	102	71	137	119	116	62
20	0,0214107	825	59	129	90	172	150	147	79
21	0,0204953	789	56	123	86	165	144	140	75
22	0,0187737	723	52	113	78	151	132	128	69
23	0,0184198	710	51	111	77	148	130	126	68
24	0,0202515	780	56	122	85	163	142	139	74
25	0,0189118	728	52	114	79	152	133	129	69
26	0,0191975	739	53	115	80	154	135	131	70
27	0,0261701	1008	72	157	109	210	184	179	96
28	0,0188271	725	52	113	79	151	132	129	69
29	0,0178918	689	49	108	75	144	126	122	66
30	0,0135382	521	37	81	57	109	95	93	50
31	0,0175797	677	48	106	73	141	123	120	64
32	0,0195131	752	54	117	82	157	137	134	72
33	0,0161936	624	44	97	68	130	114	111	59
34	0,015873	611	44	95	66	128	111	109	58
35	0,0175637	677	48	106	73	141	123	120	64
36	0,0114901	443	32	69	48	92	81	79	42
37	0,0164668	634	45	99	69	132	116	113	60
38	0,0159957	616	44	96	67	129	112	109	59
39	0,0156976	605	43	94	66	126	110	108	58
40	0,0167984	647	46	101	70	135	118	115	62
41	0,0182632	703	50	110	76	147	128	125	67
42	0,0183081	705	50	110	76	147	129	125	67
43	0,0185409	714	51	111	77	149	130	127	68
44	0,0185803	716	51	112	78	149	131	127	68
45	0,0187812	723	$\boxed{52}$	113	78	151	132	128	69
46	0,0200481	772	55	121	84	161	141	137	74
47	0,0213549	823	59	128	89	172	150	146	78
48	0,0225066	867	62	135	94	181	158	154	83
49	0,0336681	1297	92	202	141	271	237	231	124
50	0,0380141	1464	104	229	159	$\boxed{306}$	267	260	139
51	0,0335125	1291	92	202	140	270	235	229	123
52	0,0196482	757	54	118	82	158	138	135	72
		38520	2745	6011	4179	8039	7025	6842	3669
			0.0713	0.1561	0.1085	0.2088	0.1824	0.1777	0.095274

Appendix D) Vehicle Characteristics

The entire vehicle characteristics table is given here in the Appendix. Some information on load and unload time does not exist and can hardly be derived from data. These are estimated guesses by us, to try and make it as realistic as possible. The loading time per parcel for the S-ADV and R-ADV are guessed at 60 seconds, since that is what it probably would take back in the depot to load one parcel. Loading the UAV takes a bit longer because it makes use of a docking station. Unloading time of the UAV is also considered as the longest, since these UAVs make use of a droid that drops down to deliver the package. We assume that this will only happen once the customer is present at a drop-off location. The unload time of the R-ADV is assumed to be shorter than the S-ADV because the R-ADVs are designed for delivery to a single customer, which might increase the time needed for unlocking the S-ADV.

Vehicle	Characteristic	Value	Source or Extra Explanation
S-ADV	Speed	*3 km/h*	*STARSHIP DELIVERY ROBOT* (n.d.)
	Load time	*60 seconds/parcel*	Estimated guess
	Unload time	*120 seconds/parcel*	Estimated guess
	Max service range	*Operating radius 3.2km*	Starship Technologies (2024b)
	Storage capacity	*1 unit*	Starship Technologies (2024b)
	Noise	*0 (db)*	*Industry - Starship Deliveries* (2024)
	Battery Capacity	*1260 (Wh)*	*Industry - Starship Deliveries* (2024)
	Energy Usage	*35 (Wh/km)*	12 hours of driving 3 km/h -> 36 km max 1260 Wh/36 = 35 (Wh/km)
	Charge speed	*200 (W)*	(*STARSHIP DELIVERY ROBOT*, n.d.)
R-ADV	Speed	*15 km/h*	Max 25 (*Express Robot*, 2025), average around 15 km/h in urban area
	Load time	*60 sec per parcel*	Estimated guess
	Unload time	*60 seconds*	Estimated guess
	Max range	*80 km*	100 in perfect conditions (*Express Robot*, 2025), realistic 80
	Storage capacity	*24 units*	(*Express Robot*, 2025)
	Noise	*56 (db)*	Legal limit (*Waarom Een Elektrische Auto Geluid Maakt | ANWB*, n.d.)
	Battery Capacity	*7200 (Wh)*	*(Express Robot, 2025)*
	Energy Usage	*90 (Wh/km)*	7200Wh / 80km = 90 (Wh/km)
	Charge speed	*1440 (W)*	From 20% to 90% in 3.5 hours *(Express Robot, 2025).* (0,7*7200)/3,5 = 1440
UAV	Speed	*35 km/h*	Max 112 km/h (*Zipline Fact Sheet | Zipline UAV Delivery & Logistics*, n.d.). Estimated average 35 km/h in small urban area
	Load time	*120 seconds/parcel*	Estimated guess
	Unload time	*300 seconds/parcel*	Estimated guess
	Max service range	*16 km*	(*Zipline Fact Sheet | Zipline UAV Delivery & Logistics*, n.d.)
	Storage capacity	*1 unit*	(*Zipline Fact Sheet | Zipline UAV Delivery & Logistics*, n.d.)
	Noise	*60 (db)*	Youtube video (Marques Brownlee, 2025)
	Battery Capacity	*1333 (Wh)*	Battery of UAV 3 times as heavy 4000 Wh (*DJI DB2000 Intelligent Flight Battery*, n.d.). 4000/3 = 1333 Wh
	Energy Usage	*42 (Wh/km)*	Max fly range of 16*2 = 32 km
	Charge speed	*661 (W)*	Charging speed also 1/3 (*DJI DB2000 Intelligent Flight Battery*, n.d.)

Appendix E) Model Documentation

This Appendix includes all the extra interesting information on the simulation model. It starts with an overview of the four entities (Vehicle, Location, Route, Trip) and Action class.

Vehicle	
Properties	**Description**
Type	S-ADV, R-ADV, UAV or electric van
Load_Time	Load time of the vehicle (seconds)
Unload_Time	Unload time of the vehicle (seconds)
Speed	Average speed (km/h)
Noise_Pollution	Noise Pollution of the vehicle (dB)
Storage_Capacity	Storage Capacity of the vehicle (units)
Battery_Capacity	Battery Capacity of the vehicle (kWh)
Energy_Usage	Energy Usage of the vehicle (kWh/km)
Charge speed	Charge speed of the vehicle per hour (kW/h)

Location	
Properties	**Description**
Type	Warehouse or Customer Location
Georeference	Latitude and Longitude of the location
Time_Window_Open	Opening time of the service window
Time_Window_Close	Closing time of the service window
Demand	Demand for deliveries at this location
Service_Time	Time required for service at this location

Action	
Properties	**Description**
Type	Charge, Load, Unload, Wait, Move
Lifecycle	The lifecycle status of an action: "Requested,
Trip	The overall trip associated with the action
Route	The route associated with the action
Location	The location associated with the action
from	Starting location of the action (in case of move)
to	Destination of the action (in case of move)
Expected Duration	Expected Duration at the start of the day
Real Duration	Real Duration during the day

Trip	
Properties	**Description**
Locations	List of locations to visit
Vehicle	The vehicle assigned to the trip
Actions	List of all the actions included in the trip
Expected Duration	Expected Duration at the start of the day
Real Duration	Real Duration during the day

Route	
Properties	**Description**
Origin	Starting location of the route
Destination	Ending location of the route
Length	Length of the route
Coordinates	Coordinates

We will also go over the simulation model in a chronological sequence. Covering the manual mode first and then the simulation mode.

Input tab

The simulation model's "Input" tab is displayed below. The user can choose which vehicles will be used in the simulation, how many, and with what characteristics. Before customers can be generated, the input values should be determined.

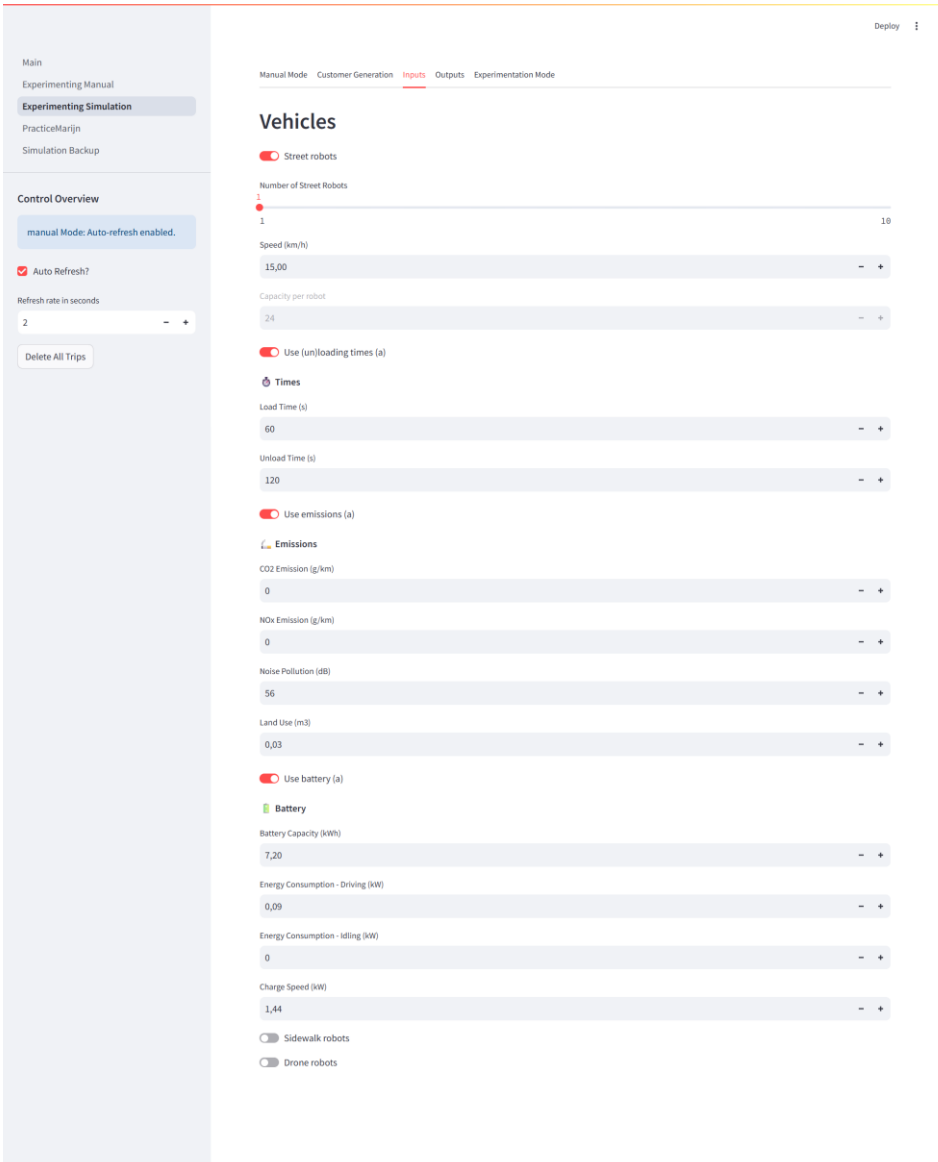

Customer Generation tab

The customers can be generated under the customer generation page once the user has entered the simulation model's correct input parameters. By entering the location Query (standard: Universiteit Twente), the relevant and available housing tags are displayed. The corresponding buildings will appear on the map below if specific tags are selected. Once the user is satisfied with the selection of tags/buildings, the customers can be generated. The user can choose the type of demand per customer, the number of customers, and whether there will be time windows using the sliders and check boxes. If all done, the user could simply generate the customers by pressing the button, and the map with the customers are shown on this screen.

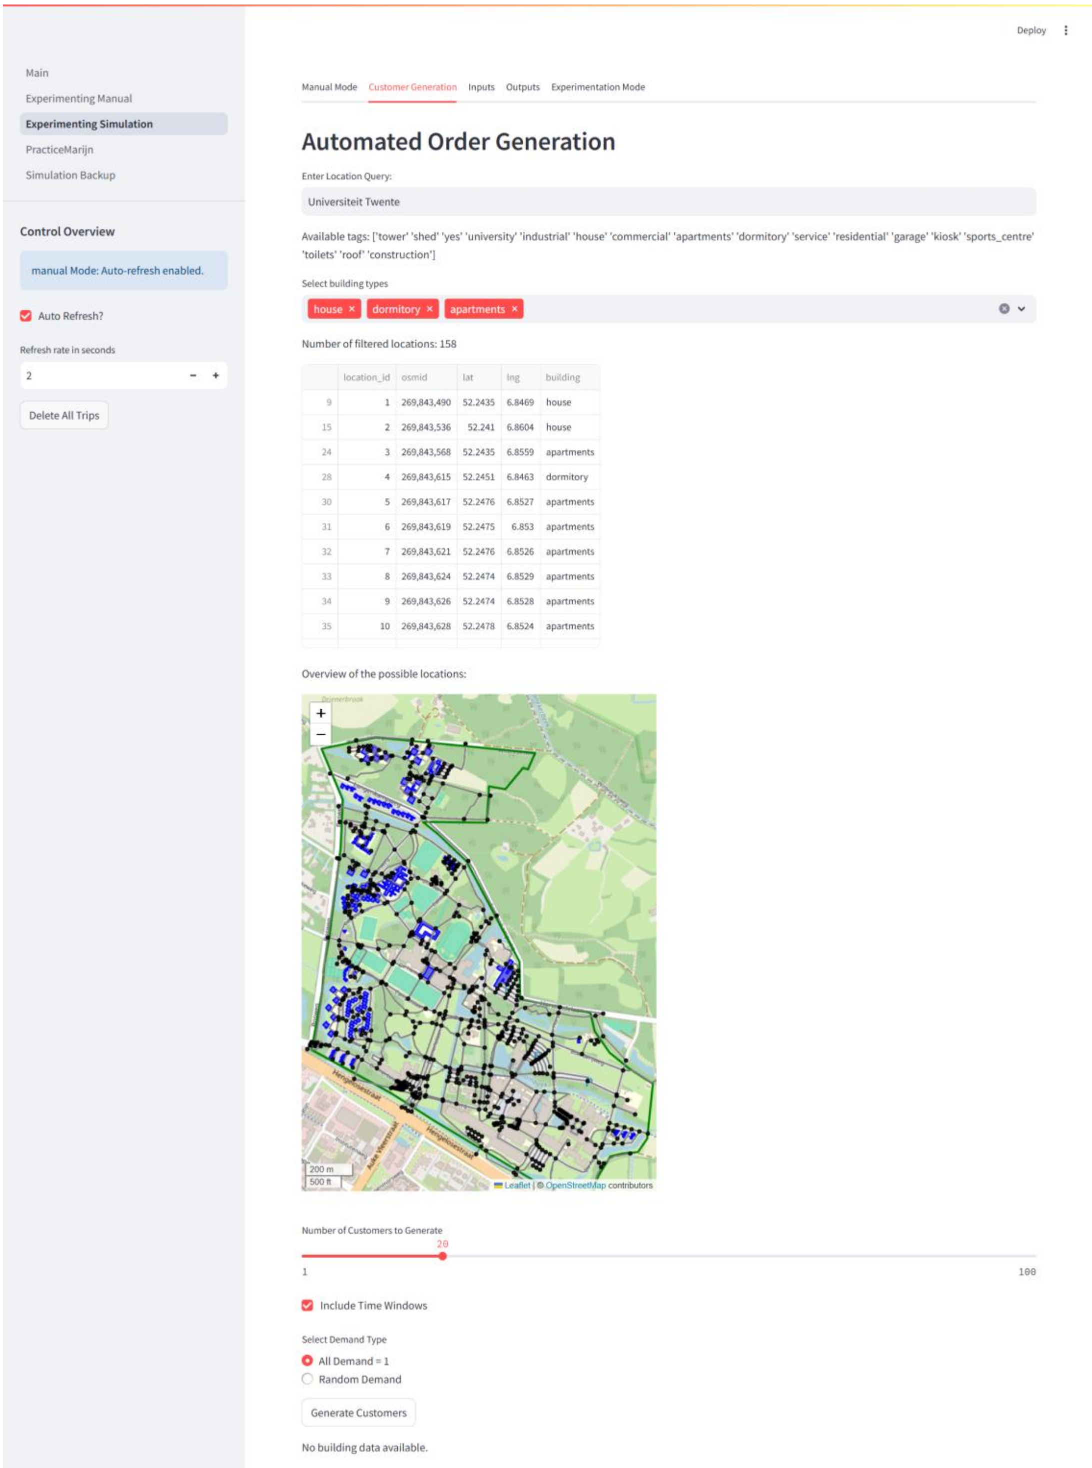

Manual Mode (Main tab)

Once the user has customized the simulation with their own inputs and generated the customers, the manual mode can be used. The user will see the customers' locations on the delivery map, and by pressing on the button *Allocate Vehicles to Customers*, the simulation will allocate the available vehicles to the customers (using the Solomons Nearest Neighbor allocation). This creates the 'Trips' of the model, which can be viewed under *Detailed Trip View*.

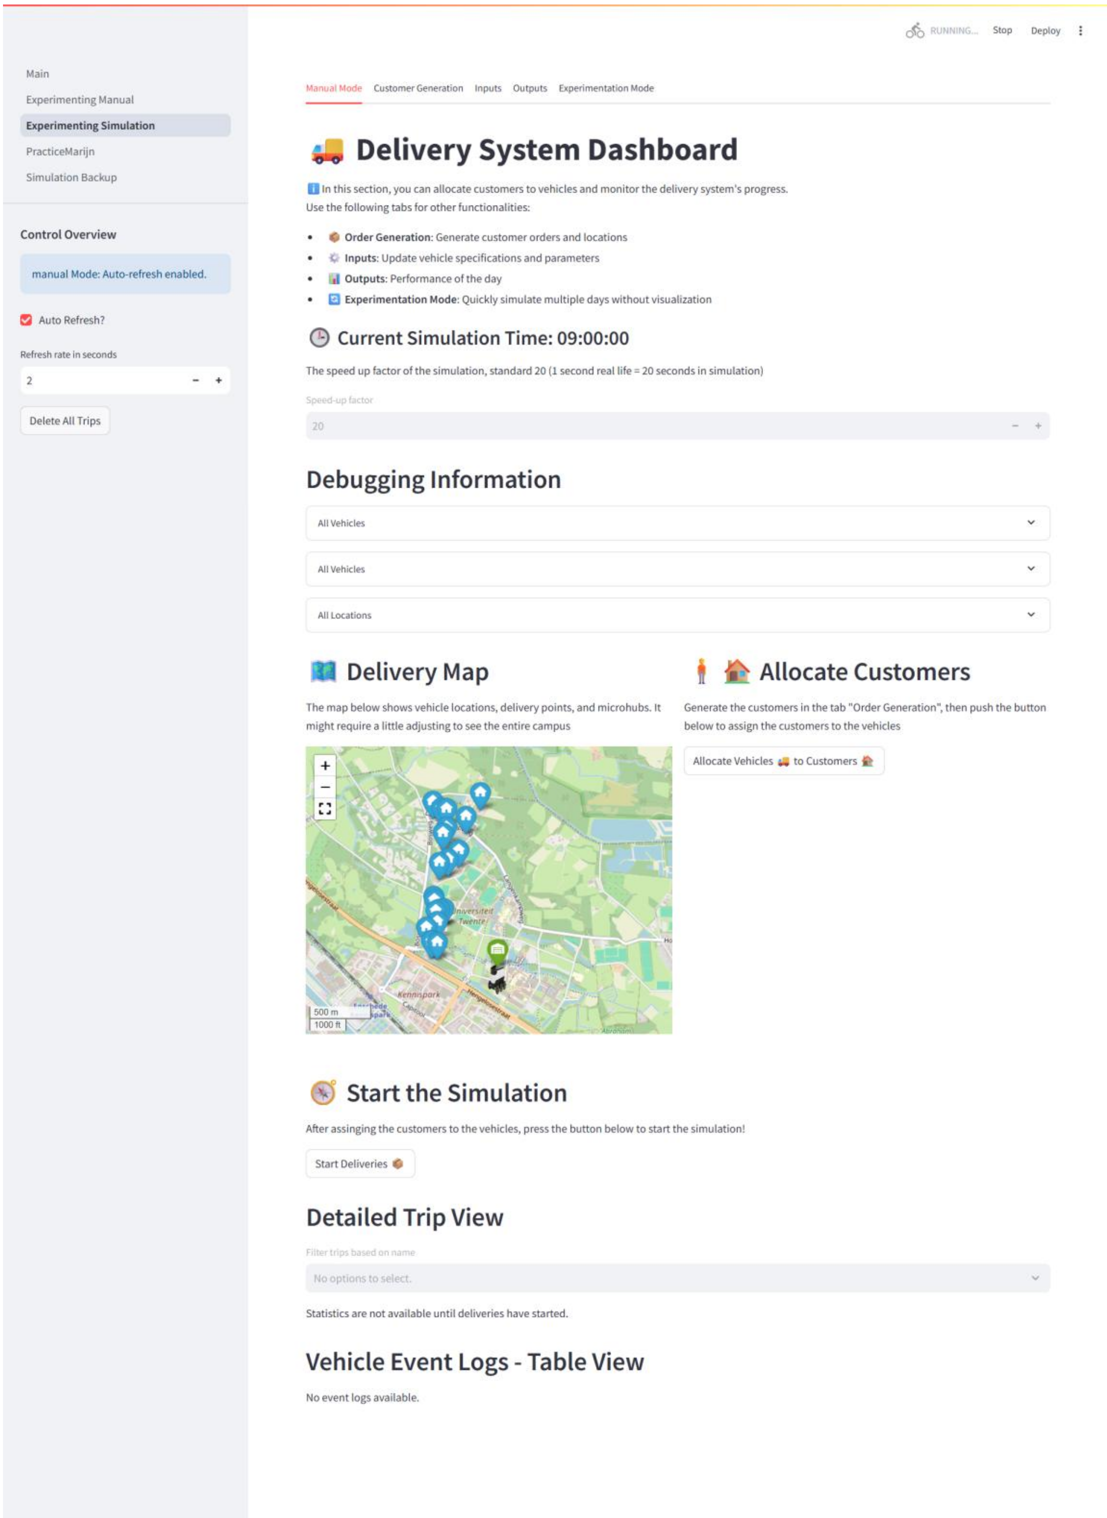

This creates the ‘Trips’ of the model, which can be viewed under *Detailed Trip View*. Here, the user can highlight a specific trip and see the actions and their expected durations. The route is also highlighted on the map itself. By pressing the button *Start Deliveries,* the trips will be started.

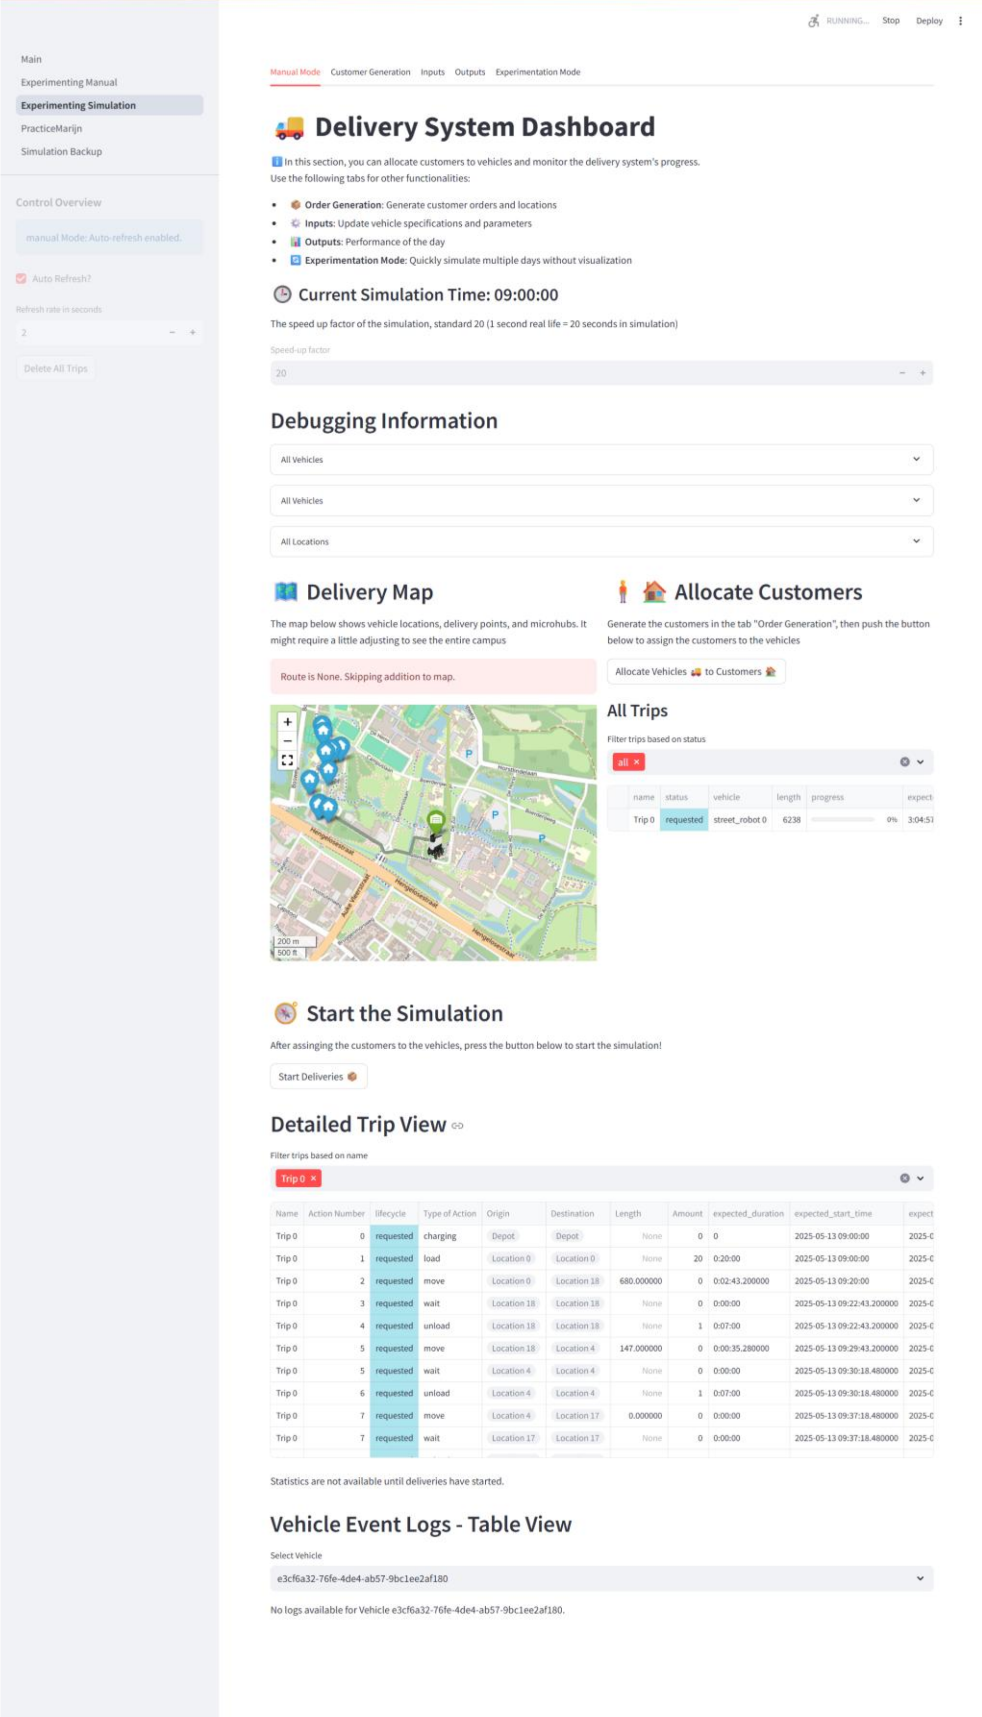

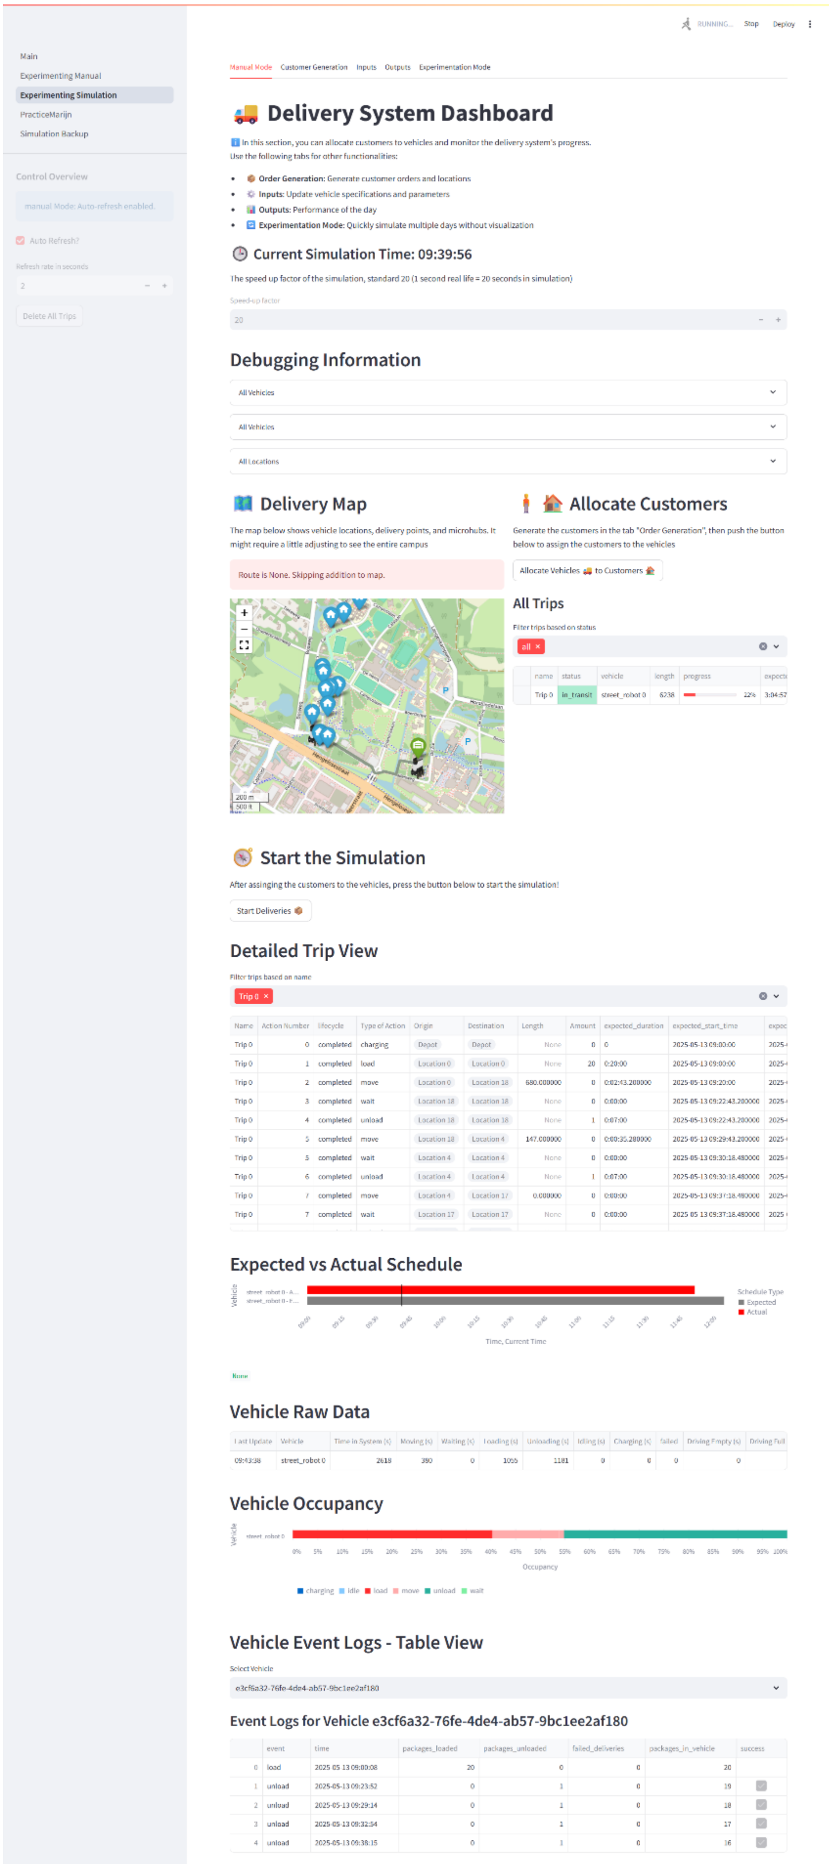

At periodic intervals, the location of the vehicles and the other information below are updated. The KPIs will appear on the *Outputs* tab at 17:00, at the end of the day.

Output tab

As said before, at 17:00 the KPIs will become visible in the *Output tab*.

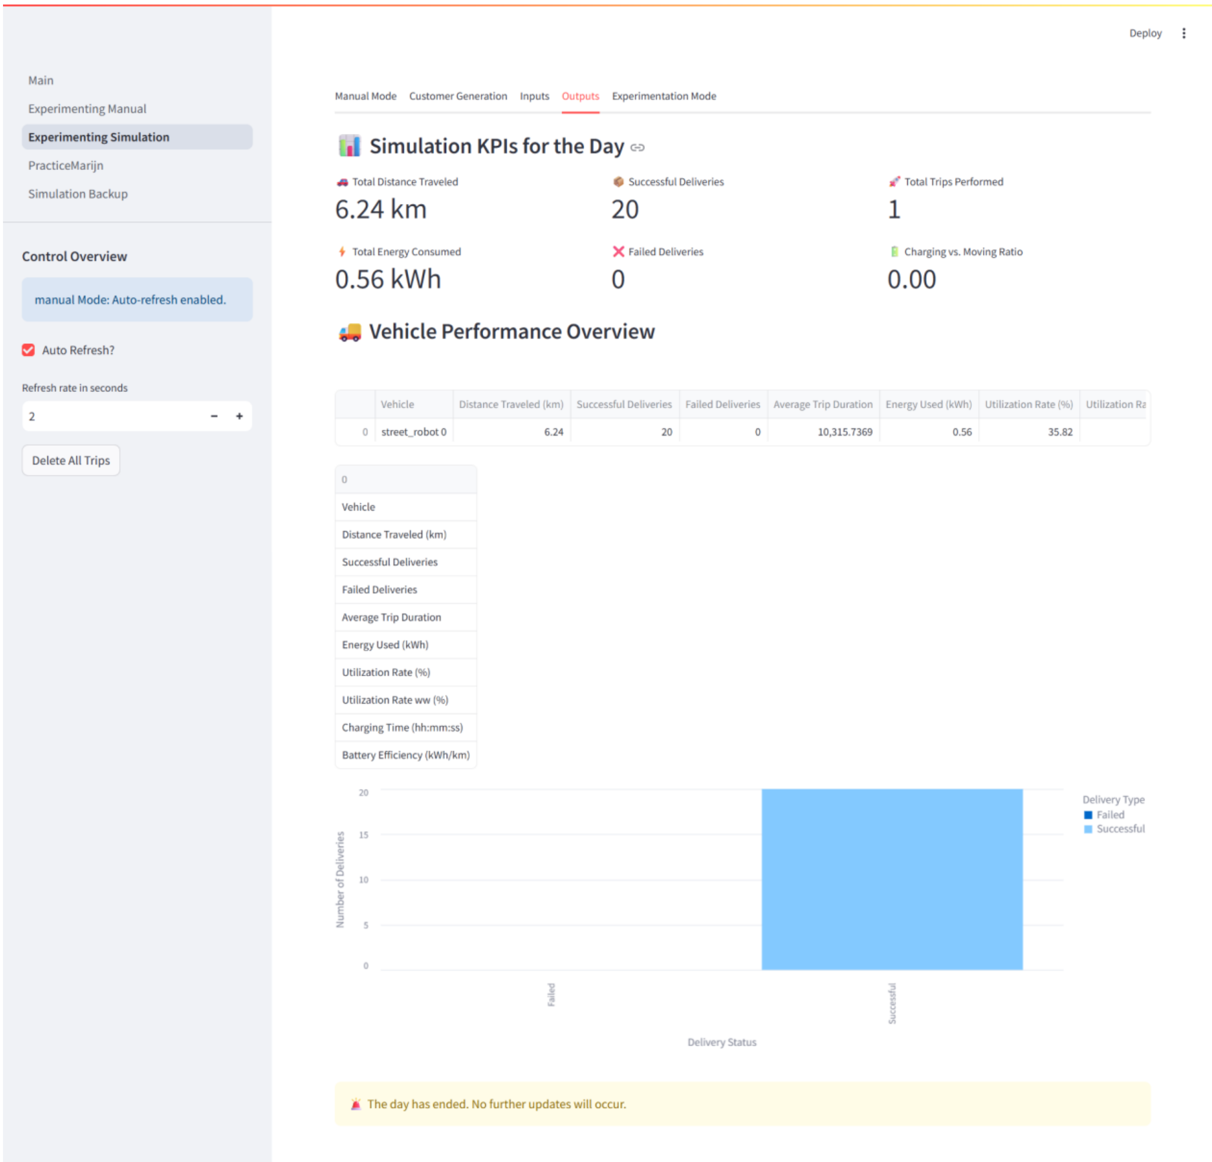

Appendix F) Detailed Experiments Section

In this Appendix, we will go over the simulation model in a chronological sequence. Covering the manual mode first and then the simulation mode

Phase 1

The main goal of phase 1 is to gain a rough understanding of the number of vehicles we need to cover the demand on a regular day (09:00-17:00), without considering time windows constraints. This phase gives us the basis we need for further experimentation. Due to the extensive number of experiments conducted, most detailed results and tables are provided in the Appendix. In phase 1, we focus exclusively on the key performance indicator **Successful Deliveries**, aiming to identify the minimum number of vehicles needed to fulfill all deliveries. Phase 1 is structured as follows: we first determine the number of required R-ADVs to cope with low, average and high demand, followed by the same analysis for S-ADVs and UAVs.

We start this phase by first analyzing the scenario Average demand. An initial experiment was conducted using a fleet size of 5, 10, 15, and 20 street Vehicles. The results, shown in Table 15, reveal that the key performance indicators (KPIs) were nearly identical across all configurations. All setups successfully completed all deliveries with no unplanned stops and minimal variation in delivery time or energy consumption. Notably, the configuration with only 5 street Vehicles already achieved full delivery coverage, with a utilization rate of 35.35%.

Exp	# of R-ADVs	Total Distance	Successful Deliveries	Trips	Util. (%)	Avg Delivery Time
1	5	19.11	106.0	5.0	26.64	00:06:01
2	10	19.13	106.0	5.0	13.46	00:06:05
3	15	18.86	106.0	5.0	8.72	00:05:55
4	20	19.42	106.0	5.0	6.51	00:05:53

Since the configuration with 5 Vehicles showed only ~35% utilization, we investigated whether even fewer street Vehicles could achieve full delivery capacity. A second experiment tested configurations with 1 to 5 Vehicles. The results below (Table 16)demonstrate that two Vehicles were sufficient to handle all regular-day deliveries.

Exp	# of R-ADVs	Total Distance	Successful + Failed Deliveries	Trips	Util. (%)	Avg Delivery Time
1	1	14.86	81.33	4.0	98.74	00:05:50
2	2	20.4	106.0	5.0	66.96	00:06:03
3	3	20.06	106.0	5.0	44.6	00:06:02
4	4	20.61	106.0	5.0	33.92	00:06:09
5	5	19.84	106.0	5.0	26.52	00:05:59

This indicates that two street Vehicles can comfortably meet daily delivery demands, and a single robot has a maximum delivery capacity of around 60 parcels (regular day, no time windows). Whether this holds on high-demand days are explored in the following section. Next, the experiments for low demand (Table 17) and high demand (Table 18) are shown in tables on the next page.

Exp	# of R-ADVs	Total Distance	Successful Deliveries	Trips	Util. (%)	Avg Delivery Time
1	1	11.66	43	2	53.19	00:05:56
2	2	11.8	43	2	28.42	00:06:23

Even a single street robot was able to handle all 43 deliveries comfortably, with a utilization of **73%**. Which, interestingly, does match what we saw earlier, that the street robot is able to deliver around 60 parcels a day (43/60 is around 73%).

Exp	# of R-ADVs	Total Distance	Successful Deliveries	Trips	Util. (%)	Avg Delivery Time
1	2	25.42	167	8	101.83	00:05:52
2	3	26.7	192	8	77.59	00:05:50
3	4	28.1	192	8	58.53	00:05:52

With 4 street Vehicles, all 192 deliveries were completed. At 3 Vehicles, some unplanned deliveries remained. Again, notice the fact that this is conform to the 60 packages a day.

Number of S-ADVs

Unlike street Vehicles, sidewalk robot configurations showed significant variation in performance depending on fleet size. This is expected, given their lower carrying capacity and reduced travel speed.

Exp	# of S-ADVs	Total Distance	Successful Deliveries	Trips	Util. (%)	Avg Delivery Time
1	5	97.28	66.0	66.0	100.43	00:36:42
2	10	204.0	102.33	102.33	100.18	00:47:01
3	15	209.08	106.0	106.0	68.0	00:46:11
4	20	209.2	106.0	106.0	51.17	00:46:20

From these results, it appears that a fleet of 10 S-ADVs is close to optimal, balancing capacity and utilization. To fine-tune this further, additional experiments were conducted with 10, 11, and 12 vehicles

Exp	# of S-ADVs	Total Distance	Successful Deliveries	Trips	Util. (%)	Avg Delivery Time
1	10	203.43	103	103.33	99.55	00:46:22
2	11	213.41	106.0	106.0	95.09	00:47:21
3	12	218.8	106	106	89.15	00:48:44

The results indicate that **11 sidewalk Vehicles** offer the best performance with near-complete deliveries, minimal unplanned stops, and excellent utilization.

S-ADVs (Low Demand)

Exp	# of sidewalk Vehicles	Total Distance	Successful Deliveries	Unplanned Deliveries	Energy (kWh)	Trips	Utili (%)	Avg Delivery Time
1	3	59.41	35.33	7.67	2.08	35.33	99.86	00:40:44
2	4	82.45	40.33	2.67	2.89	40.33	100.17	00:47:46
3	5	80.89	43.0	0	2.83	43.0	79.72	00:44:29

S-ADVs required 5 vehicles to ensure full delivery without failure, with a utilization of 83%. Lower numbers resulted in unplanned deliveries.

S-ADVs (High demand)

Exp	# of sidewalk Vehicles	Total Distance	Successful Deliveries	Unplanned Deliveries	Energy (kWh)	Trips	Utili (%)	Avg Delivery Time
1	15	301.86	167.67	24.33	10.57	167.67	99.85	00:43:07
2	16	324.97	170	22	11.37	170	100.24	00:45:27
3	17	347.22	179	13	12.15	179	100.69	00:46:02
4	18	365.8	189	3.0	12.8	189	100.02	00:45:53
5	19	384.95	189.67	2.33	13.47	189.67	98.89	00:47:43
6	20	394.65	192	0	13.81	192	95.94	00:48:02
7	21	394.55	192	0	13.81	192	91.76	00:48:14

With 20 sidewalk Vehicles, all 192 deliveries were able to be completed. The utilization of 95.57 shows us that there is not a lot of room left for more parcels.

Number of UAVs (Average demand)

For the UAV Vehicles, the initial experiment using 5 to 20 UAVs revealed little variation in KPIs. Even 5 UAVs could complete all deliveries efficiently, suggesting overcapacity in larger fleet sizes.

Exp	# of UAVs	Total Distance	Successful Deliveries	Trips	Util. (%)	Avg Delivery Time
9	5	161.02	106.0	106.0	56.16	00:12:42
10	10	164.15	106.0	106.0	25.61	00:11:35
11	15	160.57	106.0	106.0	17.33	00:11:50
12	20	156.66	106.0	106.0	12.85	00:11:38

To determine the **minimum viable fleet size**, experiments were also run with 1 to 5 UAVs.

Exp	# of UAVs	Total Distance	Successful Deliveries	Trips	Util. (%)	Avg Delivery Time
1	1	38.4	41	41	100.34	00:11:50
2	2	89.51	74	74	100.45	00:13:01
3	3	150.56	102	102	101.74	00:14:21
4	4	161.38	106	106	76.43	00:15:52
5	5	162.07	106	106	57.31	00:12:58

The data shows that **at least 4 UAVs** are required to ensure full delivery capacity. Although 3 UAVs cover a large portion, they fall just short.

UAV Vehicles (Low demand)

Exp	# of UAVs	Total Distance	Successful Deliveries	Trips	Util. (%)	Avg Delivery Time
1	1	48.34	35.67	35.67	104.61	00:14:07
2	2	61.99	43	43	55.24	00:12:519

Two UAVs were required to meet low demand fully. One UAV left about **11 deliveries unfulfilled**.

UAVs (High Demand)

Exp	# of UAVs	Total Distance	Successful Deliveries	Trips	Util. (%)	Avg Delivery Time
1	5	242.18	175.33	175.33	101.31	00:13:53
2	6	294.45	192	192	96.39	00:14:27
3	7	289.07	192	192	78.37	00:13:44

Seven UAVs were required to fulfill all deliveries at high demand, though delivery times rose significantly.

Conclusion Phase 1

The same type of experiments was conducted with the S-ADVs and the UAVs. All these experiments and their outcomes can be found in Appendix X. The results are shown in Table 19.

	Number of Vehicles necessary to fulfill entire demand		
Vehicle	**Low (43)**	**Average (106)**	**High (192)**
S-ADVs	5	11	20
R-ADVs	1	2	4
UAVs	2	4	6

Based on our own observations, we can state that the number of vehicles necessary to fulfill the entire demand of a regular day (09:00-17:00, no time windows) have a linear growth. For the specific conditions of our research, we can formulate the following equation:

$$\#\ of\ Vehicles\ needed\ for\ regular\ day \approx 60 * \alpha + 10 * \beta + 30 * \gamma$$

Where:

$\alpha = Number\ of\ R - ADVs$

$\beta = Number\ of\ S - ADVs$

$\gamma = Number\ of\ Drones$

In Phase 2, we will test and try this equation by creating heterogeneous fleets.

Phase 2

Now, we are interested in the heterogeneous fleet scenario, where a different combination of vehicles is allowed to deliver parcels. From this point on, we begin using the KPIs discussed earlier, since a lot of combinations will satisfy the demand of the parcels. We do not consider time windows in this phase, so social KPIs will not be measured. Again, more detailed information on the experiments is given in the Appendix. To estimate sufficient heterogeneous combinations, we will apply two guiding principles:

- Avoid Redundant Capacity: Combinations in which a single vehicle type already meets the entire demand on its own are excluded. This helps focus on true mixed-fleet strategies.
- Use of estimation formula: The equation derived in Phase 1 is used as a reference to guide the search for effective fleet configurations:

$$Demand_{Regular\ Day} \approx 80 * \alpha + 10 * \beta + 35 * \gamma$$

- In general, we use the following table for creating our heterogeneous combinations:

S-ADVs	R-ADVs	UAV
Only S-ADVs	-	-
-	Only R-ADVs	-
-	-	Only UAVs
High S-ADVs	Low R-ADVs	-
Low S-ADVs	High R-ADVS	-
-	Low R-ADVs	High UAVs
-	High R-ADVs	Low UAVs
High S-ADVs	-	Low UAVs
Low S-ADVs	-	High UAVs

Low Demand Performance (Homogeneous & Heterogeneous, Regular Day, no TW)

Combo ID	Demand	S-ADV	R-ADV	UAV	Reps	Del	Tot. Energy	Util.	Distance	Del. time
1	Low (43)	5	0	0	10	43	2.98	83.45	85.26	46:34
2	Low (43)	0	1	0	10	43	1.06	54.92	11.74	06:07
3	Low (43)	0	0	2	10	43	2.66	56.1	63.25	12:32
4	Low (43)	2	0	1	10	43	2.84	78.53	73.01	26:21
5	Low (43)	1	0	1	10	42.8	2.77	97.7	69.34	22:01

In the low demand scenario, the heterogeneous combinations are limited by the small number of vehicles available, especially R-ADVs (only 1) and UAV (only 2). As such, there is only one useful heterogeneous combination (Combo 4) involving both S-ADVs and UAVs. This configuration matches the demand capacity closely (2*10 + 1*35 ≈ 55, 43/55 ≈ 78%) and performs relatively well across the KPIs.

Average Demand Performance (Homogeneous & Heterogeneous, Regular Day, no TW)

The results for heterogeneous fleet with average demand are shown in Table 21:

Combo ID	Demand	S-ADV	R-ADV	UAV	Del.	Tot. Energy	Util.	Distance	Del. time
1	Average (106)	11	0	0	106	7.54	95.37	215.37	47:30
2	Average (106)	0	2	0	106	1.86	67.14	20.66	06:04
3	Average (106)	0	0	4	106	6.87	75.82	163.74	13:46
4	Average (106)	3	1	0	106	3.53	96.44	75.12	17:31
5	Average (106)	0	1	1	106	3.3	91.22	60.99	08:18
6	Average (106)	8	0	1	106	6.99	89.89	190.96	36:44

7	Average (106)	1	0	3	106	6.86	97.71	166.83	17:44

With the average demand, we observe that a lot more options seem to be viable due to the increase in available ADVs. The utilization rates indicate that these options are the most 'efficient' ones. Since adding an S-ADV or UAV to a certain combination would mean that they essentially have overcapacity.

Example Calculation using demand estimation formula

Let's verify combo 6 using the demand formula:

$$Demand_{Regular\ Day} \approx 80 * 1 + 10 * 0 + 35 * 1 \approx 120$$

Thus, Combo 6 has a total capacity of 120, giving it +-8% overcapacity (106/120) compared to actual demand.

Phase 3

Now, we are interested in the performance of the system with night delivery (09:00-21:00) instead of normal hours (09:00-17:00). Just like the previous phase, we start with the low demand case, continue with the average demand case and finally the high demand case. Since we have a longer delivery time, the expected performance is that we need fewer vehicles for each demand scenario.

Low Demand Performance (Homogeneous, Night Delivery, No TW)

The night delivery scenario (See Table 24) creates the possibility to deliver every parcel by either one UAV, one R-ADV or four S-ADVs. These values also show that a heterogeneous fleet is not possible/necessary for the low demand scenario (combinations will always lead to overcapacity).

Combo ID	Demand	S-ADV	R-ADV	UAV	Del	Late cust	Tot. Energy	Util.	Distance	Del. time
1	Low (43)	5	0	0	43	0	2.91	54.67	83.21	45:45
2	Low (43)	4	0	0	43	0	3.09	72.09	88.18	48:16
3	Low (43)	3	0	0	42.78	0.3	3.12	97.99	89.01	49:36
3	Low (43)	0	0	2	43	0	2.76	38.67	65.61	13:02
4	Low (43)	0	0	1	43	0	2.76	90.44	65.75	15:20
5	Low (43)	0	1	0	43	0	1.08	37.47	11.95	06:18

Average Demand Performance (Homogeneous and heterogeneous, Night Delivery, No TW)

The night delivery scenario with average demand (See Table 25) creates opportunities for some heterogeneous fleets. In this night delivery scenario, one R-ADV is still able to deliver all the packages on its own. Because of this, there is no efficient heterogeneous fleet combination with R-ADVs. This leaves only combinations with S-ADVs and UAVs.

Combo ID	Demand	S-ADV	R-ADV	UAV	Del	Late cust	Tot. Energy	Util.	Distance	Del. time
~~1~~	~~Avg (106)~~	~~9~~	~~0~~	~~0~~	~~106~~	~~0~~	~~7.53~~	~~77.72~~	~~215~~	~~47:41~~
2	Avg (106)	8	0	0	106	0	7.54	87.5	215.34	47:41
3	Avg (106)	0	0	3	106	0	6.81	72.5	162.21	14:46
~~3~~	~~Avg (106)~~	~~0~~	~~0~~	~~2~~	~~97.7~~	~~1.2~~	~~5.85~~	~~101.15~~	~~139.33~~	~~14:57~~
~~4~~	~~Avg (106)~~	~~0~~	~~2~~	~~0~~	~~106~~	~~0~~	~~1.81~~	~~44.47~~	~~20.14~~	~~06:02~~
5	Avg (106)	0	1	0	106	0	1.8	88.8	20	06:01
		~~6~~	~~0~~	~~2~~	~~106~~	~~0~~	~~7.09~~	~~55.2~~	~~186.15~~	~~30:02~~
		~~3~~	~~0~~	~~1~~	~~98~~	~~1.2~~	~~6.08~~	~~102.33~~	~~159.84~~	~~30:11~~
		~~5~~	~~0~~	~~2~~	~~106~~	~~0~~	~~7.01~~	~~59.67~~	~~182.45~~	~~28:31~~
		2	**0**	**2**	**106**	**0**	**6.86**	**83.87**	**172.01**	**22:49**
		4	~~0~~	~~1~~	~~105.9~~	~~0.9~~	~~6.98~~	~~98.32~~	~~186.03~~	~~33:28~~
		5	**0**	**1**	**106**	**0**	**7.08**	**84.76**	**190.37**	**34:35**

High Demand Performance (Homogeneous and heterogeneous, Night Delivery, No TW)

In the high demand scenario, it is no longer possible for the single R-ADV to deliver each parcel. Two R-ADVs are now necessary to meet the demand of 192 parcels, with a utilization rate of 80%. The S-ADVs do not scale well, even in the night delivery scenario. Instead of the 20 S-ADVs we needed, we now still need 15 of them. The R-ADVs can now also be considered in the heterogeneous fleet combinations. Besides the combinations between two ADVs, this high demand scenario also has one heterogeneous fleet combination with all three ADVs.

Combo ID	Demand	S-ADV	R-ADV	UAV	Del.	Tot. Energy	Util.	Distance	Del. time
1	High (192)	15	0	0	192	13.5	84.03	385.76	47:22
2	High (192)	0	2	0	192	2.45	77.61	27.27	05:49
3	High (192)	0	0	5	192	12.26	79.52	291.98	14:58
4	~~High (192)~~	~~8~~	~~1~~	~~0~~	~~192~~	~~7.98~~	~~79.03~~	~~208.13~~	~~26:44~~
~~5~~	~~High (192)~~	~~0~~	~~1~~	~~3~~	~~192~~	~~6.92~~	~~65.13~~	~~149.46~~	~~09:47~~
~~6~~	~~High (192)~~	4	~~1~~	~~2~~	~~192~~	~~7.52~~	~~61:49~~	~~174.96~~	~~16:08~~
7	High (192)	~~9~~	~~0~~	~~2~~	~~192~~	~~13.02~~	~~83.1~~	~~348.44~~	~~34:20~~
8	High (192)	~~3~~	~~0~~	4	~~192~~	~~12.41~~	~~79.75~~	~~307.82~~	~~20:58~~
9	High (192)	**7**	**1**	**0**	**192**	**7.65**	**82.35**	**191.84**	**24:48**
10	High (192)	**0**	**1**	**2**	**192**	**6.11**	**82.91**	**126.52**	**09:20**
11	High (192)	**8**	**0**	**2**	**192**	**12.6**	**87.44**	**335.96**	**32:49**
12	High (192)	**2**	**0**	**4**	**192**	**12.5**	**87.18**	**306.61**	**19:41**
13	High (192)	**2**	**1**	**1**	**192**	**5.65**	**88.33**	**125.49**	**13:17**

Conclusion Phase 3

From Phase 3, we can conclude that the following number of ADVs is necessary to meet the demand of parcels:

	Number of Vehicles necessary to fulfill entire demand		
Vehicle	**Low (43)**	**Average (106)**	**High (192)**
S-ADVs	4	8	15
R-ADVs	1	1	2
UAVs	1	2	5

From the experiments, we can still observe some linear relationships between the number of vehicles and the demand level. Where the R-ADV and UAV really improve in the potential number of parcel deliveries on a day-to-day basis, the S-ADV does not really show the improvement you would want in the night scenario. Just like in Phase 1, we formulate another number of vehicles demand equation.

$$Demand_{Night\ Delivery} \approx 120 * \alpha + 12 * \beta + 50 * \gamma$$

Where:

- $\alpha = Number\ of\ R - ADVs$
- $\beta = Number\ of\ S - ADVs$
- $\gamma = Number\ of\ UAVs$

Example Calculation using demand estimation formula

Let's verify *Average demand (2 S-ADVS, 0 R-ADV, 2 UAV)* using the demand formula:

$$Demand_{Regular\ Day} \approx 120 * \ 0 + 14 * \ 2 + \ 50 * 2 \ \approx 128$$

This combination had a utilization of 83.5% according to the experiments, 106/124 results in 82.8%.

Let's verify *High demand (2 S-ADVS, 1 R-ADV, 1 UAV)* using the demand formula:

$$Demand_{Regular\ Day} \approx 120 * \ 1 + 14 * \ 2 + \ 50 * 1 \ \approx 198$$

This combination had a utilization of 88.3% according to the experiments, 192/198 results in 96.9%.

So, the equation works quite well in the average case and underestimates the performance in the high demand case.

Phase 4

For the final phase, we will check the compatibility of the autonomous delivery system with customer specific time windows. Which means that the customers can choose their own preferred time window. For example, on a 09:00-17:00 day, customers can select morning delivery 09:00-13:00 or afternoon delivery 13:00-17:00. One of the most interesting KPIs right now are the number of late customers (and the total minutes being late). Also, we are interested in the KPI *Average time after the start of TimeWindow* which, as the name suggests, is the average amount of time customers must wait after the start of the TimeWindow. A lower *Average Time After Start of TimeWindow* is preferred, since customers won't have to wait as long. Before we start experimenting, we will have to take another look at our Solomons nearest neighbor's cost function. This cost function, as we have used (and talked about) before, has the following weights:

$$C_{i,j} = w_1(0.025) * d_{i,j} + w_2(0.025) * T_{i,j} + (0.95)w_3 * \max\left(0, v_{i,j}\right) + PenaltyCosts * \max\left(0, -v_{i,j}\right)$$

The reason that we chose for this distribution of weights now becomes useful. Earlier, we only had to deal with distance and time. Where now, we must deal with TimeWindows. So, why would we choose for weights distributed like this? Please look at the following example:

Low Demand Performance (Homogeneous and heterogeneous, Regular Day, TW)

From this point on, the important information to save in the tables becomes a bit much, so we use some abbreviations in the tables (LM = Total Late Minutes, LC = Average Late Customers, ATAT = Average Time After TimeWindows). In the low demand scenario, it takes 8 S-ADVs and 4 Drones to deliver to each customer on time in their TimeWindow. As one can see, the R-ADV scenario does not get to the point of achieving zero late customers. Due to our Solomons Nearest Neighbor algorithm, the R-ADVs are being filled with the 'best' possible customers. So, in the lower demand scenario, 43 parcels must be delivered to the customers, which can be done in two trips (2*24 = 48). Because the trips now take a bit longer (since we prioritize time windows instead of distance), it is no longer possible for 1 R-ADV to meet the demand (41.47 deliveries). Two R-ADVs are thus necessary to meet the demand. The first R-ADV is filled with the best 24 customers, after which the following R-ADV is filled with the other 19 customers. In the unfortunate case of having a lot of morning deliveries relative to afternoon deliveries (for example 26/17) in combination with (morning) customers located all around the campus, the R-ADVs are not able to deliver to each customer in the morning timeslot. This means that on certain days, the morning deliveries are postponed to late in the afternoon (after the afternoon deliveries). So even though the Late Customers KPI is not equal to zero, it is a bit of a distorted view, because if on 2/10 days 4 customers are late, it will tell us 0.8 customers on average late (10 runs).

ID	S-ADV	R-ADV	UAV	Del.	Dist.	LM	LC	Energy	Util.	ATAT	Del. time
1	5	0	0	42	83.21	260	1.4	2.91	95.94	96.32	54:59
2	6	0	0	42.87	82.52	278	1.93	2.89	88.01	95.82	59:08
3	7	0	0	43	87.63	137	0.87	2.98	81.24	79.9	01:03:35
4	8	0	0	43	88.83	0	0	3.09	75.37	69.55	01:07:18
5	0	1	0	41.47	15.12	1546	8.93	1.36	90.48	151.7	10:32
6	0	2	0	43	15.7	186.3	3	1.41	64.81	79.49	14:30
7	0	3	0	43	16.12	76.62	1.07	1.45	40.75	68.71	13:40
8	0	4	0	43	15.7	67.89	1.53	1.41	31.32	70.24	14:01
9	0	0	2	42.66	64.69	540.8	3.67	2.72	87.17	106.9	19:47
10	0	0	3	43	64.93	10.76	0.47	2.73	68.85	56.34	23:03
11	0	0	4	43	65.05	0	0	2.73	62.68	37.36	28:03

Average Demand Performance (Homogeneous and heterogeneous, Regular Day, TW)

In the average demand scenario, the number of S-ADVs necessary to deliver the parcels to get zero late customers is 18, and the number of UAVs is 9. The R-ADVs do not improve after 5, since they drive a total of 5 trips which is enough to deliver each parcel (5*24 = 120, 120>106).

ID	D	S-ADV	R-ADV	UAV	Dist.	LM	LC	Energy	Util.	ATAT	Del. time
1	106	17	0	0	210	102.54	1.1	7.37	78.02	76.94	01:00:23
2	106	18	0	0	218.82	0	0	7.66	77.48	71.25	01:03:11
3	106	0	5	0	26.84	114.43	2.5	2.42	37.91	72.29	12:53
~~4~~	~~106~~	~~0~~	~~6~~	~~0~~	~~29.31~~	~~864.58~~	~~14~~	~~2.64~~	~~48.26~~	~~99.24~~	~~13:07~~
~~5~~	~~106~~	~~0~~	~~7~~	~~0~~	~~29.74~~	~~740.86~~	~~12.4~~	~~2.68~~	~~40.24~~	~~95.42~~	~~12:44~~
~~6~~	~~106~~	~~0~~	~~8~~	~~0~~	~~29.76~~	~~701.27~~	~~13~~	~~2.68~~	~~36.7~~	~~96.51~~	~~13:20~~
~~7~~	~~106~~	~~0~~	~~9~~	~~0~~	~~29.86~~	~~665.61~~	~~12.2~~	~~2.69~~	~~33.39~~	~~94.46~~	~~13:41~~
8	106	0	0	7	162.45	55.02	1.1	6.82	71.01	64.11	22:35
9	106	0	0	8	164.42	1.64	0.2	6.9	66.28	50.79	24:08
10	106	0	0	9	163.14	0	0	6.85	63.79	41.21	26:02

High Demand Performance (Homogeneous and heterogeneous, Regular Day, TW)

In the high demand scenario, just like in the low and average scenario, the R-ADVs do not improve after 8, since they drive a total of 8 trips which is enough to deliver each parcel (8*24 = 192). We observe that the 30 S-ADVS are sufficient for time window delivery without late customers, as well as 11 drones.

ID	D	S-ADV	R-ADV	UAV	Dist.	LM	LC	Energy	Util.	ATAT	Del. time
1	192	25	0	0	383.87	1103	6	13.44	60.98	98.27	55:30
2	192	28	0	0	381.16	442.83	2.5	13.34	55.34	83.29	58:14
3	192	29	0	0	389.25	135.94	1	13.87	53.3	78.05	01:00:09
4	192	30	0	0	382.94	0	0	13.44	53.57	74.52	01:00:11
5	192	31	0	0	381.38	0	0	13.35	52.69	74.24	01:01:30
6	192	0	8	0	35.79	41.49	0	3.22	36.88	66.18	11:05
7	192	0	0	15	288.01	0	0	12.09	43.22	44.43	24:27
8	192	0	0	14	294.02	0	0	12.35	44.52	48.08	23:24
9	192	0	0	13	292.09	0	0	12.27	45.88	60.32	22:21
10	192	0	0	12	292.06	0	0	12.27	47.58	58.5	21:36
11	192	0	0	11	288.68	0	0	12.12	48.92	55.63	20:11
12	192	0	0	10	290.86	320.44	2.33	12.22	51.39	70.96	19:18

Conclusion Phase 4

From Phase 4, we can conclude that the following number of ADVs is necessary to meet the demand of parcels:

	Number of Vehicles necessary to achieve 0 late customers		
Vehicle	**Low (43)**	**Average (106)**	**High (192)**
S-ADVs	4	18	30
R-ADVs	**2**	**5**	**8**
UAVs	1	9	11