UiT The Arctic University of Norway

Faculty of Science and Technology
Department of Physics and Technology

Attention-Based Neural Network for Solving the Green Vehicle Routing Problem in Waste Management

Marit Utheim Breimo

EOM-3901 Master's Thesis in Energy, Climate and Environment - June 2023

Abstract

The transport sector is a major contributor to the emission of greenhouse gases and air pollution. As urbanization and population growth continue to increase, the demand for transportation services grows, emphasizing the need for sustainable practices. Therefore, incorporating sustainability into the transport sector can effectively reduce its negative impacts on the environment and optimize the utilization of resources.

This thesis aims to address this issue by proposing a novel method that integrates neural networks into the development of a green vehicle routing model. By incorporating environmental considerations, particularly fuel consumption, into the optimization process, the model seeks to generate more sustainable route solutions. The integration of machine learning techniques, specifically an attention-based neural network, demonstrates the potential of combining machine learning with operations research for effective route optimization.

While the effectiveness of the green vehicle routing problem (GVRP) has been demonstrated in providing sustainable routes, its practical applications in real-world scenarios are still limited. Therefore, this thesis proposes the implementation of the GVRP model in a real-world waste collection routing problem. The study utilizes data obtained from Remiks, a waste management company responsible for waste collection and handling in Tromsø and Karlsøy.

The findings of this study highlight the promising synergy between machine learning and operations research for further advancements and real-world applications. Specifically, the application of the GVRP approach to waste management issues has been shown to reduce emissions during the waste collection process compared to routes optimized solely for distance minimization. The attention-based neural network approach successfully generates routes that minimize fuel consumption, outperforming distance-optimized routes. These results underscore the importance of leveraging the GVRP to address environmental challenges while enhancing decision-making efficiency and effectiveness. Overall, this thesis provides insights for developing sustainable and optimized routes for real-world problems.

Acknowledgements

Throughout the work on this master's thesis and during my studies, I am incredibly grateful for all the support I have received.

First and foremost, I would like to express my heartfelt thanks to my supervisors, Michael Kampffmeyer and Chiara Bordin, for their guidance and expertise throughout the project. Thank you, Michael, for your vast knowledge and precise guidance throughout the entire project. Thank you, Chiara, for your expertise and kindness. I appreciate their belief in the project, which allowed us to bring this idea to reality.

I am also thankful to my fellow student, Astrid, for her valuable contributions to this project. The collaboration and discussions we had emphasized the importance of teamwork in our academic journey.

I want to thank Remiks for providing insights into their operations and granting access to data, as well as the drivers who allowed me to accompany them on waste collection routes.

I want to express my gratitude to my family for their unwavering support. To my fellow students in the master's office, thank you for the engaging discussions. Special thanks to my friend Silje, who has been a companion in the field of sciences from the very beginning. Your support has been invaluable. Finally, I want to thank my partner, Peter, for his constant encouragement throughout my studies.

Marit Utheim Breimo
Tromsø, June 2023.

Contents

List of Figures

Abbreviations

OR	Operations Research
LP	Linear Programming
ILP	Integer Linear Programming
TSP	Travelling Salesman Problem
VRP	Vehicle Routing Problem
CVRP	Capacitated Vehicle Routing Problem
ACO	Ant Colony Optimization
GVRP	Green Vehicle Routing Problem
TBL	Triple Bottom Line
PRP	Pollution Routing Problem
ML	Machine Learning
RNN	Recurrent Neural Network
GNN	Graph Neural Network
GAT	Graph Attention Network

1

Introduction

1.1 Emissions From the Transport Sector

Globally, the transport sector is crucial, as it not only provides accessibility, mobility, and social connectivity but also promotes economic growth by connecting people, communities, and businesses, thereby facilitating communication and commerce. Unfortunately, the industry is a significant contributor to the massive amounts of carbon emissions that occur in the world today.

Transportation accounts for about one-quarter of the world's energy-related carbon dioxide emissions, and this development is projected to increase in the future [1]. Road transportation, mainly commercial and industrial vehicles such as delivery trucks, is a significant source of greenhouse gas emissions [2]. These vehicles are essential to our daily lives, but they rely heavily on traditional combustion engines, which result in high fuel consumption and pollution. To meet the Paris Agreement's objective of limiting global warming to 2°C, exploring more sustainable transportation solutions is essential.

It is undeniable that we must transition to a more sustainable transport sector. In recent years, the word "sustainability" has lost some meaning and has become overwhelming for many [3]. So, the problem is where to begin and what is required to accomplish this objective. A solid starting point is to concentrate on the repetitive daily transportation processes.

The waste management industry is a repetitive, everyday emitter due to the emissions caused by waste collection processes. Adopting electric or biofuel-powered collection vehicles, promoting waste reduction and recycling to reduce the quantity of waste requiring collection and disposal, and optimizing waste collection routes to reduce fuel consumption can reduce these emissions. This thesis will emphasize the latter approach, as it allows organizations to gain a competitive advantage without incurring additional costs while reducing pollution and fuel costs.

1.2 Operations Research

Operations Research (OR) is a promising field for optimizing waste collection routes and reducing fuel consumption. The field of OR has become increasingly important in recent years due to the rapid pace of technological change, globalization, and environmental challenges. Utilizing mathematical and quantitative techniques, OR analyzes complex problems, develops models, and identifies solutions that enhance operational performance and decision-making across various industries [4]. Therefore, this field has the potential to help solve some of the challenges related to the transportation sector. Allocating resources optimally is crucial for organizations to remain competitive and efficient, especially in complex and rapidly changing environments. The power of OR is exemplified by its transformative impact on industries such as airline scheduling [5] and energy power planning [6], which have been transformed by the application of OR and mathematical optimization.

In sustainable transport, OR can be applied to various areas such as transportation planning, logistics, and supply chain management. For instance, OR techniques can analyze traffic patterns and optimize public transportation routes to reduce emissions, energy consumption, and travel time. It can also be used to develop more sustainable freight transportation networks, for example, by minimizing the number of empty truck journeys and increasing the use of rail or water transport, which can result in significant carbon emission reductions.

The Vehicle Routing Problem (VRP) is a classic optimization challenge in OR that aims to find the most efficient set of routes for a fleet of vehicles to serve a set of locations [7]. Applying OR techniques to optimize routing decisions offers cost savings and emission reductions, providing organizations with a competitive edge. VRPs have applications in diverse sectors such as transportation, logistics, healthcare, and waste management. In this thesis, we investigate VRPs potential for determining optimized waste collection routes.

1.3 Machine Learning

In the current era of technology and digitalization, the amount of data being produced is growing at an exponential rate, and it has become critical for businesses to make sense of this data to gain a competitive edge. And it is in this context that the power of machine learning becomes particularly interesting. Machine learning is a sub-field of artificial intelligence. It involves algorithms and models that allow computers to learn from data and improve their performance on specific tasks without being explicitly programmed to do so [8]. These algorithms can help organizations make sense of their data and extract valuable insights, which can be used to improve business operations, develop new products and services, and make better decisions.

One of the exciting developments in the field of operations research is the integration of machine learning techniques to solve complex optimization problems [9] [10][11] [12]. Although this symbiosis has been explored since 1985 [13], it has only recently gained significant attention and interest. Over the years, a growing interest has been in incorporating machine learning techniques into the field of OR. Different approaches have been explored, including attention-based machine learning techniques [14] [9], supervised learning models [15] [10], and reinforcement learning approaches [16] [17].

This thesis specifically focuses on the application of reinforcement learning as a promising approach for solving routing problems, offering the advantage of not requiring a training dataset. Reinforcement learning techniques have demonstrated great potential in optimizing complex optimization problems by leveraging trial and error learning and adapting strategies based on environmental feedback. By harnessing the power of reinforcement learning, we aim to develop versatile routing solutions that can be applied to various problem domains, as this approach does not limit the possibilities due to the lack of a training dataset.

1.4 Key challenges

This thesis primarily focuses on the Green Vehicle Routing Problem (GVRP), a variant of the VRP that aims to minimize emissions while optimizing transportation operations. In the context of the growing emphasis on sustainability in the transport sector, the GVRP has gained attention due to its potential to reduce emissions and promote sustainable practices. Nonetheless, the GVRP presents several challenges that require consideration to develop effective and practical solutions.

The definition of the GVRP is a subject of ongoing research, accompanied by several challenges. An essential consideration is how to effectively incorporate environmental concerns into the problem formulation. Researchers have explored various approaches, such as introducing penalties or constraints, treating it as the primary objective, or employing a multi-objective optimization framework that considers factors like cost or time. The choice of formulation depends on the specific context, goals, and constraints of the problem. Striving to strike a balance between environmental considerations and other factors remains an active research area, aiming to develop robust models that accurately capture and address environmental factors.

The applicability of the GVRP to real-world applications presents a challenge due to the difference between theoretical models and practical considerations. For instance, most studies in the literature simplify the problem by assuming Euclidean distance instead of road distance, which does not fully capture the complexities and constraints of real transportation networks. To overcome this challenge, improving the applicability of GVRP models to real-world scenarios and enhancing their effectiveness in supporting practical decision-making processes is important. Addressing these challenges is crucial for successfully integrating sustainable transportation practices into real-world applications.

1.5 Contribution

This thesis explores the symbiotic relationship between machine learning and operations research, specifically in the context of the GVRP. The primary contributions of this research can be summarized as follows:

Firstly, methodologically, we propose a novel method that tailors the attention model introduced by Kool et al. [9] specifically for the GVRP. Our approach involves adapting the model to address the requirements and constraints of the GVRP, by incorporating fuel consumption as the environmental consideration. By integrating fuel consumption into a multi-objective optimization framework that simultaneously considers road gradient and distance, our objective is to minimize emissions and promote sustainable transport practices.

Secondly, analytically, we validate the practical applicability of the proposed attention-based method by applying it to a real-world scenario. Going beyond theoretical concepts, we conduct a case study that addresses the real-world GVRP faced by Remiks, a waste management company. By optimizing Remiks' waste collection routes using our proposed method, we showcase the model's ability to deliver practical solutions that effectively address the specific needs and operational constraints of the company.

Moreover, this research encompasses extensive data gathering and manipulation, ensuring the availability of reliable and accurate information for analysis. Furthermore, close communication and collaboration with the relevant company were maintained throughout the research process, ensuring the alignment of objectives and the study's relevance.

1.6 Outline

This thesis consists of nine chapters. Chapter 1 provides an introduction to the topics and the work in this thesis. The following chapters are divided into three parts: background, method and data, and results and experiments.

Background:
Chapter 2 covers operations research and its relevance to optimization models. Chapter 3 explores the green vehicle routing problem and sustainability in transportation. Chapter 4 focuses on machine learning, specifically neural networks. The existing research on neural networks in the context of operations research is reviewed, highlighting promising approaches that can be applied to optimization models. Chapter 5 delves into attention mechanisms and their recent application to routing problems.

Method & Data:
Chapter 6 presents the proposed methodology for solving waste management's green vehicle routing problem using an attention-based neural network. Chapter 7 presents a real-world application of this method to a specific real-world scenario involving the waste company Remiks, highlighting how the model can generate more sustainable waste collection routes. The chapter, therefore, consists of the required data to do this.

Results & Experiments:
Chapter 8 shows the results and analysis of the proposed method. Chapter 9 includes a discussion, future directions, and a conclusion.

An appendix with plots and results is also included.

Part I

Background

2

Operations Research

Operations Research (OR) is a broad, interdisciplinary field that employs a variety of mathematical and quantitative techniques to analyze and optimize complex decision problems. It emerged during World War II, when the efficient allocation of scarce resources became crucial [18]. In today's rapidly changing and resource-constrained business landscape, the importance of OR has only grown. Its focus on using analytical techniques to improve decision-making is a powerful asset for addressing real-world issues in a world that is increasingly reliant on data-driven solutions.

Whether in finance, business management, or any other industry, effectively allocating resources remains a significant challenge. In such scenarios, human intuition is often insufficient, and mathematical models are frequently required to discover optimal solutions. Consequently, OR provides a robust framework for optimizing complex decision-making processes.

This section of the master's thesis draws heavily on the project thesis [19], which provided a comprehensive examination of the field of OR. It offers an overview of the fundamental theories and key concepts that are central to OR. This section also covers optimization tasks such as the Travelling Salesman Problem and the Vehicle Routing Problem.

2.1 An Introduction to Operations Research

Operations Research (OR) is a discipline that aims to improve decision-making by employing suitable analytic techniques [4]. It focuses on solving strategic problems in real-world operations to achieve optimal results. Multiple disciplines, including business management, finance, decision support systems, information technology, and tourism management, are surrounded by OR [20]. It has revolutionized the way industries like airline scheduling and energy power planning approach problem-solving by leveraging mathematical optimization and OR techniques [5] [6].

When solving a real-world problem as an operational researcher, it is essential first to understand the problem in its entirety and analyze the available data [4]. The next step is to create a suitable mathematical model that simplifies the problem while still incorporating its key features. The solution derived from the mathematical model must be interpreted and validated to ensure its relevance to the real problem.

Model validation is critical, as it ensures a more practical and applicable solution by evaluating and verifying the result [21]. This is important because the simplified version of the problem must retain its essential features. The process of evaluating and validating the result is crucial, particularly when costly and irreversible decisions must be made. Finally, the best possible course of action is determined through mathematical optimization algorithms, which provide an optimal solution to the decision-maker. Overall, operations research aims to solve real-world problems in a way that benefits the entire organization through the search for optimality.

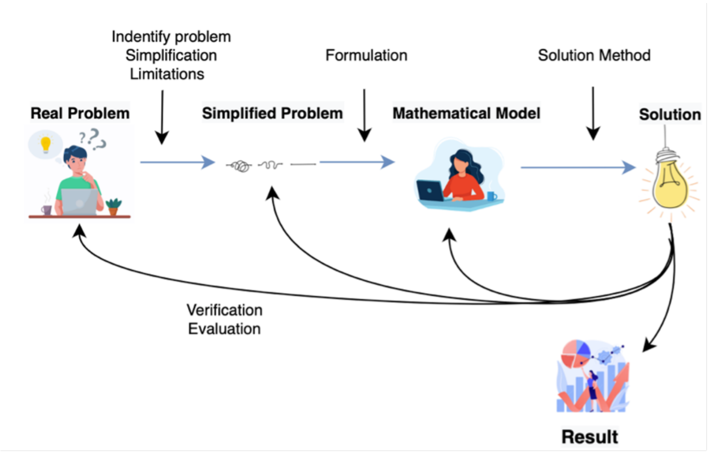

Figure 2.1: This figure illustrates a simplified process of mathematical modeling, which involves problem formulation, model constructing, model solving, and interpreting results. It also emphasizes model validation.

2.2 Mathematical Optimization

Operations research uses mathematical models and techniques to optimize business processes, improve efficiency, and reduce costs. Mathematical optimization is a key concept in OR, involving finding the best solution by determining optimal values for variables within constraints [22]. It provides a rigorous framework for solving problems, allowing OR to identify optimal solutions despite limited resources and competing goals. By using mathematical optimization, OR can provide insight into complex real-world problems and inform decision-making.

Mathematical optimization involves constructing a mathematical model that integrates relevant factors and constraints and then utilizing available data to identify the optimal solution. This can involve minimizing or maximizing an objective function. An objective function expresses the goals to be achieved and what the mathematical model will optimize according to. For instance, the objective function could be to minimize costs or risk or maximize profit or performance. The time horizon for decision-making can vary from short-term, intermediate-term, or long-term, depending on the problem and context.

In addition to the objective function, mathematical optimization models have two other main components: decision variables and constraints [23]. Constraints refer to a set of conditions that restrict the set of feasible solutions to a given problem. These conditions specify the limitations and requirements that must be satisfied by the solution of the optimization problem. Constraints can be of different types, including equality constraints, inequality constraints, and bound constraints. Manpower, resources, and operation time are all examples of constraints because they limit the capacity of a system or process.

A decision variable is a changeable quantity used to optimize a system or process. It is crucial to find the optimal solution for tasks such as product production, inventory management, or resource allocation. In energy optimization, decision variables can represent the power output of generators, while in airline scheduling, they can represent flight departure and arrival times. A general minimization problem can be expressed as:

$$min \quad f(x), \qquad x \in X, \tag{2.1}$$

where $f(x)$ represents the objective function, which is a mathematical function that measures the quality of a particular solution to the optimization problem. The objective function depends on the decision variable x, which represents the values that we want to optimize. X is a set of feasible values for x, as following:

$$g_i(x) \leq b_i, \qquad i = 1, ..., m. \tag{2.2}$$

$g(x)$ is a function that is dependent on the decision variable x and, together with the parameter b_i, forms a constraint. A solution that satisfies the constraints is called a feasible solution. In other words, a feasible solution is a set of values for the decision variables that meet all of the requirements and limitations of the optimization problem. A solution that also optimizes the objective function is called an optimal solution. The optimal solution satisfies $x \in X$ that minimizes $f(x)$ and is typically denoted by x^*. Thus, the optimal objective function value is given by $f(x^*) = z^*$.

2.3 Linear Programming

Linear Programming (LP) is a technique used to optimize a linear objective function subject to linear constraints. It is a type of mathematical optimization used widely in OR for modeling and solving complex problems in a simplified manner. Given that multiple solution algorithms are based on LP, it is advantageous first to establish LP concepts to lay the groundwork for more complex methods. Terminology and notation are used further in the context of more complex solutions methods.

In LP problems, both the objective function and all the constraint functions are linearly expressed. In addition, all the variables are continuous. The benefit of using an LP approach is that a local optimum will also be the global optimum. A local minimum is the point at which the value of a function is less than the neighborhood values. A global minimum, on the other hand, is the lowest point of a function across its entire domain.

The goal of LP is to attain the optimal solution to a given problem. The results of LP can guide a decision-maker in deciding how to employ the available resources in the most effective way. In this way, LP techniques are used to improve the quality of decision-making and provide companies with big competitive advantages [4]. The general linear programming problem [4] can be formulated with the following objective function:

$$min \quad z = \sum_{j=1}^{n} c_j x_j. \tag{2.3}$$

z is the objective function value, which is the quantity that we want to maximize or minimize. c_j represents the coefficients of the decision variables. These determine the weighting or the importance of each decision variable in the objective function. The decision variables x_j are the unknown quantities that we want to determine in order to optimize the objective function. And the objective is here to minimize z.

The general linear programming problem can be formulated with the following set of constraints:

$$\sum_{j=1}^{n} a_{ij}x_j \leq b_i, \quad i = 1, ..., m \;\; and \;\; x_j \geq 0, \quad j = 1, ..., n. \tag{2.4}$$

In optimization problems, the constraint coefficients $a_i j$ and the right-hand side coefficients b_i provide crucial information for identifying feasible solutions. In particular, the former indicates the amount of resource i required for each unit of activity j, while the latter represents the total quantity of resource i available for all activities. A feasible solution can be obtained by ensuring that the solution satisfies all of the constraints in a problem [4]. Conversely, an unfeasible solution violates at least one constraint.

The feasible region contains all feasible solutions, and the optimal solution is a feasible solution that produces the most favorable objective function value. In a maximization problem, the optimal solution corresponds to the maximum value of the objective function, such as maximum profit or maximum budget impact, whereas in a minimization problem, it corresponds to the minimum value, such as minimizing cost or emissions [4].

Typically, the influence of changing the parameters in the LP formulation is analyzed. This technique for evaluating the robustness of a solution is known as sensitivity analysis [24]. This can help to identify critical variables that should be carefully monitored or controlled to ensure the optimality of the solution. Sensitivity analysis is an important tool for decision-making, as it helps to understand the implications of different scenarios and make more informed decisions.

Overall, linear programming is a highly versatile optimization technique that can be applied to a wide range of problems across various fields. For example, it can be used to solve routing problems like the Travelling Salesman Problem and the Vehicle Routing Problem, providing valuable insights and solutions to complex real-world problems. By applying linear programming, decision-makers can optimize their operations and resources, leading to improved efficiency, cost savings, and enhanced competitiveness.

2.4 Graph Theory

Graph theory is a powerful branch of mathematics that deals with the study of graphs, which are structures made up of vertices or nodes connected by edges or arcs. It has numerous real-world applications in fields such as operations research, mathematical optimization, and LP. In OR, graph theory is used to model and analyze complex systems like transportation networks, communication networks, and supply chains. It can optimize the routing of vehicles, scheduling of tasks, or the allocation of resources, among other things.

A graph is composed of nodes and edges that connect node pairings [25]. Nodes and edges are typically represented as points or circles with lines or curves connecting them. Figure 2.2 illustrates that each edge connecting two elements represents a relationship or connection between them. A graph is represented by a square matrix known as an adjacency matrix. The matrix's dimensions correspond to the number of vertices in the graph, and the entry in row i and column j indicates whether there is an edge between nodes i and j. This matrix is an indispensable instrument for graph analysis and manipulation.

Subgraphs are smaller graphs that exist within a larger graph. A subgraph is created by selecting a subset of the vertices and edges from the original graph. Subgraphs are useful for analyzing the structure of a graph and for identifying patterns or substructures within the graph. Node degrees represent the number of edges that are connected to a node in a graph. In undirected graphs, the degree of a node is simply the number of edges connected to it. In directed graphs, node degrees can be split into indegree and outdegree, representing the number of incoming and outgoing edges, respectively.

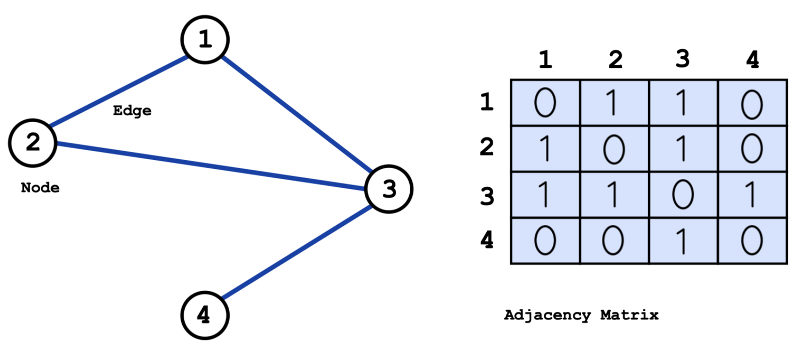

Figure 2.2: This figure illustrates an undirected graph with nodes represented by circles and edges represented by lines connecting the circles. Additionally, the corresponding adjacency matrix, a square matrix that represents the connections between nodes, is also depicted.

2.5 Travelling Salesman Problem

The Travelling Salesman Problem (TSP) is a well-known optimization problem in the field of OR. The problem involves finding the shortest possible route that visits a given list of cities exactly once and returns to the starting city, given the distances between the cities [26]. While the problem is simple to state, it is computationally difficult to solve, especially for large instances. Despite this challenge, the TSP has numerous applications in areas such as logistics, transportation, and network optimization. As a result, researchers have developed a variety of solution methods to tackle the TSP, making it a popular area of study in the field of optimization.

The original problem that the TSP aims to solve is to find the shortest possible route for a salesman to visit a given set of cities and return to the starting city while visiting each city once. One of the key constraints of the problem is that it prohibits sub-tours, meaning that the route must not be broken into multiple trips that go back to the starting point [26]. However, extensions can be made to the problem to enable a wider application. Figure 2.3 provides an illustrative example of a solution route for the TSP.

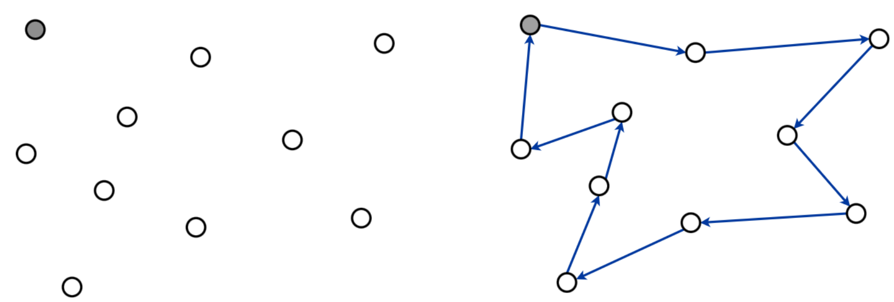

Figure 2.3: The solution illustrated in this figure for the Traveling Salesman Problem visits each city once and then returns to the starting point. The optimization problem determines the shortest route for a salesman to visit several cities and return to the starting point (illustrated in grey).

Problem Statement
The problem can be formulated as a graph, where each node n represents a city, and the binary variable x_{ij} takes on a value of 1 if the corresponding edge is used and 0 otherwise. An edge is considered used if the salesman chooses to travel along it. The cost of traveling from city i to city j is denoted as c_{ij} and represents the distances or costs of traveling from one city to another.

The TSP can be modified to include additional constraints, such as limited capacity or time constraints. However, the objective is to find the shortest possible route that the salesman can take. The objective function can then be defined as the sum of the costs of all the edges in the route as follows:

$$min \quad \sum_{i=1}^{n} \sum_{j\neq i, j=1}^{n} c_{ij} x_{ij}. \tag{2.5}$$

The objective function is optimized subject to the following constraints:

$$\sum_{i=1, i\neq j}^{n} x_{ij} = 1, \qquad j = 1, ..., n, \tag{2.6}$$

$$\sum_{j=1, j\neq i}^{n} x_{ij} = 1, \qquad i = 1, ..., n, \tag{2.7}$$

$$\sum_{i\in S}^{n} \sum_{j\neq i, j\in S}^{n} x_{ij} \leq |S| - 1, \qquad S \subset \{1, ..., n\}, |S| \geq 2. \tag{2.8}$$

The first constraint requires that each city is visited exactly once. This constraint ensures that the salesman must travel to every city in the problem exactly once and cannot skip any cities or visit any city more than once.

The second constraint ensures that the salesman leaves and returns to the starting city. This constraint requires the salesman to return to the starting city at the end of the tour. This ensures that the solution is a complete cycle rather than a path that ends somewhere in the middle of the route.

The third constraint prevents the use of sub-tours. This constraint prevents the salesman from taking shortcuts or revisiting cities already visited in a way that would break the overall cycle. In other words, the solution must be a single, continuous tour through all the cities, with no backtracking or skipping of cities. Additional constraints can be added to fit specific applications, such as limited capacity or time constraints.

The TSP received considerable attention due to its simple conceptual framework, which contrasts significantly with the difficulty of actually solving these problems. They are NP-hard problems, where NP means nondeterministic polynomial time and describes the complexity of such problems [27]. This means that as the size of the problem increases, so does the time required to solve it. This makes the TSP a difficult one to solve, especially as the number of cities to be visited grows.

Despite the difficulty of the problem, various solution methodologies exist that provide approximate solutions with a trade-off between accuracy and computational time. The simplicity of the TSP makes it a more theoretical problem than practical for real-world application. However, it is often used as a starting point or a benchmark for many optimization methods.

2.6 Vehicle Routing Problem

The Vehicle Routing Problem (VRP) is an extension of the TSP. While the TSP involves finding the optimal route for a single salesman to visit a set of cities, VRP is a combinatorial optimization problem that focuses on determining the optimal set of routes for a fleet of vehicles to visit a set of customers with varying demands, starting and ending at a depot [28] [29]. The origins of this problem date back to 1959 when Dantzig and Ramser published the first research paper on the topic [30]. In 1976, Christofides established the term VRP [31]. The main objective of the problem is to minimize the total distance traveled by all vehicles while fulfilling all the constraints.

The VRP arises in various real-world applications such as logistics, transportation, distribution, and service operations. For example, in the logistics industry, companies need to deliver goods to their customers while minimizing transportation costs and ensuring timely delivery. In such cases, the VRP can help in determining the optimal set of routes for delivery trucks.

The VRP is one of the most widely studied topics in the OR field [32]. It is an important and fundamental problem in the domain of transport and logistics. Solving the VRP is challenging because it belongs to the class of NP-hard problems, which means that finding the exact solution for large instances is not computationally feasible. As the number of customers increases, the number of possible routes increases tremendously.

In addition, complexity is increased when adapting the models to particular problems by including constraints. For a logistics company with several vehicles and many customers to visit, the number of routes to choose between becomes overwhelming. In such cases, VRP can provide powerful decision-making support to guide a manager to choose the route that provides, for example, maximum profit.

2.6.1 Capacitated Vehicle Routing Problem

The VRP is available in numerous variants, each with its own constraints and optimization objectives due to the wide range of operating principles and constraints encountered in real-world applications [33]. The Capacitated Vehicle Routing Problem (CVRP), Time-Dependent VRP (TDVRP), Multi-Depot VRP (MDVRP), and Stochastic VRP (SVRP) are among the most common variants [7]. The CVRP is a highly studied and popular variant of the classic VRP formulation, which incorporates capacity constraints. In this problem, goods or services are transported to customers using vehicles with limited capacity. The goal is to minimize the total distance traveled while satisfying multiple constraints, including vehicle capacity.

The practical applications of CVRP have attracted substantial interest from both academics and industry specialists. For instance, package delivery companies use CVRP to optimize their delivery routes, resulting in lower transportation costs and increased productivity. Likewise, waste collection companies utilize the CVRP to optimize their waste collection routes, leading to a more streamlined waste management process[34]. Thus, the CVRP can be used as an important tool in enhancing operational efficiency and reducing costs for businesses across industries.

Although the CVRP is commonly used as a benchmark problem in academic research, it also has practical applications in various industries, such as transportation, logistics, and manufacturing. Furthermore, it serves as an excellent starting point and guide, providing insight into the current state of affairs in the field of vehicle routing and transportation optimization. The CVRP problem statement is commonly used as a baseline when structuring a range of vehicle routing problems. By customizing the problem through the addition of constraints specific to the application, the CVRP can be adapted to various real-world scenarios, making it a valuable tool for optimizing routing and transportation solutions.

Problem Statement
The CVRP is a mathematical optimization problem with the goal of minimizing the total distance a fleet of vehicles travels. Specific mathematical formulations have been developed for the CVRP. Nonetheless, these formulations can be modified and expanded better to explain other VRP variants.

The problem statement below defines the customers $N = 1, 2, ..., n$ who need to be visited by a fleet of vehicles in the CVRP. Each customer has a demand, denoted by $q_i \geq 0$. If the vehicles in the fleet are similar in terms of capacity, we refer to them as homogeneous. Each vehicle has a capacity of $Q > 0$, and they operate at identical costs in this case.

When a vehicle travels from one customer to another (from i to j), it incurs a travel cost of c_{ij}. A binary variable x_{ij} is used to indicate whether the edge is used ($x_{ij} = 1$) or not ($x_{ij} = 0$). The objective function is formulated in a similar way as in the TSP, as follows:

$$min \quad \sum_{i=1}^{n} \sum_{j \neq i, j=1}^{n} c_{ij} x_{ij}. \tag{2.9}$$

The objective function is optimized subject to the following constraints:

$$\sum_{i=1, i \neq j}^{n} x_{ij} = 1, \qquad j = 1, ..., n, \tag{2.10}$$

$$\sum_{j=1, j \neq i}^{n} x_{ij} = 1, \qquad i = 1, ..., n, \tag{2.11}$$

$$\sum_{i \in S}^{n} \sum_{j \neq i, j \in S}^{n} x_{ij} \leq |S| - 1, \qquad S \subset \{1, ..., n\}, |S| \geq 2. \tag{2.12}$$

The first set of equations in this problem statement closely resembles the equations used in the TSP. Equation 2.10 ensures that each node is visited exactly once, while equation 2.11 requires the vehicle to leave and return to the depot. To avoid sub-tours, equation 2.12 is used.

In addition to the equations described earlier, the CVRP can include a capacity constraint that is related to the capacity of the vehicle. To enforce this constraint, the following equation is utilized:

$$\sum_{j=1}^{n} q_j x_{ij} \leq Q_i, \quad \forall_i \in V, i \neq 0. \tag{2.13}$$

It states that the total demand q_j of node j, visited by a vehicle starting from node i cannot exceed the capacity of the vehicle, Q_i. This equation ensures that the vehicle does not exceed its maximum capacity when servicing the customers.

The problem statement typically starts with the essential constraints derived from the TSP, which can be further expanded to accommodate specific problem requirements and constraints. For example, if there is a predetermined number of vehicles, the following constraint can be incorporated using the provided formulation:

$$\sum_{j=1, j \neq i}^{n} x_{oj} = |K|. \tag{2.14}$$

However, the addition of multiple constraints may result in no feasible solutions if there exists no solution that satisfies all of them.

In conclusion, the CVRP plays a crucial role in the field of transportation and logistics, with practical applications in various industries. By implementing suitable algorithms, businesses can optimize their routing strategies, leading to reduced transportation costs, enhanced efficiency, and improved overall customer satisfaction.

2.7 Solution Methodologies

As stated previously, the VRP is an NP-hard problem, which is computationally difficult to solve for large instances. As the number of customers or locations increases, the problem becomes increasingly complex and difficult to solve, posing a challenge for efficient solution methods and industries that rely on efficient routing. This has driven extensive research in the field of OR, developing a wide range of solution methods and algorithms.

2.7.1 Exact algorithms

The method of exact algorithms aims to provide optimal solutions to optimization problems through a systematic and exhaustive search of all possible solutions [29]. These algorithms use mathematical optimization techniques to determine an optimal solution based on a set of constraints and an objective function. Given that the exact algorithms seek an optimal solution, these solution methods are computationally intensive.

Integer Linear Programming

The VRP can be solved using a variety of exact algorithms, including variants that employ Integer Linear Programming (ILP). ILP is an extension of LP that adds the requirement that some or all of the decision variables must take on integer values. This restriction is what sets ILP apart from LP, which allows variables to take on any real value [35]. ILP problems are similar to LP problems in that they involve maximizing or minimizing a linear objective function subject to a set of linear constraints.

ILP is a valuable tool for solving problems that involve selecting discrete numbers [4]. The Integer Linear Programming method is a good fit for the VRP because such problems often involve selecting discrete or integer values for decision variables. In the Vehicle Routing Problem, for example, the decision variables might include the assignment of vehicles to specific routes or the order in which customers are visited. These decisions can be intuitively expressed as integer variables, making Integer Linear Programming a natural method for optimizing the problem.

Branch and Cut

The branch-and-cut algorithm is one of the most widely used ILP techniques. This method combines two techniques: branching and cutting. The branching step involves dividing the problem into smaller sub-problems by assigning a value to one of the integer variables. The cutting step includes adding new linear constraints to exclude previously explored regions of the solution space [36]. Combining these techniques helps to efficiently explore the solution space and find the optimal integer solution. Various variants of this technique have been proposed [37] [38] [39]. However, exploring the entire solution space is still computationally intensive, and the approach's success in solving larger instances of problems is limited [29].

Branch and Bound

Branch and Bound is another exact algorithm commonly used in optimization. It is similar to Branch-and-cut in that it systematically explores a problem's feasible region. Still, it does not require a linear programming relaxation or cutting planes. This algorithm is an established approach for solving a variety of optimization problems. Studies applying it to the VRP date back to 1969 [40], and it has been implemented in numerous contexts.

The Branch and Bound algorithm divides the problem into smaller sub-problems known as branches, each representing a potential solution. Using lower bounds, the algorithm can eliminate sub-optimal branches by subdividing the remaining branches into even smaller sub-problems and thoroughly investigating all feasible solutions. As the algorithm continues to examine the most promising branches, these lower bounds are refined until an optimal solution is found [41]. Thus, this method is a technique that, with enough time and resources, can guarantee to find an optimal solution.

Overall, exact algorithms are not practical for large-scale problems, and their use is limited in realistic-sized instances due to their high computational requirements. For example, the most sophisticated exact algorithm for the VRP can only solve instances of up to about 100 nodes [29]. However, they provide the optimal solution that satisfies all the constraints and valuable information and is a good guide for solving optimization problems.

2.7.2 Heuristics

Heuristics are methods used to find good solutions to problems without guaranteeing their optimality [42]. Unlike exact algorithms, heuristics are generally faster and less computationally expensive. They are based on rules of thumb, intuition, or experience rather than mathematical guarantees. Because of this, heuristics are often used to solve problems that are too large or complex to be solved by exact methods. While heuristics do not provide optimal solutions, they can often find high-quality solutions quickly and efficiently, making them valuable tools in many real-world applications [4].

One of the primary advantages of heuristics is their speed and efficiency in finding solutions, especially for large-scale problems. In addition, heuristics can be used as an initial step to identify potential solutions before applying more sophisticated algorithms to find optimal solutions.

Nevertheless, heuristics are typically problem-specific and significantly dependent on problem structure and problem-solving experience[43]. This means that they are intended to address a specific problem and may not be effective for other problems. Consequently, devising effective heuristics that can be applied to a wide variety of problems is a challenging task that requires problem-specific knowledge and expertise.

Moreover, they are frequently satisfied with discovering a locally optimal solution when exploring the solution space. Therefore, these methods do not always ensure the optimal solution, as the sub-optimal solutions they generate may not be the global optimal [4]. Therefore, the heuristics are inappropriate for problems requiring high accuracy, such as financial forecasting or medical diagnosis.

Different heuristics are suited to different types of problems, and some problems may require a combination of different heuristics to find the best possible solution. Additionally, new heuristics are constantly being developed and refined, making it an interesting field of research.

The Clarke and Wright algorithm
The Clarke and Wright algorithm is the best-known instance of a constructive heuristic and is used in vehicle routing problems to determine efficient routes for a fleet of vehicles to service a set of customers [44]. The algorithm generates initial routes starting and ending at the depot and one additional customer to optimize routes. The algorithm then calculates potential distance savings by merging two routes and sorts the results in descending order. Routes with the highest savings are merged, and the process continues until no feasible and profitable merges are possible [45].

Beam Search
Beam search is a widely used heuristic algorithm that efficiently searches for the best solution by generating and evaluating a set of possible successor solutions based on a given evaluation function. At each iteration, the algorithm prunes the candidate set, retaining only a limited number of the most promising solutions, called the beam width, to balance search space exploration with exploitation and discover the best solution [46].

To overcome the time-consuming nature of branch-and-bound, beam search was developed as an adapted technique that efficiently explores the search space by evaluating and selecting the most promising nodes at each search level based on an evaluation function [47]. Beam search is commonly used in various applications, such as optimization problems and machine learning methods, such as natural language processing. It is also useful in artificial intelligence because it efficiently investigates a large search space while consuming minimal computational and memory resources.

Heuristics can be effective for a wide range of optimization problems. Still, they require careful selection and tuning to ensure they are well-suited to the specific problem. They are also subject to the risk of getting stuck in local optima or sub-optimal solutions. Heuristics are a powerful tool for tackling complex optimization problems and providing approximate solutions with a trade-off between accuracy and computational time.

2.7.3 Metaheuristics

Metaheuristics are higher-level problem-solving methods used to guide the search for solutions to complex optimization problems. Unlike exact algorithms, which guarantee optimal solutions, metaheuristics are approximate methods that aim to find good enough solutions within reasonable time [43].

Metaheuristics can be thought of as a type of heuristic that works at a higher level of abstraction, combining multiple heuristics and strategies to improve the search process. They are used when exact algorithms are not practical or feasible, such as in large-scale or complex problems where finding an exact solution is computationally infeasible. While classic heuristics are more problems specific, metaheuristics can be used to solve a wider range of problems. They provide a general framework or policies for finding good solutions, and they are considered general optimization algorithms that can be applied to a wide range of problems.

Metaheuristics differ from heuristics in that they often incorporate randomness to effectively explore the search space and avoid getting stuck in local optima. They are able to tolerate temporary solutions that are worse than or similar to the currently obtained solutions, allowing them to search beyond the first local minimum [29].

Tabu Search
Tabu Search is a widely recognized metaheuristic that utilizes a common-sense approach to assist the search process in avoiding local optima [4]. Employing memory-based mechanisms, this type of local search is inspired by the workings of human memory and begins with a set of initial candidate solutions. Each solution is evaluated using an objective function. The algorithm then modifies the solution slightly to explore its neighborhood. Certain moves are prohibited or tabu for a fixed number of iterations to avoid revisiting solutions or becoming trapped in local optima.

Genetic Algorithms
Genetic algorithms are a metaheuristic that involves searching for a population of potential solutions [29]. Unlike Tabu Search, which is based on the local search approach, Genetic Algorithms use an evolutionary approach inspired by the biological process of natural selection. The genetic algorithm involves creating a population of potential solutions and applying genetic operators such as selection, crossover, and mutation to evolve the population toward better solutions. The algorithm continues to iterate until a stopping criterion is met or a satisfactory solution is found [4].

Ant Colony Optimization
Ant Colony Optimization (ACO) is a metaheuristic algorithm that takes inspiration from the behavior of ants. It mimics the way ants leave scent trails to indicate favorable paths and use them to navigate their search. The algorithm utilizes a probabilistic decision rule to determine the next step and iteratively updates the trials based on the quality of the solutions found [48]. ACO is a popular method for solving optimization problems, such as TSP and VRP.

2.7.4 Striving towards the best solution method

Metaheuristics are preferred for solving real-world problems because they efficiently search through solution spaces and find high-quality solutions for larger instances. Metaheuristics are highly adaptable and practical because they can escape local optima, explore search spaces effectively, and accept suboptimal solutions. Although they may not always find the optimal solution, they often come very close while remaining computationally efficient. As a result, they offer a promising approach to solving complex problems with practical applications across various fields.

Searching for the best solution method for the VRP is ongoing. The ideal solution methods should be computationally efficient, generalize well to various real-world applications, and provide solutions close to optimal. Unfortunately, no single solution method currently meets all these criteria perfectly. However, ongoing research is continuously striving towards achieving these objectives.

Exact algorithms can provide optimal solutions for the VRP. However, due to their high computational complexity, they are not practical for solving real-world problems. Although research has been conducted to improve these algorithms, their applicability is still limited. For instance, Laporte reported that the best exact algorithm showed inconsistent performance and could only solve problems with up to 100 nodes. Despite these limitations, exact algorithms play a crucial role as baselines or starting points for other solution methods.

Utilizing heuristics can help in addressing the challenge of computational time. Generating solutions that closely approximate the optimal solution effectively balances solution quality with computational time. However, heuristics also have their limitations. For instance, heuristics require significant effort for customization to specific problems. Therefore, developing more generalized solution methods that can provide near-optimal solutions is desirable while remaining adaptable to various problems. In recent years, different heuristic solution methods have emerged, and promising research directions have been explored using machine learning to build improved solution methods.

2.8 The Potential of Machine Learning for Solving the Vehicle Routing Problem

As previously mentioned, the VRP poses a significant challenge for diverse solution methodologies due to its classification as an NP-hard problem. Consequently, obtaining near-optimal solutions in a reasonable period is a difficult task. Potential synergies between machine learning, known for its efficient computation and innovative research, and optimization problems have been explored in this context. The idea is to leverage the strengths of both fields to create more generalized solution methods that produce near-optimal results while minimizing computational time. As the field of machine learning continues to advance, it holds great potential for building effective solution methods to solve optimization problems, including the VRP.

()ield of artificial intelligence that enables computer systems to improve their performance on tasks through experience with data, allowing for the automation of complex tasks and accurate predictions. It offers frameworks and techniques for automating data-driven learning processes without explicit programming. ML algorithms analyze input data to make predictions or decisions and help identify patterns and relationships within large datasets. The paradigm is widely used in various fields, including autonomous vehicles [49], cancer detection [50], and stock market forecasting [51].

The synergy of ML and OR can lead to good solution methods to solve optimization problems, such as the VRP. Traditional optimization techniques rely on assumptions about the problem to define the objective function and constraints and often struggle with computational time and solution quality. However, ML techniques can analyze vast amounts of data and identify patterns that may not be apparent through traditional optimization methods.

Recent research has demonstrated that integrating ML with OR methods shows great potential for solving solve optimization problems, such as the VRP [9] [10] [14]. Traditional optimization techniques rely on assumptions about the problem to define the objective function and constraints and often struggle with computational time and solution quality. The synergy between OR and ML is an emerging area of investigation, and researchers have already achieved promising results in this field. While incorporating ML will not overcome the fact that these problems are NP-hard, it can help build solution methods that approach established heuristics and metaheuristics. As such, this approach represents a promising direction for addressing routing problems in practical applications.

3

Green Vehicle Routing

The Green Vehicle Routing Problem (GVRP) is gaining increasing recognition as a crucial problem in the transport and logistics sectors. Global warming is causing more frequent extreme weather events, heatwaves, and the annual vanishing of Arctic sea ice, all of which indicate that climate change is already beginning to affect us [52]. As a result, businesses and governments are now forced to prioritize sustainability and reduce their carbon footprints. The GVRP is a variant of the traditional VRP and is designed to minimize the environmental impact of transportation activities. Its focus is on reducing harmful pollutant emissions and greenhouse gases [7]. Consequently, the GVRP has become an essential tool for promoting sustainable transportation and logistics, offering significant potential to mitigate negative environmental effects.

This thesis aims to shed light on the significance of sustainable transportation and its potential to reduce negative environmental impacts through the implementation of green vehicle routing strategies. In the upcoming chapter, we will introduce the concept of GVRP and elaborate on the significance of sustainability in the transport sector. We will discuss how frameworks such as the Triple Bottom Line can be employed to integrate sustainability in transportation and logistics. Moreover, we will delve into different approaches to integrate environmental considerations into routing problems and provide examples of their applications in real-world scenarios. The content of this chapter is influenced by the research conducted in the preceding project thesis [19].

3.1 Sustainability in the Transport Sector

Transportation is essential to society as it facilitates the movement of people and goods, driving economic growth and development. For example, in 2019, the volume of goods transported globally by sea surpassed 11.1 billion tons, emphasizing the critical role of transportation in enabling international trade and commerce [53]. With ongoing globalization, the transportation sector is expected to become even more significant in the coming years. Consequently, the importance of the transportation industry in sustaining and promoting economic growth and social welfare cannot be overstated.

However, the transportation sector faces several challenges that must be addressed to reduce its impact on the environment and promote sustainable development. One of the most significant challenges is the heavy dependence on fossil fuels, particularly in road transportation. This reliance has resulted in the sector being responsible for 37% of global carbon dioxide emissions from end-use sectors in 2021, making it one of the largest sources of greenhouse gas emissions worldwide [1].

The latest IPCC report highlighted fossil fuel combustion in land transportation as one of the top three largest contributors to global greenhouse gas emissions on a 100-year time scale [54]. This dependency on fossil fuels not only contributes significantly to greenhouse gas emissions but also makes the sector vulnerable to price volatility and supply disruptions.

The transportation sector must prioritize sustainability and reduce its environmental impact. Urgent action is necessary to limit global warming to 1.5°C above pre-industrial levels, mitigating severe climate change impacts [55]. It is imperative that policymakers and stakeholders work together to address these challenges and promote sustainable transportation solutions to protect the environment and human health. With the urgency to address climate change, businesses and governments are searching for ways to reduce their environmental impact, and the GVRP provides a practical solution to this challenge.

In GVRP, the objective is to optimize a route in terms of cost minimization by selecting the cheaper route according to various cost functions coupled with environmental impact minimization, such as carbon emissions and fuel consumption. This is achieved by incorporating additional constraints into the classic VRP or through a broader objective function. The goal is to find a set of routes that minimize the total environmental impact of the fleet while satisfying all of the constraints. This problem is particularly relevant in modern transportation systems, where there is increasing pressure to reduce greenhouse gas emissions and other environmental impacts associated with transportation.

3.2 The Triple Bottom Line

The pursuit of sustainability has become an increasingly significant objective for many businesses. However, the ambiguity surrounding the terminology of "green" and "sustainability" often leads to a lack of clarity in goals and purpose. Therefore, it is crucial for companies to utilize a proper framework to achieve meaningful sustainability objectives [56].

Triple Bottom Line (TBL) is a framework that provides a comprehensive approach to sustainability by taking into account three dimensions: social, environmental, and economic [3]. This approach recognizes that achieving sustainable development requires a balance of environmental, social, and economic considerations, as illustrated by Figure 3.1. By incorporating these dimensions into their decision-making processes, companies can identify ways to align their operations with sustainability goals and objectives.

The TBL framework serves as a guide for companies to pinpoint their objectives and establish clear goals towards sustainability. By balancing the three dimensions, companies can achieve a sustainable business model that addresses the environmental impact, social equity, and economic viability of their operations. This model leads to the creation of sustainable value, which contributes to a company's long-term success rather than short-term gain.

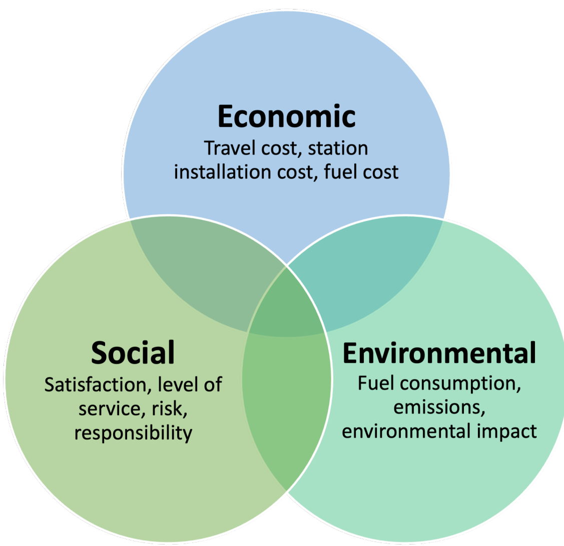

Figure 3.1: This figure illustrate the triple bottom line framework which represents the interplay between three essential dimensions of sustainability: economic, social, and environmental. These dimensions are depicted as overlapping circles, highlighting the crucial importance of achieving a harmonious balance among them.

The environmental dimension of sustainability involves minimizing the negative impact of human activities on the natural environment. This includes reducing carbon emissions, protecting biodiversity, and conserving natural resources. The social dimension of sustainability focuses on ensuring that human well-being is protected and enhanced.

This includes promoting social equity, ensuring access to basic needs such as clean water and healthcare, and protecting human rights. The economic dimension of sustainability involves ensuring that economic growth is compatible with sustainability goals. This includes promoting sustainable business practices, investing in renewable energy, and ensuring fair and equitable distribution of resources.

This framework is also applicable to the GVRP. Figure 3.1 contains examples of the three TBL factors relevant to the GVRP. Economic factors can include driver- or vehicle-related travel expenses as well as gasoline and diesel prices. Social aspects can include customer satisfaction, such as on-time package delivery. It can also include risk or responsibility concerns [57]. The environmental aspect is often associated with emissions or fuel consumption, but it can also take into account other environmental impacts. All of these concerns can be incorporated into the GVRP in a variety of methods.

This framework has become increasingly important in the field of transportation and logistics as businesses and governments seek to balance the economic benefits of transportation with the social and environmental impacts. The GVRP is an example of how the TBL framework can be applied in practice.

3.3 Solution Methods

The development of algorithms to solve the GVRP is an ongoing field, with new directions and proposals emerging frequently. To incorporate environmental considerations, there are several common approaches, four of which are described below.

1) **Including environmental considerations as the primary objective**

$$min \quad \sum_{i,j \in A} E_{ij} x_{ij} \tag{3.1}$$

The first possible method for incorporating environmental concerns into the optimization model is to make environmental pollution the primary objective rather than the conventional cost of distance [58]. The goal is to minimize total emissions in kilograms per hour, where total emissions are a function of the emissions associated with the route of each vehicle.

Optimizing for emissions alone may not always be the optimal strategy, as it may necessitate sacrificing other essential goals such as cost, safety, or reliability. Therefore, it is often necessary to consider multiple objectives to address these trade-offs. For instance, optimizing for emissions minimization can result in less cost-effective or less reliable routes. In this case, considering multiple objectives can lead to more realistic and logical routes that account for various trade-offs between emissions reduction, cost, and reliability, making it a more comprehensive approach.

2) **The multi-objective problem approach**

$$min \quad \sum_{i,j \in A} E_{ij} x_{ij} \tag{3.2}$$

$$min \quad \sum_{i,j \in A} c_{ij} x_{ij} \tag{3.3}$$

A multi-objective problem is an optimization problem in which multiple, often conflicting, objective functions must be optimized simultaneously [59]. The problem depicted in Equations 3.2 and 3.3 is an example of a multi-objective problem, which involves finding a solution that balances the cost of distance c_{ij} and the cost of emissions E_{ij}. This approach requires finding a compromise solution that balances both objective functions, considering that minimizing one objective function may come at the expense of the other. Therefore, the goal of multi-objective optimization is to identify the set of optimal solutions which represent the best trade-offs between the competing objectives without compromising the quality of any individual objective.

This method has pros and cons. It produces more authentic routes that consider various factors, but it may not be the most cost-effective or emission-efficient [60]. For businesses, considering multiple factors is crucial, and this method can help balance them. However, a multi-objective strategy can increase complexity and make interpreting results more challenging.

3) **Adding environmental pollution as a penalty to the objective function**

$$min \quad \sum_{i,j \in A} c_{ij} x_{ij} + C_e \times \sum_{i,j \in A} E_{ij} x_{ij} \tag{3.4}$$

Incorporating environmental concerns by adding environmental pollution as a penalty to the distance objective function is an alternative approach. This involves representing emissions as a cost that can be added to the cost of the routes, where C_e represents the penalty cost in dollars per kilogram of environmental pollution. While this approach can be effective, accurately estimating the cost of emissions can be difficult, and it only considers the cost of emissions. This method can be employed to avoid purchasing a portion of the CO2 emission quota [61].

4) **Adding environmental pollution as a constraint**

$$\sum_{i,j \in A} E_{ij} x_{ij} \leq E_{limit} \tag{3.5}$$

One way to optimize with regard to environmental concerns is to add an environmental pollution constraint to the original cost of distance problem, as shown in Equation 3.5. This constraint limits total emissions to a predetermined level, allowing for the inclusion of specific goals that a business or government may have. For instance, the EU2020 target of 147 g/km can be utilized as an upper bound while optimizing [62].

Sustainable Solution Routes for Routing Problems

The real world is intricate, and it is always possible to add more constraints to modify the model. However, there will always be a balance between not over-constraining the model to the point of no feasible solutions and avoiding excessive computational time. For instance, when utilizing the simplex method to solve VRPs, the number of functional constraints has a much greater impact on computational time than the number of decision variables. However, it is worth noting that solving the dual problem can provide a way to reduce the number of constraints in the primal problem while still finding a feasible solution. This is an important consideration to ensure that the model is not overly constrained or too computationally intensive.

In addition to cost and environmental factors, the optimization model can also include social considerations. For example, route consistency may be incorporated. It has been demonstrated that this results in safer routes, as the driver becomes highly familiar with the route [57]. Other social factors, including consumer satisfaction, can also be taken into account.

Creating more comprehensive models that take into account a larger number of real-world factors can lead to more sustainable solutions. Sustainable solutions refer to the balance of the TBL framework, which aims to achieve an optimal balance between social, environmental, and financial considerations. This comprehensive approach strives to provide improved long-term solutions for the future.

3.4 The Pollution Routing Problem

The environmental factor E in the mentioned solution methods for the GVRP can represent different environmental concerns, but fuel consumption is a particularly promising and intriguing consideration. The Pollution Routing Problem (PRP) is a variant of the GVRP that addresses this topic. The PRP focuses on reducing the pollution emissions produced by vehicles during transportation by optimizing the route and minimizing the total quantity of pollutants emitted [63]. This is achieved by incorporating environmental constraints or penalties into the objective function of the optimization model.

Addressing the PRP is particularly important because transportation activities are a significant source of pollution and greenhouse gas emissions. By developing efficient and effective solutions to the PRP, we can help reduce the negative impact of transportation on the environment and move towards more sustainable and responsible transportation practices. This can have a significant effect on public health, as air pollution is linked to various respiratory and cardiovascular diseases. Additionally, addressing the PRP can help companies comply with environmental regulations and reduce their carbon footprint.

Fuel Consumption
When fuel is burned in an engine, it produces emissions such as CO2 and NOx, which are GHGs that contribute to global warming and air pollution. The amount of emissions produced is directly related to the fuel burned, as stated by Equation 3.6 [63].

$$E = \delta_1 F + \delta_2 \tag{3.6}$$

The emission rate of the engine-out is represented by E, while F denotes the rate of fuel use, and δ_1 and δ_2 are the GHG-specific emission index parameters. By reducing the amount of fuel used, we can also reduce the amount of emissions produced. Therefore, minimizing fuel consumption is an important strategy for reducing GHG emissions and addressing climate change.

Some might wonder why minimizing distance is insufficient, contending that doing so reduces fuel consumption. In reality, however, we can achieve superior results in minimizing fuel consumption and emissions by incorporating a more comprehensive objective into the optimization problem. Therefore, we must evaluate factors other than the distance in the GVRP.

Several studies have examined these factors and their impact on fuel consumption. They have investigated how road gradient, speed, and load affect fuel usage. Additionally, they have looked into the effects of factors such as travel time when the engine is cold and traffic congestion [64]. Including traffic congestion as a concern in the GVRP is an example of how the solution route can drive a long distance to avoid congestion and reduce emissions.

In real-world applications, environmental concerns are often included as constraints or penalties but are not typically the primary objective function. This is because prioritizing environmental factors results in the least fuel consumption, but the routes must also be practical and realistic. For instance, if the optimization algorithm solely focuses on road gradients and zigzags over multiple altitude meters while traversing many different streets, the resulting routes may become impractical.

3.5 Smart Waste Collection Routing Problem

Waste management is a challenging field with significant environmental, economic, and public health implications. In 2020, the European Union generated approximately 2.15 million tonnes of waste, averaging 4.8 kg per person [65]. Inadequate recycling practices lead to environmental degradation and resource depletion. Landfills and incineration facilities, commonly used for waste disposal, occupy valuable land and emit greenhouse gases, contributing to climate change. Moreover, conventional waste collection methods involving diesel vehicles result in high emissions.

Waste collection processes are a consistent contributor to emissions of greenhouse gases. Conventional waste collection methods typically involve diesel combustion engines, which emit high levels of pollutants. Of this, a substantial portion can be attributed to waste collection processes. Adopting electric or biofuel-powered collection vehicles, promoting waste reduction and recycling to reduce the quantity of waste requiring collection and disposal, and optimizing waste collection routes to reduce fuel consumption can reduce these emissions. By optimizing waste collection routes, it is possible to reduce these emissions and mitigate the impact of waste on the environment.

The GVRP problem is a compelling topic with immense potential for companies to reduce emissions, gain a competitive edge, and work towards sustainable goals. Although the field is growing, research on effective solving methods for the GVRP is still limited, especially for real-world applications. Many industries can benefit from using GVRP as a decision-making tool, and an interdisciplinary approach involving energy use and environmental impact, public policy, engineering, and transportation system management is crucial to develop the field further.

While some research has examined the practical applications of VRPs [66] [67] [68], there is a clear need to prioritize real experimental data and real-world case studies from diverse backgrounds, especially regarding GVRPs, which still lack sufficient exploration. Incorporating real-world cases is essential for obtaining valuable insights and developing a comprehensive understanding to pursue optimal methods. Ensuring that the solution methods accurately represent real-world scenarios is crucial; accounting for variations in coordinate positions, scaling, and other relevant factors to enhance their accuracy and practicality is vital.

With this in mind, machine learning is an exciting area of research to investigate further how the power of machine learning can be utilized to build good solution methods for the GVRP. Overall, the GVRP has significant potential, and prioritizing it in the future could yield many positive outcomes.

Considering this, exploring the application of machine learning in developing effective solution methods for the GVRP is an exciting research direction. Emphasizing the importance of the GVRP in future studies can lead to numerous beneficial outcomes, given its substantial potential for addressing environmental and logistical challenges in transportation.

4

Neural Networks

Machine learning is a subfield of artificial intelligence that uses algorithms to enable computers to learn from data, improving their performance over time. Neural networks are an effective algorithm for machine learning that mimics the structure and function of the human brain. They consist of interconnected nodes, known as neurons, that process and transmit information in a way that can be trained to perform a specific task.

Neural networks have emerged as a powerful tool with a broad range of applications. In recent years, there has been an increased focus on investigating the potential of neural networks for solving optimization problems, including the Vehicle Routing Problem (VRP). The remarkable success that neural networks have shown in solving such problems indicates that they may have the power to revolutionize the field of operations research. Neural networks are particularly useful when dealing with large datasets and complex problem structures, as they can learn patterns and relationships in the data that are difficult to detect with traditional optimization approaches.

In this section, we will examine neural networks, a powerful tool in machine learning. Before discussing the foundations of neural networks, we will explore key concepts in machine learning, such as supervised, unsupervised, and reinforcement learning. We will then cover the foundational elements of neural networks, including gradient descent, backpropagation, loss functions, and activation functions. Finally, we will examine how neural networks can be applied to solve optimization problems, such as the VRP.

4.1 Machine Learning

Neural networks are a type of machine learning algorithm that has gained increasing popularity in recent years due to their ability to learn from complex and large datasets. Before diving into the foundations of neural networks, it is important to understand the concept of machine learning and its different subfields. Machine learning is a type of artificial intelligence that enables computers to learn from data and make predictions or decisions. Three subfields of machine learning are supervised, unsupervised, and reinforcement learning. Although each of these subfields is different, they all share the goal of training machines to perform specific tasks.

4.1.1 Supervised Learning

Supervised learning is the most common and maps input data to output data based on labeled examples provided during training [69]. The training data consists of a set of input and output pairs, which are split into training and test sets. During training, the algorithm iteratively adjusts its weights to minimize the difference between predicted and actual outputs in the training set. Once the training is complete, the algorithm can make predictions on new, unseen data by generalizing the patterns learned during training. Test data is used to evaluate the performance of the trained model on new data, ensuring that it can generalize well beyond the training set.

Supervised learning can be used for various tasks, such as classification, regression, and prediction. This approach has found successful applications in diverse fields such as computer vision and natural language processing [70] [**NLP**]. Its flexibility and generalization ability have made it popular, and it provides a clear and direct feedback mechanism to the learning algorithm. However, to achieve optimal performance, supervised learning requires high-quality training data that is accurately labeled.

4.1.2 Unsupervised Learning

Unsupervised learning is when algorithms learn patterns and relationships in input data without being explicitly given output data to learn from [69]. Unlike supervised learning, there are no labeled examples in unsupervised learning. The goal of unsupervised learning is to find structure in the input data, such as clustering similar data points together or identifying underlying patterns in the data. This is typically done using techniques such as clustering [71], anomaly detection [72], and dimensionality reduction [73].

In clustering, the algorithm groups similar data points together based on some similarity metric. Anomaly detection involves identifying data points that are significantly different from the rest of the data. Dimensionality reduction involves reducing the number of features in the data while preserving the most important information.

One challenge with unsupervised learning is the difficulty of measuring the performance of the model because there is no objective metric to evaluate the output. Unlike supervised learning, where the correct output is provided for each input. Instead, unsupervised learning relies on assumptions about the structure of the data to guide the learning process. However, one of the benefits of unsupervised learning is that it enables learning without the need for labeled training data, which can be time-consuming and expensive to obtain.

4.1.3 Reinforcement Learning

Reinforcement learning is another type of machine learning that involves an agent interacting with an environment to learn how to take actions that maximize a cumulative reward [74]. Unlike supervised learning, where the correct output is provided for each input, and unsupervised learning, where the model learns from unlabeled data, reinforcement learning involves learning from feedback in the form of rewards or punishments.

In reinforcement learning, the agent takes actions based on its current state and the rewards it has received in the past. The environment responds to the agent's actions by transitioning to a new state and providing a reward or penalty based on the action taken. The goal is to learn a policy that maximizes the total reward obtained over time [75].

One of the key challenges in reinforcement learning is the trade-off between exploration and exploitation. The agent must balance the desire to explore new actions and learn about the environment with the need to exploit its current knowledge to maximize reward. Reinforcement learning has many applications, including game playing [76], robotics [77], and self-driving cars [78]. In addition, it has been shown to be effective in solving optimization problems [9], a topic on which the following chapters of this thesis will elaborate.

4.2 Introduction to Neural Networks

Now that we have explored the key concepts in machine learning, we will examine the foundations of neural networks. Neural networks are models composed of interconnected processing nodes, called neurons, arranged in multiple layers [79]. They are inspired by the structure of the brain, where biological neurons are connected in a network. Neurons are the fundamental building blocks of neural networks and operate as data processing and transmission units. They receive input signals from other neurons or input layers, perform a function on those inputs, and then transmit an output signal to other neurons.

Neural networks consist of multiple interconnected neurons, and each connection between neurons is associated with a weight that determines the strength of the connection. These weighted connections play a crucial role in transmitting the output signal from one neuron to another throughout the network [69]. During the training process, the weights associated with these connections are learned and adjusted to optimize the performance of the network.

In addition to weight and learning, bias and activation functions are two other key concepts that are critical to understanding neural networks. Bias is a constant value that is added to the input of each neuron. The bias allows the neural network to adjust the output of the neuron independently of its inputs. Activation functions are applied to the output of individual neurons to introduce non-linearity into the model. This enables models to perform more complex tasks. Neural networks use weights and biases to process input data and make predictions or decisions, and through the process of training, they learn to adjust the weights to improve their performance in various tasks.

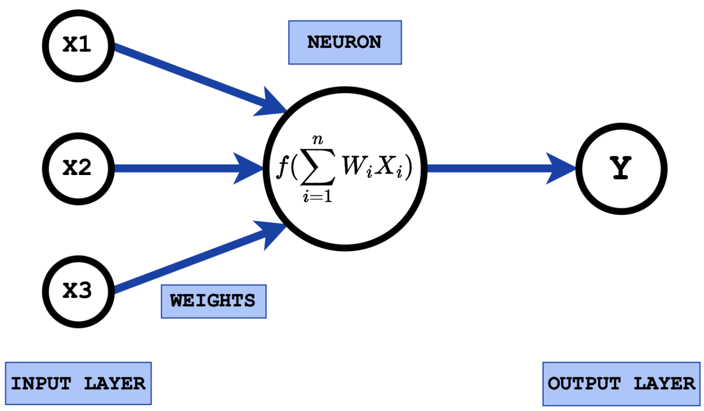

Figure 4.1: This figure illustrates a neuron for a neural network, which consists of an activation function f and the sum of the weights W and input X. The output Y of the neuron is determined by these elements.

4.3 Feedforward Neural Networks

Feedforward neural networks are a type of neural network that processes information in a single direction, from input to output, through a series of interconnected neurons without any feedback loops. Each neuron in a feedforward neural network receives inputs, applies weights and bias, and produces an output that becomes the input to the next layer of neurons until the final output is generated.

While various specialized and complex neural network architectures have been developed in recent years, feedforward neural networks remain a fundamental and widely used type of neural network. At the core of all neural networks is the basic structure, where the input layer receives data from the outside world, and the output layer produces the final result. Between the input and output layers, hidden layers can perform complex computations on the input data, allowing the network to learn and extract features from the data.

$$y = f(\sum_{i=1}^{n} W_i X_i) \tag{4.1}$$

Equation 4.1 represents the computation of a single neuron in a feedforward neural network. This equation combines the key components of the forward pass, which is the flow of information through the network in the forward direction. A forward pass refers to the process of computing the output of the network given an input. During the forward pass, the input is passed through the network one layer at a time, with each layer applying a linear transformation to the input followed by a non-linear activation function. The output of the final layer of the network is the prediction of the network for the given input.

In Equation 4.1, X_i represents the input to the neuron, W_i represents the weight assigned to each input, and the function f is a non-linear activation function that is applied to the sum of the weighted inputs and a bias term to produce the output y. The equation is also referred to as the transfer function and is a fundamental unit of a neural network. The transfer function is essential for the network to perform complex computations by allowing the neurons to model complex non-linear relationships between inputs and outputs [80].

A single-layer neural network is the simplest type of feed-forward neural network, consisting of a single layer of output neurons connected directly to the input neurons. This type of network is typically used for simple classification tasks, where the input data is linearly separable.

A multi-layer neural network, on the other hand, consists of multiple layers of neurons, including one or more hidden layers between the input and output layers. These hidden layers allow the network to perform nonlinear computations on the input data, enabling it to solve more complex problems that cannot be solved by a single-layer neural network.

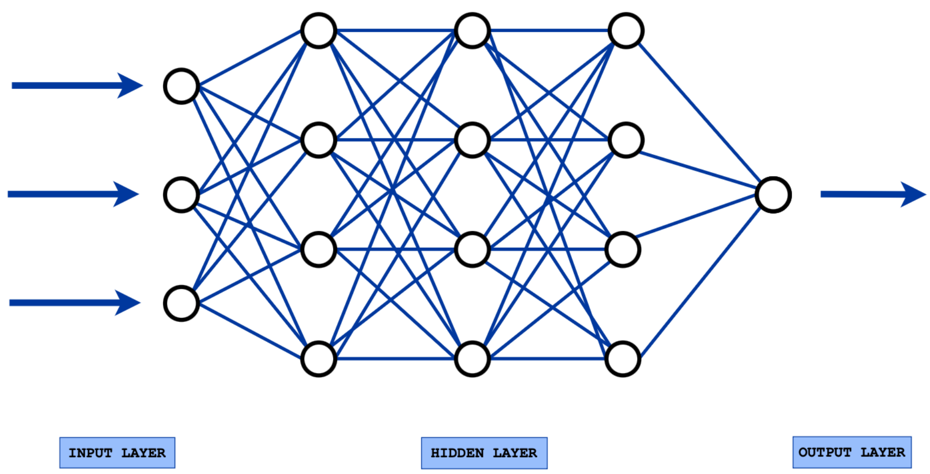

Figure 4.2: This figure illustrates a fully connected neural network with the input layer, three hidden layers with four neurons each, and an output layer with one neuron. The input layer is fully connected to the first hidden layer, which is fully connected to the second hidden layer, and so on, until the output layer.

A fully connected neural network is a type of multi-layer neural network where every neuron in one layer is connected to every neuron in the next layer, shown in Figure 4.2 [75]. This type of network can learn very complex and nonlinear relationships between the input and output data, but it requires a large number of parameters and can be prone to overfitting, becoming too specialized.

While the key foundations of neural networks include neurons, activation functions, weights, and biases, these alone are not enough to train a neural network to perform a specific task. To train a neural network, we also need to define a loss function, which quantifies the difference between the predicted output of the network and the actual output. This is a crucial component of the training process, which we will discuss in the following section.

4.4 Loss function

Loss functions play a critical role in training neural networks. The loss function serves as a measure of how well the learning model is performing a specific task and is used to guide the optimization process during training by updating the model's parameters in the direction that minimizes the loss.

For supervised learning, the loss function takes the neural network's predicted outputs and compares them to the true outputs or labels of the training data. The function then calculates a value that represents the difference between the predicted outputs and the true outputs. The goal of training the neural network is to minimize this error. Different types of loss functions can be used depending on the type of problem being solved. For example, in classification problems, the cross-entropy loss function is commonly used [81]. In regression problems, the mean squared error loss function is often used [8]. Equation 4.2 represents the mean squared error (MSE) between the predicted values $\hat{y}_i$ and the actual values y_i of a dataset with n examples.

$$\epsilon = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \tag{4.2}$$

In unsupervised learning, the selection of an appropriate loss function is essential to learn the underlying structure of the data without explicit supervision. The choice of the loss function is task-specific and can be based on several factors. For instance, in clustering, the loss function is usually based on the distance between points, while in dimensionality reduction, the loss function may aim at preserving pairwise distances between points.

In reinforcement learning, the reward function is the most common loss function, which assesses the performance of the agent in the environment. The primary objective of the agent is to maximize the reward function over a sequence of actions. The selection of the reward function is crucial and requires careful consideration since it directly influences the behavior of the agent.

It is imperative to choose an appropriate loss function for effective neural network learning. An inappropriate loss function may lead to slow convergence or poor performance, while a well-designed loss function can enhance the accuracy and generalization of the neural network on new data. Therefore, careful consideration of the loss function is crucial in determining the effectiveness of the neural network.

4.5 Optimization

Optimization is a key aspect in enhancing the performance of neural networks. Gradient descent is a fundamental optimization method widely employed in neural networks for training and enhancing performance. It iteratively adjusts parameters based on the negative gradient of the loss function. Stochastic gradient descent (SGD) is a variation that randomly selects subsets or individual training examples for parameter updates.

To further improve performance, optimization methods like Adam and batch normalization are often used in conjunction with gradient descent. Adam adapts the learning rate for each parameter based on gradient moments, accelerating convergence. Batch normalization normalizes layer inputs, stabilizing training and preventing overfitting.

4.5.1 Gradient Descent

In neural networks, the iterative process of finding optimal weights and biases that minimize the loss function is crucial for learning. Gradient descent is an optimization algorithm frequently used for this purpose [69]. It adjusts model parameters by following the negative gradient of the loss function, which is essentially the rate of change of the loss function with respect to the model parameters.

This derivative provides information on how much the loss changes with a small modification of the model parameters. By gradually adjusting the model parameters in the direction of the negative gradient, the loss can be reduced.

The process of changing model parameters iteratively to minimize the loss function is repeated until a minimum is achieved or a predetermined number of iterations is reached. Gradient descent computes the gradient of the objective function with respect to the model parameters at each iteration. The gradient is a vector that points in the direction of the maximum increase of the function. The algorithm updates the model parameters by subtracting a fraction of the gradient from the current values of the parameters to optimize the function. The fraction used for this purpose is called the learning rate, and it controls the step size of the parameter update.

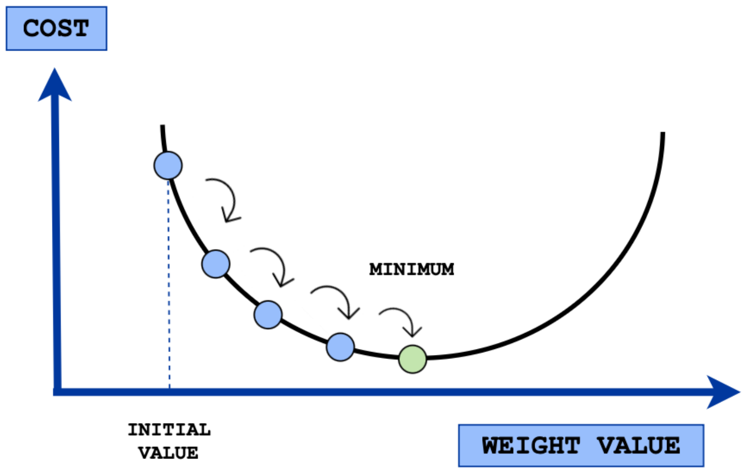

Figure 4.3: This figure illustrates how the gradient decent algorithm uses the concept of derivatives to the optimal value and improves results.

The direction of the gradient indicates the direction of the steepest ascent. To minimize the function, we must therefore move in the opposite direction of the gradient. This procedure is illustrated in Figure 4.3, in which the weight value iteratively approaches the minimum. By repeatedly modifying the parameters in the opposite direction of the gradient, we can approach the function's minimum in an iterative manner.

The weight update rule for a single weight w_r is given by:

$$w_j^r(new) = w_j^r(old) + \Delta w_j^r. \tag{4.3}$$

Gradient descent is an effective optimization algorithm that iteratively optimizes the parameters of a neural network by utilizing derivatives, slopes, and small changes. By fine-tuning the learning rate, we can control the parameter update's step size and accomplish a more rapid convergence to the optimal solution.

Stochastic Gradient Descent

There are various variants of gradient descent, including batch gradient descent, stochastic gradient descent, and mini-batch gradient descent. Stochastic gradient descent introduces the concept of updating the parameters using a small subset of samples at a time rather than the entire dataset [81]. This approach allows for more frequent updates and faster computation, making it especially beneficial for large-scale datasets.

Adam
Adam is a highly effective optimization algorithm utilized for weight updates in neural networks during training [82]. It dynamically adjusts the learning rate for each parameter based on estimations of gradient moments. Additionally, Adam incorporates bias correction to compensate for momentum effects in the initial training stages. Known for its computational efficiency and faster convergence compared to other optimization algorithms, Adam is a popular and recommended choice for neural networks. It is widely used, robust, and has proven to deliver excellent results.

$$\theta_{t+1} = \theta_t - \frac{\eta}{\sqrt{\hat{v}_t + \epsilon}}\hat{m}_t \tag{4.4}$$

Equation 4.4 represents the weight update rule in the Adam optimization algorithm, where θ_{t+1} denotes the updated weight at time step $t+1$, θ_t represents the current weight at time step t, η is the learning rate, $\hat{v}_t$ and $\hat{m}_t$ are estimates of the first and second moments of the gradients, and ϵ is a small value, added for numerical stability.

Batch normalization
In the context of optimizers, batch normalization can be used to increase the effectiveness of optimization. Batch normalization is a method for normalizing the input of each layer of a neural network to have a zero mean and one variance. This is done by normalizing the outputs of each layer using the mean and variance of the activations in the current mini-batch. The normalization is applied before the activation function, which helps prevent the output of the activation function from becoming too large or too small. Batch normalization has been shown to give more stable training and improve the gradient flow through the network. It also helps to regularize the network and can reduce overfitting.

4.6 Backpropagation

Neural networks learn to predict outcomes from input data through a process known as forward propagation, where input data is passed through neurons with specific weights and activation functions. To optimize the network's performance, gradient descent is employed to minimize the loss function by updating the network parameters. This process necessitates the computation of gradients, which is where backpropagation comes into play.

Backpropagation is an algorithm that computes the gradients of the loss function with respect to the network parameters. This is achieved by propagating the error back through the network and computing the gradients layer by layer. The algorithm begins at the output layer and moves backward through the network to the input layer.

The backpropagation process starts from the last layer of the neural network, denoted as $r = L$, and progresses backward through the layers until it reaches the input layer. The algorithm is named after this backward propagation process. To calculate the weight gradients, we first simplify the expression for the gradient of the output in the last layer h_j^L, which we denote as δ for simplicity, as follows:

$$\frac{\partial \epsilon}{\partial h_j^L} = \delta_j^L. \tag{4.5}$$

This gradient expression is then used to calculate the derivative of the loss with respect to the weights [69]. Here, h_j^L denotes the preactivation in the jth neuron in layer L, which is the last layer, and ϵ is the loss function.

To compute the gradient of the loss with respect to the preactivation in the hidden layers, backpropagation propagates backward from the last layer to the earlier layers, where $r < L$, and employs the chain rule twice to rewrite the expressions. The resulting gradient calculation is given by:

$$\frac{\partial \epsilon}{\partial h_j^{r-1}} = \frac{\partial \epsilon}{\partial h_j^r} \frac{\partial h_j^r}{\partial h_j^{r-1}} = \delta_j^r w_j^r f'(h_j^{r-1}). \tag{4.6}$$

The goal is to compute the gradient of the loss function with respect to the weights. This is achieved by applying the chain rule, which allows for the decomposition of the gradient calculation. Specifically, the chain rule is utilized as follows:

$$\frac{\partial \epsilon}{\partial w_j^r} = \frac{\partial \epsilon}{\partial h_j^r} \frac{\partial h_j^r}{\partial w_j^r}. \tag{4.7}$$

To simplify the equation, we make two substitutions. First, we replace $\frac{\partial h_j^r}{\partial w_j^r}$ with y^{r-1}. This substitution allows us to express the derivative in terms of the previous layer's output. Second, we use the previously introduced δ to substitute $\frac{\partial \epsilon}{\partial w_j^r}$. The resulting equation is as follows:

$$\frac{\partial \epsilon}{\partial w_j^r} = \frac{\partial \epsilon}{\partial h_j^r} \frac{\partial h_j^r}{\partial w_j^r} = \delta_j^r \frac{\partial h_j^r}{\partial w_j^r} = \delta_j^r y^{r-1}. \tag{4.8}$$

By applying the previously derived equation to all inputs, we can express the gradient of the weights as:

$$\Delta w_j^r = -\mu \sum_{i=1}^{N} \delta_j^r y^{r-1}. \tag{4.9}$$

The resulting expression serves as the gradient for weight updates during the iterative learning process of the neural network [69], following the principles of gradient descent. The learning rate μ is a hyperparameter incorporated into the update equation, determining the step size for weight adjustments during training. By controlling the magnitude of these adjustments, it influences the convergence and speed of the learning algorithm.

To update the weights and biases, these equations are iteratively applied to each layer of the network, starting from the output layer and progressing backward to the input layer. Once the gradients for all weights are computed, they are utilized in an optimization algorithm such as gradient descent, as discussed previously, to update the parameters of the network and improve its performance.

4.7 Activation Functions

Activations are the output values of neurons in each layer of a neural network, which are obtained by applying an activation function to the weighted sum of the inputs. The activation function is a non-linear function that introduces non-linearity to the network, allowing it to capture complex patterns and relationships in the data. As a result, activations play a crucial role in determining the output of each neuron and the overall output of the network.

Figure 4.4 illustrates some of the most frequently used activation functions, including the Sigmoid function, the ReLU and leaky ReLU functions, and the tanh function. The choice of activation function depends on the specific task and the characteristics of the dataset. In addition, the derivative of the activation function is a crucial factor in the learning process of a neural network, which will be discussed in more detail in the upcoming chapters.

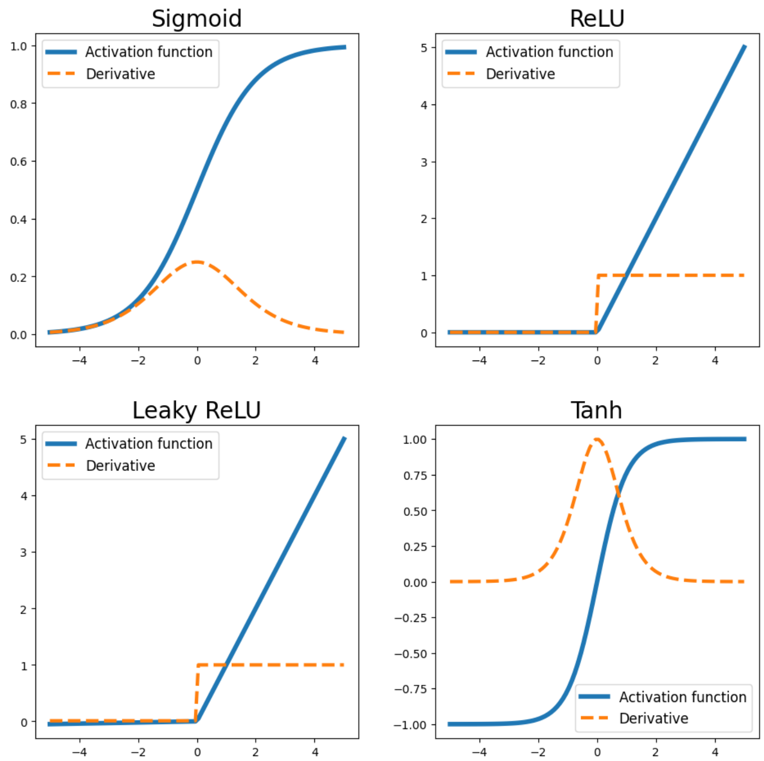

Figure 4.4: The figure illustrates the four commonly used activation functions in neural networks: sigmoid, ReLU, Leaky ReLU, and tanh. The x-axis represents the input values, while the y-axis represents the output values.

ReLU

The Rectified Linear Unit (ReLU) is a popular activation function due to its simplicity, computational efficiency, and good performance [81]. The ReLU function is defined by Equation 4.10, where the output is equal to zero for negative values and linear for positive values. This means that the function has a bounded lower end but not on the upper part.

$$f(x) = max(0, x) \tag{4.10}$$

The ReLU activation function offers several advantages over other activation functions, such as reducing the common problem of vanishing gradients. This problem will be discussed in more detail in the subsequent sections. This is due to the fact that the ReLU function has a derivative of 1 for positive input values, which helps in gradient stabilization during backpropagation. In addition, the ReLU function is simple and quick to compute, making it the most used activation function.

Softmax
Softmax is another activation function that is widely used in multiclass classification tasks in neural networks. It transforms an input vector into a probability distribution, where each output element represents the probability of the corresponding class. The softmax function applies the exponential function to each element of the input vector and then normalizes the resulting vector so that the sum of the elements is equal to 1, as shown in Equation 4.11 [81].

$$\text{softmax}(x_i) = \frac{e^{x_i}}{\sum_{j=1}^{n} e^{x_j}} \tag{4.11}$$

The Vanishing Gradient Problem
The vanishing gradient problem occurs in neural networks when gradients become too small as they propagate backward through the network during training. This makes it difficult to update the weights of early layers, which can cause those layers to not learn meaningful representations. To mitigate the vanishing gradient problem, various techniques have been developed. One such technique involves using activation functions, like the ReLU function, that have larger gradients. As shown in Figure 4.4, ReLU has a constant derivative of 1 for all positive values, allowing gradients to flow more easily through the network. In addition, initializing the weights of the network and using skip connections can also help gradients flow more easily through the network.

Skip Connections
In neural networks, skip connections, also known as residual connections, are a technique that enables information to bypass one or more network layers [83]. In typical neural networks, each layer processes the output of previous layers using weights and activation functions to generate a new output. Skipping connections enables output to pass directly from one layer to the next without intermediate processing. This can help prevent the vanishing gradient problem, which occurs when gradients become too small, and weight updates occur slowly or not at all.

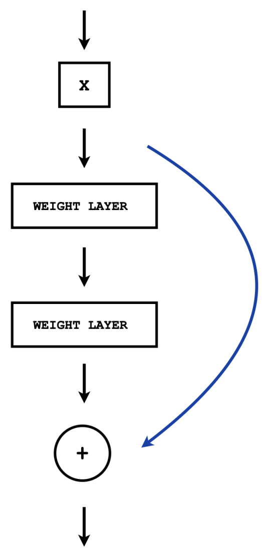

Figure 4.5: The blue arrow in the figure denotes the skip connection, which allows the input to bypass the layer without any changes. The output of the weight layers and the input from the skip connection are combined using residual addition at the plus sign. This process promotes faster convergence and improved network performance.

Figure 4.5 illustrates a skip connection, which is added to the output of two weight layers, producing the network's final output. By modifying the current representation as opposed to learning an entirely new one, the model can converge more quickly and yield more accurate results. In the ResNet architecture for image classification, it has been demonstrated that skip connections enhance deep neural networks with over 100 layers [83]. Additionally, various neural network architectures have utilized skip connections.

4.8 Neural Networks for Solving Routing Problems

Techniques from the field of Machine Learning (ML) can be used for various approaches, and one particularly exciting area is in Operations Research (OR). As discussed in Chapter 2, combining ML with OR has the potential to enhance solution methods for problems like the Vehicle Routing Problem (VRP). ML algorithms offer several advantages over traditional optimization techniques, including the ability to handle large amounts of data and flexibility. ML algorithms can adapt to different optimization problems and create more general solution methods, such as learned heuristics. These heuristics can improve the efficiency and effectiveness of optimization algorithms, leading to improved solution methods in the field of OR.

In the domain of ML, researchers have specifically focused on using neural networks for tackling the VRP. As this problem is known to be NP-hard, finding an optimal solution is extremely difficult. Balancing computational power with close-to-optimal solutions has been identified as the primary challenge. This has led to the use of heuristics, as discussed in Chapter 2. Innovative approaches have been explored by leveraging the intersection of different fields to improve these solution methods. Therefore, the primary objective of using neural networks to solve optimization problems like the VRP is to address this computational power and optimal solution balancing issue.

In previous studies, researchers have shown uncertainty about using neural networks to solve the VRP due to unsuccessful earlier attempts [84] [85]. However, as both fields have made significant progress in recent years, there is renewed interest in exploring this synergy [9] [10] [14]. Currently, researchers are focused on identifying the most effective ways to incorporate neural networks into the problem-solving process. It is important to note that neural networks are not the ultimate solution for the VRP but rather a tool to be used strategically for improved results.

The idea behind using neural networks for solving optimization problems is to quickly find the best solution to a problem based on input information. The goal is to provide a good enough solution within a reasonable time rather than finding the optimal solution, which can be time-consuming. Much like the heuristic, but now also utilizing the power of machine learning [13]. Neural networks can provide a powerful and fast computational solution to such optimization problems by taking advantage of the extensive connectivity among neurons and their ability to handle non-linear relationships and complex interactions between variables, which can be particularly useful for problems with many variables or constraints.

4.8.1 The Hopfield Network

The Hopfield Network, introduced in 1985 [13], was an early attempt to utilize neural networks for solving routing problems, specifically small instances of the Travelling Salesman Problem (TSP). Trained with city coordinates, this recurrent neural network settles into a stable state representing a good TSP solution. Although not necessarily optimal, the resulting solution is often highly effective for small problem instances. This pioneering work inspired further exploration of neural networks for tackling challenging optimization problems.

4.8.2 The Pointer Network

Although the Hopfield network was among the early adopters of exploring the combination of neural networks and routing problems, little progress was made in this field for a considerable amount of time. However, in 2015, the Pointer Network achieved promising results [86]. The Pointer Network is a neural network model that uses an attention mechanism to output a permutation of the input. This approach is supervised learning because it uses training examples to learn the conditional probability of an output sequence. The model is designed to address the challenge of finding a sequence of elements that satisfy certain constraints, which is difficult to solve using traditional sequence-to-sequence models.

4.8.3 Advancements in Neural Network-based Routing Optimization

In recent years, significant progress has been made in the neural network-driven routing optimization field. The introduction of the Pointer Network, which employs an attention mechanism to produce a permutation of the input, was a breakthrough. Furthermore, recent models have demonstrated the potential for attention mechanisms to enhance the performance of routing optimization models, indicating a promising research direction.

Researchers have also investigated the possibility of training networks without labeled data using unsupervised or reinforcement learning techniques. This approach can substantially reduce the time and cost required for model training in routing optimization problems, leading to more efficient and accurate models.

Reinforcement Learning
Supervised learning is the most commonly used, but applying it to routing problems is challenging due to the difficulty in obtaining optimal labels for new routes. As a result, alternative approaches, such as reinforcement learning, have been explored for optimization problems [17, 87]. In this context, reinforcement learning evaluates solution quality and provides reward feedback to the learning algorithm [88].

Recent work has shown the effectiveness of reinforcement learning in solving route optimization problems. This approach offers improved solution quality compared to current deep learning approaches, faster computation time, scalability to different-sized problem instances, and generalization to different route optimization problems [89].

Attention Mechanism
Another promising technique for constructing an effective neural network to solve routing problems is through utilizing attention mechanisms. The attention mechanism is a technique that allows neural networks to focus on specific parts of the model selectively. The solution space for routing problems can become enormous, as discussed in Chapter 2. It is essential to employ a method that does not exhaustively examine the entire solution space. Utilizing the attention mechanism was already a part of the Pointer Network [13], and it has since become a common component of neural networks designed to solve routing problems [9] [11] [90]. Attention mechanisms allow the model to focus on specific parts of the problem, and further details on this will be provided in subsequent sections of the thesis.

Neural networks have made remarkable strides in solving optimization problems, from producing unsatisfactory results and requiring substantial computational resources to rivaling meta-heuristics like simulated annealing regarding solution quality. However, promising research areas still demand further exploration, specifically enhancing solution quality and optimizing computation time.

The untapped potential of neural networks in solving real-world optimization problems remains a promising research avenue. Furthermore, the under-researched application of neural networks in solving the Green Vehicle Routing Problem (GVRP) presents opportunities for advancements. This thesis introduces methods that leverage neural networks' potential and contribute to addressing these gaps.

5

Attention Mechanism

The attention mechanism is a technique used in neural networks to selectively focus important parts of the input data while reducing the influence of less relevant parts. By focusing on important parts, the model can capture underlying patterns and dependencies in the data, resulting in improved performance across various tasks. Consequently, the attention mechanism enables networks to handle complex and diverse data effectively.

The attention mechanism is a promising technique for solving routing problems, such as the Vehicle Routing Problem (VRP), due to its ability to handle complex and sequential decision-making tasks. By selectively focusing on important regions of the input data, this technique can efficiently manage large-scale and high-dimensional data encountered in the VRP. Studies have shown promising results when incorporating the attention mechanism into neural network models for solving routing problems [9] [91]. This offers a promising direction for solution methods to tackle optimization problems like the VRP.

This chapter introduces the concept of attention mechanism, discussing soft and hard attention. The Transformer architecture is then presented as an example of an architecture that utilizes the attention mechanism to its fullest potential. The chapter then delves into the application of attention in the context of graphs, discussing the concept of graph attention networks and how the attention mechanism is utilized for solving optimization problems.

5.1 Introduction to the Attention Mechanism

To effectively utilize attention mechanisms in solving optimization problems such as the VRP, it is important to grasp the fundamental concepts involved. Traditional neural networks process each input feature independently, with uniform weights applied across all features. Attention mechanisms improve on this by enabling the network to learn which parts of the input data to focus on. The attention mechanism selectively highlights important parts of the input data while minimizing the impact of less important parts, thereby allowing the network to focus on smaller but critical regions of the input data. There are various methods of implementing attention mechanisms in neural networks, and the upcoming sections will introduce some of these methods.

5.1.1 Soft Attention

Soft attention is the most common type of attention mechanism in neural networks that assigns weights to each input feature based on its relevance to the output. To achieve this, a score is assigned to each element in the input data based on the relevance to the current state of the model. The scores are then used to compute weights for each input element using a softmax function. Softmax is utilized to get a probability distribution that determines the degree of attention to be paid to each element [92].

A benefit of soft attention is that it computes a continuous distribution of attention over the input sequence, making it differentiable and allowing the entire system to be trained using standard back-propagation methods [87]. This makes it easier to optimize the model and can lead to better performance on complex deep-learning tasks.

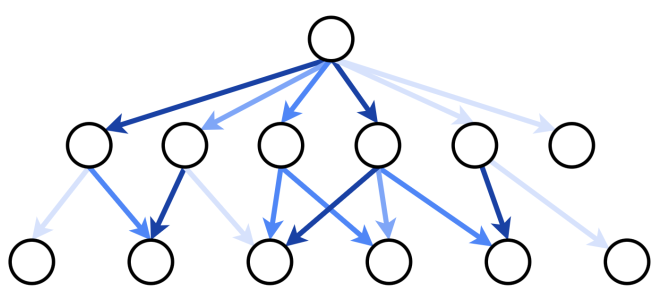

Figure 5.1: The figure illustrates soft attention, a mechanism that allows models to focus on specific parts of input data selectively. The strength of the weighting is represented by the color of the arrows connecting the nodes, with darker colors indicating higher weights.

5.1.2 Hard Attention

Hard attention is another type of attention mechanism in which the model selects a single item from the input sequence to focus on at each time step [93]. This means that the model explicitly decides which part of the input sequence to pay attention to and ignores the rest. This can be useful when the model needs to make discrete decisions, such as selecting the most relevant piece of information from a set of inputs.

However, hard attention can also be problematic because it requires the model to make a hard decision, which can be difficult to optimize using traditional gradient-based methods. Soft attention, on the other hand, allows the model to attend to multiple parts of the input sequence simultaneously by assigning a weight to each item in the sequence, which can be interpreted as the degree of the model of attention to that item.

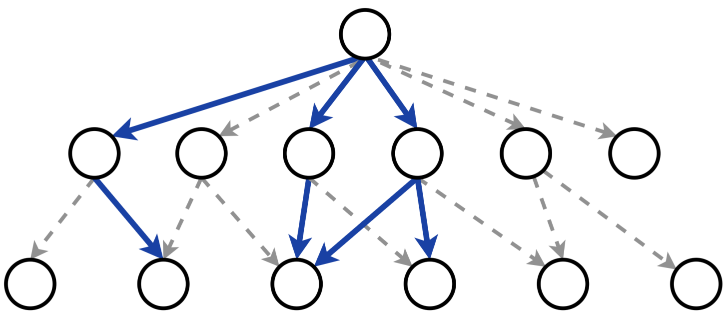

Figure 5.2: The figure illustrates hard attention, a mechanism that forces the model to choose a limited number of inputs to focus on while ignoring the others. The blue arrows represent the inputs that the model considers important, while the gray arrows represent the inputs that the model disregards.

5.1.3 Attention in Sequence-to-Sequence Problems

Sequence-to-sequence problems are a type of machine learning task where the input and output are both sequences of variable length. The goal is to learn a mapping from one sequence domain to another, and these problems are commonly addressed using neural network models. In sequence-to-sequence problems, an attention mechanism can be used to improve the ability of the model to generate a target sequence by selectively attending to relevant parts of the input sequence at each time step.

The use of attention has revolutionized the common bottleneck problem in sequence-to-sequence problems. Traditionally, the encoder is responsible for condensing the entire input sequence into a fixed-length vector, which results in a loss of information and limited capacity. This bottleneck problem can cause issues such as gradient starvation and the decoder ignoring relevant parts of the input sequence. However, by utilizing an attention mechanism, the model can selectively focus on different parts of the input sequence at different time steps, providing a flexible and targeted approach to utilizing all available information.

The attention mechanism computes a set of weights at each time step, indicating the relative importance of different positions in the input sequence. These weights can be visualized as a probability distribution over the input positions. By examining the attention weights, insights into the behavior of the model can be gained, including which parts of the input sequence it focuses on and how that attention changes over time. Attention mechanisms have become popular tools to improve the performance of sequence-to-sequence models and provide a means to interpret and understand their behavior.

5.2 The Transformers Architecture

The Transformer Architecture is a deep learning model primarily used for language tasks. It has proven to be highly effective due to its ability to utilize attention mechanisms to their fullest potential [94]. This has motivated researchers to explore its applicability beyond language tasks, including optimization problems such as the Green Vehicle Routing Problem (GVRP).

The GVRP aims to find the most efficient routes for a fleet of vehicles while minimizing their environmental impact, making it a crucial problem for sustainable transportation. The Transformer Architecture has shown promising results in solving routing problems and other optimization problems, making it a versatile tool for various fields beyond language tasks. Hence, it offers a promising but unexplored solution method for the GVRP.

5.2.1 Encoder - Decoder

In the context of the Transformers architecture, the encoder and decoder are two fundamental components of the model. The encoder is responsible for processing the input sequence and generating a set of hidden state vectors that capture the important information in the input. The input sequence is typically first embedded into a set of dense vectors, which are then processed by a series of stacked encoder layers. Each encoder layer employs an attention mechanism that enables the model to focus on different parts of the input sequence selectively. Finally, the encoder is equipped with a feedforward neural network that applies nonlinearity to the model.

The decoder, on the other hand, is responsible for generating the output sequence based on the information captured in the hidden state vectors generated by the encoder. Like the encoder, the decoder consists of a series of stacked layers. Each layer in the decoder is composed of three sub-layers that also utilize attention mechanisms so that the model can focus on different parts of the output sequence generated so far. Also, the decoder is a feedforward neural network that applies nonlinearity.

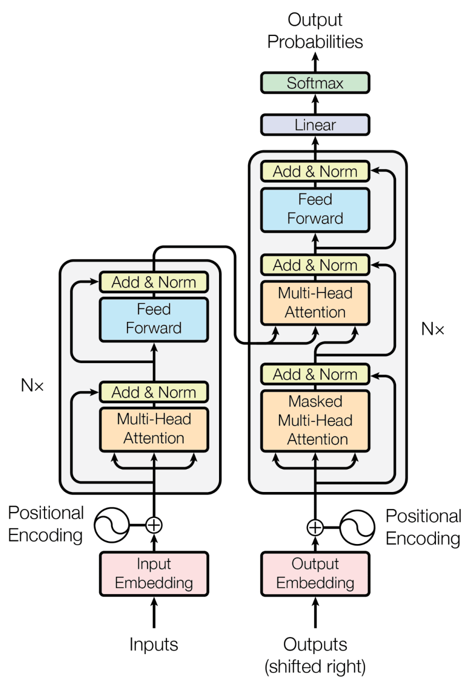

Figure 5.3: The figure illustrates the architecture of the Transformer model, which incorporates attention mechanisms in both the encoder and the decoder. These mechanisms allow the model to focus on relevant parts of the input sequence while diminishing the influence of irrelevant parts. *Figure from [94].*

5.2.2 Self-Attention

The self-attention mechanism is the key innovation of the Transformer architecture. Unlike other attention mechanisms, self-attention allows the model to capture dependencies between different parts of the input sequence, even when they are far apart. Enabling the model to learn long-range dependencies effectively. Additionally, self-attention enables the Transformer to process all input elements in parallel, making it more efficient for longer sequences. This is achieved by mapping the input sequence to itself and using three key elements: queries, keys, and values.

Queries, Keys, and Values

In the Transformer architecture, the self-attention mechanism relies on three key elements: queries Q, keys K, and values V. These elements work together to determine the relevance of each input element with respect to every other element in the sequence.

Queries play a role in capturing information from the input sequence. They serve as the basis for asking questions about the input sequence. In the self-attention mechanism, each query is used to calculate a score, indicating the significance of each element in the sequence relative to that particular query. Keys represent the different elements in the input sequence that the queries are trying to learn from. The keys are used to determine how well each element in the sequence matches the query. Values are the final input vector in the self-attention calculation. They represent the information that the self-attention mechanism should use to construct the output sequence.

Scaled dot-product Attention

Scaled dot-product attention is a fundamental component of the Transformer architecture, specifically used for self-attention. It calculates the similarity between queries and keys, providing a measure of relevance. This similarity score determines the weight assigned to each value, indicating its importance in the context of the queries. To calculate the similarity between queries and keys, scaled dot-product Attention uses a dot product between the query vector and the key vector. In scaled dot-product attention, the dot product between the query vector and key vector is computed, and the result is scaled down by the square root of the dimension of the key vector. This scaling is done to prevent the dot product from growing too large or too small, which can make the softmax function unstable. The scaled dot-product attention mechanism, denoted as $Attention(Q, K, V)$, is computed as:

$$Attention(Q, K, V) = softmax(\frac{QK^T}{\sqrt{d_k}})V, \tag{5.1}$$

where Q, K, and V represent the query, key, and value matrices, respectively.

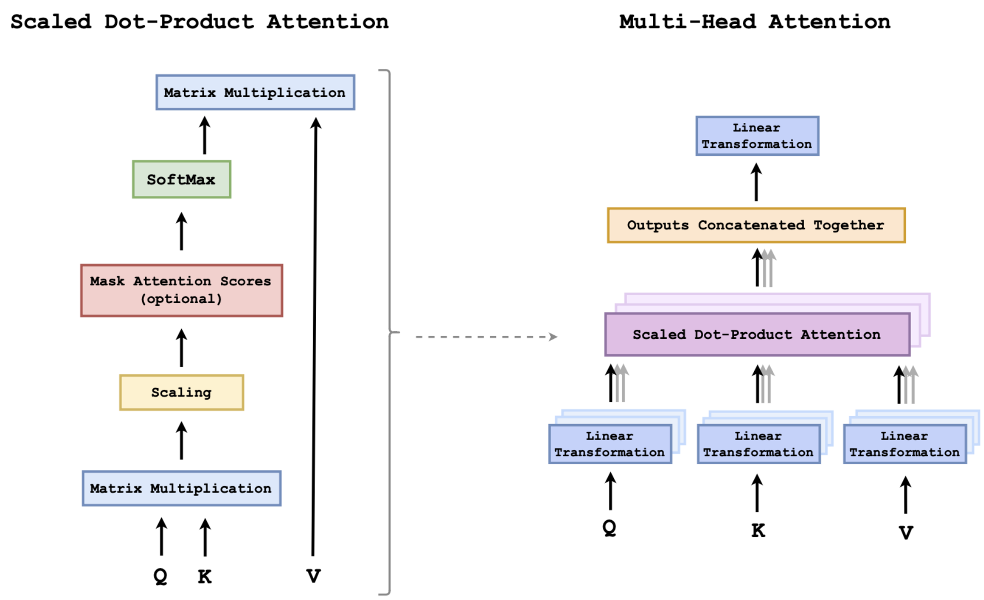

Figure 5.4: Left: Illustration shows the scaled dot-product attention mechanism, which computes the attention scores between the input sequence and the query vector. The attention scores are used to weigh the value vectors and produce the output.
Right: Illustration shows the multi-head attention mechanism, which splits the input sequence into multiple heads and applies a scaled dot-product attention mechanism to each head.

5.2.3 Multi-head Attention

Multi-head Attention is an extension of the attention mechanism that allows for capturing different relationships between different parts of the input sequence. It allows the model to focus on different aspects of the input sequence simultaneously. It consists of multiple attention heads that operate in parallel, each focusing on a different part of the input sequence. Each attention head takes in the same query, key, and value inputs but learns a different set of parameters through training. This schematic is illustrated in Figure 5.4.

The output of each attention head is then concatenated and passed through a linear layer to produce the final output. This allows the model to capture multiple aspects of the input sequence and learn different representations that can be combined to improve performance. Overall, Multi-head Attention ability to capture multiple aspects of the input sequence through parallel attention heads and potential to improve model performance.

5.3 The Attention Mechanism on Graph Structured Data

The attention mechanism is well-suited for graph-structured data due to the variable number of nodes and the need to consider permutation and variance. This is especially relevant in routing problems, which are often represented as graphs, where the attention mechanism can be used to improve accuracy and performance. Graph-structured problems involve data represented as a graph, with nodes representing entities and edges representing relationships between entities. Examples of graph-structured problems include social network analysis, molecular graph generation, and when working with geographical maps.

To process graph-structured data, a Graph Attention Network (GAT) can be used. GATs are a type of neural network that applies the attention mechanism to graph data, and they have been shown to be effective in a variety of graph-structured problems. To fully understand GATs, it is important to have a basic understanding of graph theory and Graph Neural Network (GNN). Graph theory is the study of graphs and their properties, while graph neural networks are a class of neural networks that can process graph-structured data. By combining the principles of graph theory and neural networks, GATs offer a powerful tool for processing graph-structured data.

5.3.1 Graphs Theory

Graph networks are built on the theory of graphs, which is a mathematical discipline that studies structures composed of nodes and edges, as mentioned in Chapter 2. Graphs are used to represent relationships between objects or entities. Graph theory provides the essential concepts and tools to understand and apply GNNs, which is a rapidly growing application of graph theory.

In GNNs, nodes in the graph represent entities, such as individuals, genes, or cities, and edges represent the relationships between these entities, such as distance or cost. A node feature vector is a vector that represents information about a node in a graph. It is a set of features associated with each node in the graph, as illustrated in Figure 5.5.

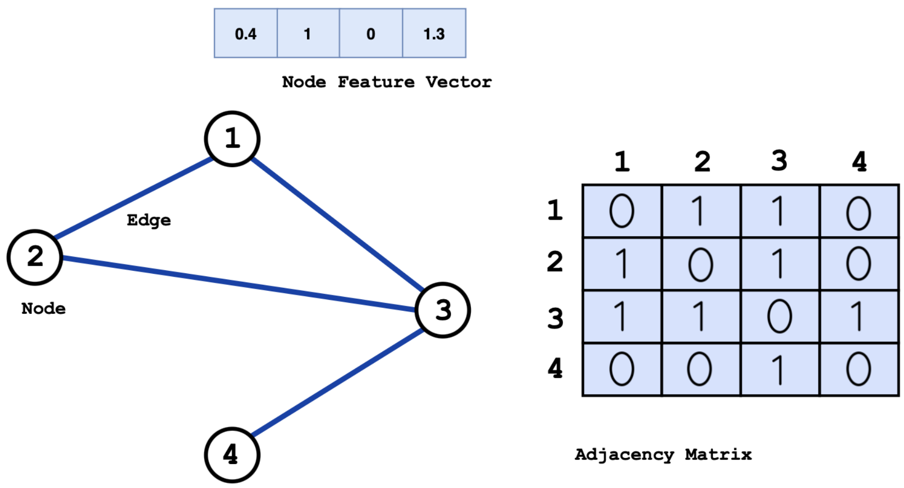

Figure 5.5: The figure illustrates a graph, its adjacency matrix, and a node feature vector. The graph is represented by nodes and edges connecting them. The adjacency matrix indicates the connections between nodes in the graph. The node feature vector is a vector containing information about each node in the graph.

An adjacency matrix is a fundamental tool in GNNs, as it provides a matrix representation of the connections between nodes in a graph, indicating which nodes are connected and helping to understand the relationships between entities. Understanding fundamental graph theory concepts, such as nodes, edges, and adjacency matrices, is essential to comprehend the workings of GNNs.

5.3.2 Graph Neural Networks

Before going into the use of attention mechanisms in GATs, it is important first to explore the foundations of GNNs. GATs are a specific type of GNN, which are neural networks designed to process graph-structured data. Hence, GNNs operate on data that can be represented as a graph, with nodes representing entities and edges representing relationships between them.

The main process of GNNs is to update the embeddings of individual nodes. Node embeddings are a type of representation learning method that is commonly used in graph neural networks. They are a way of transforming high-dimensional node feature vectors into a lower-dimensional space while preserving the essential information about the nodes and their relationships in the graph.

Node embeddings are learned during the training process of a graph neural network. In this process, each node feature vector is passed through a single-layer neural network that is applied only to the node's neighbors, and the resulting embeddings are summed up to form a single-node embedding.

$$h_i^{`} = \sigma(\sum_{j \in N(i)} w h_j) \tag{5.2}$$

Equation 5.2 represents the process of updating the embedding of a single node i, where $j \in N(i)$ represents the set of neighboring nodes of i. The new embedding $h_i^{`}$ is computed by taking a weighted sum of the embeddings of the neighboring nodes h_j using learnable weights w and passing the result through a non-linear activation function σ.

5.3.3 Graph Attention Networks

Now that we have established the basics of graph theory and GNNs, we can extend our understanding to Graph Attention Network (GAT)s. GATs are a type of GNN that implements an attention mechanism to assign different importance weights to neighboring nodes when updating a node's embedding [95].

In GATs, the attention mechanism is used to learn the importance of each neighbor node j to node i. The goal is to determine which nodes to prioritize when updating the node embeddings during training. The importance of each node is expressed through an attention coefficient, which assigns different weights to each neighboring node based on its relevance to the target node.

Attention Coefficient

The attention coefficient is a key concept in GATs and refers to the weight assigned to each neighboring node based on its relevance to the target node. It measures how much attention should be paid to each neighboring node when updating the embedding of a target node during the training process. The attention coefficient is calculated as:

$$e_{ij} = a(wh_i, wh_j), \tag{5.3}$$

where h_i and h_j are the feature vectors of nodes i and j, and w is the weight matrix. w is a learnable transformation of the node features. a is the attention mechanism used to compute the attention coefficient. There are multiple ways to calculate or learn the attention coefficient a, and one common approach is using a shared single-layer neural network. It takes as input the transformed feature vector of node i and j and returns a value for the importance of the features of node j, for node i.

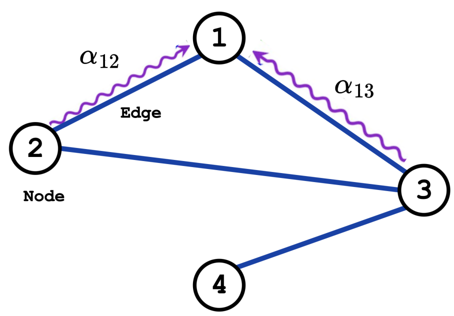

Figure 5.6: This figure illustrates attention coefficients in a Graph Attention Network. The purple lines correspond to the attention coefficients α_{12} and α_{13}, indicating the importance of the connections between node 1 and its neighboring nodes 2 and 3, respectively. These coefficients play an important role in capturing relevant information and guiding the information flow in the network. *The figure is inspired by [95].*

When calculating attention coefficients, the resulting values can vary significantly across nodes. To ensure comparability, it is essential to normalize these values. This is commonly achieved through the use of the softmax, as illustrated by the following equation:

$$\alpha_{ij} = softmax(e_{ij}) = \frac{exp(a(wh_i, wh_j))}{\sum_{k \in N(i)} exp(a(wh_i, wh_j))}. \tag{5.4}$$

Thus, the node update process for a graph attention network can be expressed as [95]:

$$h_i^{'} = \sigma(\sum_{j \in N(i)} \alpha_{ij} w h_j), \tag{5.5}$$

where the attention coefficients α_{ij} are used to assign different weights to each neighbor node j of the target node i. The weights are learned through the attention mechanism and used to update the node embedding based on the features of its neighboring nodes. The learnable weight matrix w is used to transform the neighboring node features, and the activation function σ applies non-linearity to the output.

In recent years, GATs have emerged as a prominent approach in addressing diverse challenges in graph-based learning [96] [95]. They have demonstrated their effectiveness in various domains, such as social networks [97], biology [98], and transportation [99]. By leveraging the inherent structure of graphs, GATs provide a powerful neural network architecture that is specifically designed for processing graph-structured data.

5.4 Attention Mechanism for Solving Routing Problems

Neural networks with attention mechanisms, such as GATs, are suitable for solving routing problems, such as VRPs, due to their ability to effectively model relationships between nodes in a graph. In VRPs, where nodes represent locations and edges represent distances between them, GATs are good at capturing the intricate dependencies among nodes.

By learning to assign varying weights or attention coefficients to neighboring nodes, GATs can identify critical nodes that play a significant role in optimizing the routing process. This enables efficient and accurate route planning, considering factors such as distance, traffic conditions, and delivery constraints. The unique combination of graph-based modeling and attention mechanisms makes neural networks with attention particularly well-suited for addressing routing problems and holds promising potential for operational research applications.

This thesis explores the application of attention-based neural networks in the field of operational research, with a specific focus on solving the VRP, which is characterized by its graphical nature. The proposed network shares similarities with the GAT but incorporates an attention mechanism and a model structure inspired by the Transformers architecture. In the upcoming methodology section, we will go into these ideas and provide a detailed explanation of the attention-based neural network's utilization for addressing the GVRP within waste management.

Part II

Method & Data

6

Attention-based Neural Network for Sustainable Routing in Waste Management

Expanding on the theoretical foundation, this chapter presents a method for achieving more sustainable waste collection routes using an attention-based neural network. Although neural networks have shown promising results in solving routing problems for simulated data, their application in real-world scenarios with environmental concerns is still limited. Therefore, this thesis contributes to the current research on the Green Vehicle Routing Problem (GVRP) by addressing a real-world scenario within waste management.

This chapter first provides an overview of the method employed to generate sustainable waste routing. It then goes into the GVRP model and how it is tailored to address waste management challenges. The subsequent section introduces the attention-based neural network, offering insights into the model structure and learning process. To illustrate the practical application of the model, a real-world scenario involving the waste management company Remiks is presented. This example demonstrates how the proposed method is valuable for the company to generate sustainable collection routes.

6.1 Waste Collection Route Optimizing using Attention-based Neural Network

While neural networks have shown promising results in solving routing problems [9] [10] [14], their application in real-world scenarios, particularly those involving environmental considerations, has been limited. Consequently, there is a need to explore their potential in addressing real-world problems. One field that provides an interesting avenue for exploring the practical applications of this method is the transport sector, particularly in waste management.

When examining the waste management system in Norway today, it is evident that there are several impressive operations in place. For instance, innovative waste sorting technologies, such as optical sorting equipment, are utilized to handle the collected waste efficiently. Additionally, many waste management companies use waste-to-energy technologies, which can convert waste into energy, such as electricity or heat. These technologies play a crucial role in reducing emissions by diverting waste from landfills and decreasing reliance on fossil fuel-based energy sources.

However, the daily waste collection procedure still remains a significant source of emissions due to inefficient operations and the use of diesel-powered vehicles. These inefficiencies result in high transportation costs and poor resource utilization. Therefore, it is critical for the waste management industry to address emissions from waste collection operations to promote more sustainable practices and reduce their carbon footprint.

To address emissions associated with waste collection and promote more sustainable waste management practices, various strategies can be employed. This thesis focuses on one such strategy: route optimization. By optimizing waste collection routes, we can effectively reduce emissions during the collection process. Route optimization offers a cost-effective alternative to other approaches, such as investing in electric vehicles, which often require substantial initial capital investments. In this thesis, we propose utilizing an attention-based neural network [9] to achieve route optimization, leveraging the synergies between machine learning and operations research.

The attention-based neural network utilized in this approach builds upon the transformer architecture and comprises an encoder and a decoder. This network takes a set of coordinates as input and generates an optimized solution route as output. By leveraging the capabilities of the attention mechanism, the neural network can effectively learn to allocate attention to relevant locations and make informed routing decisions.

To learn and improve these solution routes, the proposed method in this thesis employs reinforcement learning, specifically utilizing a variant called REINFORCE with baseline. This approach is well-suited for route optimization since it eliminates the need for a preexisting training dataset since the solution routes can be generated on the fly. In addition, it enables iterative improvements to the routes by leveraging feedback received during the reinforcement learning process. This iterative nature allows the model to continuously improve its routing decisions, ultimately leading to more optimized and efficient routes.

Furthermore, to address environmental concerns within the routing model, we employ [100] to incorporate fuel consumption as the objective in the optimization process. The primary goal of the optimization model is to minimize fuel consumption. In the case of Tromsø, we found that road gradient plays a significant role in fuel consumption due to the substantial elevation differences in the area. Figure 6.1 provides an overview of the proposed model. Subsequent sections will provide more detailed information on these elements of the method.

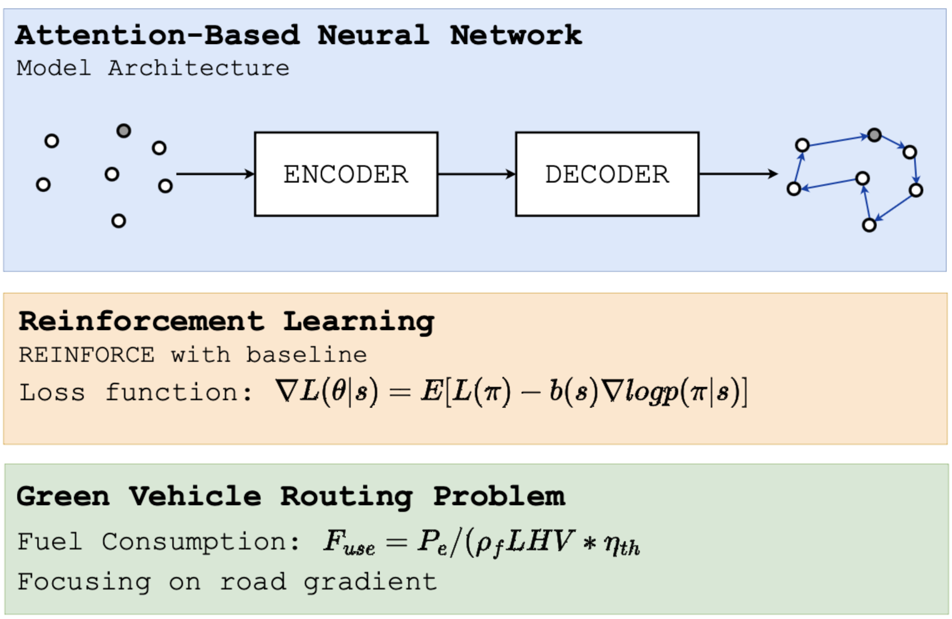

Figure 6.1: This figure presents an overview of the proposed model for route optimization. The core component of the model is an attention-based neural network, which serves as the solution method for optimizing routes. The network utilizes reinforcement learning to iteratively improve the solution routes over time. Notably, the model incorporates environmental considerations by integrating green vehicle routing principles, with a focus on minimizing fuel consumption as a cost metric for the routes. This approach ensures that the optimization process takes into account the environmental impact, aligning with sustainable practices in waste management.

6.2 The Green Vehicle Routing Problem for Optimizing Waste Collection Routes

This thesis aims to optimize waste collection routes by specifically targeting environmental concerns through the incorporation of the GVRP. The GVRP is a variant of the traditional Vehicle Routing Problem (VRP) that focuses on minimizing the environmental impact of transportation. It involves finding the most efficient routes for a fleet of vehicles to serve a set of customers while considering various environmental factors. Despite the growing importance of sustainable transportation, research on the GVRP, especially in real-world applications, remains limited.

Utilizing the GVRP for optimizing waste collection routes to reduce emissions provides a practical solution for waste management companies to achieve their environmental goals. However, one of the challenges related to the GVRP is how to incorporate environmental aspects into the problem, as discussed in Chapter 3. Therefore, exploring various modeling techniques for solving the GVRP remains an active area of research.

Applying the GVRP to optimize waste collection routes and minimize emissions offers a practical solution for waste management companies aiming to meet their environmental objectives. However, a significant challenge in the context of the GVRP lies in effectively integrating environmental considerations into the problem formulation, as discussed in Chapter 3.

We address this challenge by investigating the integration of fuel consumption as an environmental consideration in optimizing waste collection routes. Recognizing the relationship between fuel consumption and vehicle emissions, our approach focuses on reducing fuel usage to minimize emissions from the fleet of vehicles. To incorporate fuel consumption into the GVRP, a representation of fuel consumption for the waste collection vehicles needs to be established. This section will go into further detail regarding the methodology employed to incorporate and address fuel consumption within the GVRP framework.

6.2.1 Modeling Fuel Consumption

Fuel consumption is closely related to emissions, as the burning of fossil fuels is a significant source of vehicle emissions. Hence, minimizing fuel consumption is an effective way to reduce emissions from a fleet of vehicles running on fossil fuels. Many studies have investigated this approach [63] [101] [100]. However, accurately defining fuel consumption is still challenging in this area.

This thesis employs the power delivered by the vehicle engine to calculate fuel consumption [100], as fuel consumption is directly related to engine power. The engine power needs to overcome various resistive forces, including air resistance, rolling resistance, and incline resistance. The intensity of these forces is influenced by factors such as vehicle speed [101], mass [102], and characteristics of the terrain [100]. These factors collectively determine the magnitude of force required to counter the resistive forces and ultimately impact fuel use.

Equation 6.1 calculates the engine power P_e, taking into account several factors that impede a vehicle's movement [100]. This equation adds up the power associated with each force and divides the result by the mechanical efficiency of the powertrain, η_t. The total energy required to overcome these forces represents the energy consumed by the vehicle.

The equation considers rolling resistance F_r, drag force F_d, gravity F_g, and inertial forces as the main factors that impede the vehicle's motion. By considering these factors, we can accurately estimate the amount of energy required to move the vehicle and thus calculate its fuel consumption.

Engine Power:

$$P_e = \frac{(F_d + F_r + F_g + MM_{fi}a)V}{\eta_m} \tag{6.1}$$

Drag Force: $$F_d = \frac{1}{2}c_d\rho_a A_f V^2 \tag{6.2}$$

Rolling Force: $$F_r = C_r Mgcos\theta \tag{6.3}$$

Gravity: $$F_g = Mgsin\theta \tag{6.4}$$

Inertial Force: $$M_{fi} = 1 + 0.04N_{TDi} + 0.0025N_{TDi^2} \tag{6.5}$$

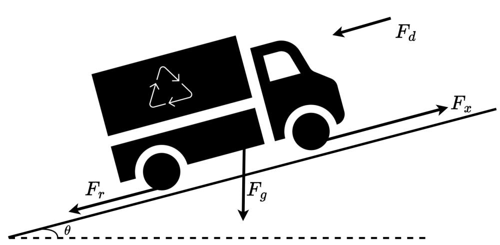

Figure 6.2: This figure is a visual representation of the physical forces that come into play when vehicles are traveling on the road with an incline. These forces highlight the factors that must be taken into consideration when estimating the fuel consumption of a vehicle. *The figure is inspired by [100].*

To estimate vehicle fuel consumption based on engine power, the method used in this thesis employs Equation 6.6. This equation utilizes the engine power P_e calculated from Equation 6.1, which sums the power from each force that opposes the vehicle's movement. Other variables in the fuel consumption equation include the fuel density ρ_f, lower heating value (LHV) of the fuel, and the engine thermal efficiency η_{th}.

By dividing the power delivered by the engine by the product of the fuel density, lower heating value, and the engine thermal efficiency, Equation 6.6 provides an estimate of the rate at which the engine consumes fuel. By utilizing this estimated fuel consumption, the waste collection routes in this thesis are optimized accordingly.

$$\textbf{Fuel Consumption} = F_{use} = P_e/(\rho_f LHV * \eta_{th}) \tag{6.6}$$

When fuel consumption is chosen as the environmental criterion in the GVRP and an equation is derived for it, it becomes possible to construct a suitable GVRP model to optimize waste collection routes and minimize fuel consumption. This approach enables waste collection companies to reduce their emissions in their day-to-day operations.

The objective of the GVRP used in this method can then be expressed by Equation 6.7, where E_{ij} represents the fuel consumption between nodes i and j, x_{ij} represents a binary decision variable indicating whether or not a vehicle travels from node i to node j. F_{ij} represents the fuel used for travel between nodes i and j, estimated using Equation 6.6. The objective is to minimize the total fuel consumption, which is equivalent to the sum of $E_{ij}x_{ij}$ over all nodes.

$$min \quad \sum_{i,j \in A} E_{ij}x_{ij} = F_{ij}x_{ij} = (P_e/(\rho LHV * \eta_{th}))x_{ij} \tag{6.7}$$

The proposed objective function optimizes the environmental concerns as the primary objective. However, a multi-objective approach can be beneficial compared to the proposed objective function, as it provides a more comprehensive approach. Introducing distance as a second objective can bring benefits beyond reducing fuel consumption, such as reducing wear and tear on roads, emissions from vehicles, and noise pollution.

Additionally, including multiple objectives can offer a more balanced and flexible optimization approach producing more realistic routes. By considering both fuel efficiency and route feasibility, a multi-objective approach can lead to more efficient and environmentally friendly waste collection routes. The two objectives used in this thesis are expressed as

$$min \quad \alpha \sum_{i,j \in A} E_{ij} x_{ij} + \beta \sum_{i,j \in A} c_{ij} x_{ij}, \tag{6.8}$$

where the first sum minimizes emissions and the second sum minimizes distance.

In multi-objective optimization, it is necessary to determine the relative importance of each objective being optimized. The influence of each objective is scaled by the weighting factors α and β. These factors allow the decision-maker to determine the trade-off between the objectives and to adjust the relative importance of each objective according to the problem at hand.

The values of α and β should be chosen based on the specific needs and requirements of the problem being solved. This may involve some trial and error to ensure that the chosen values are appropriate and lead to a satisfactory solution. By adjusting the weighting factors, the decision-maker can find the optimal balance between fuel consumption and distance and ultimately arrive at a more balanced and realistic solution for the problem.

6.3 Attention-Based Neural Network

In this thesis, we present a novel method for solving the GVRP problem by using an attention-based neural network [9]. While this approach has been successfully applied to address the VRP using artificially generated data, its adaptability, and effectiveness in real-world scenarios, particularly in the context of the GVRP, remain unexplored.

The approach adopted in this thesis utilizes an attention-based model that follows the transformer architecture [94]. The model consists of an encoder and a decoder, as illustrated in Figure 6.3. To optimize the loss function and enhance the model's performance, reinforcement learning is employed. This reinforcement learning technique enables the neural network to learn from its own actions and progressively improve over time. In this section, we introduce the encoder-decoder architecture and elaborate on the reinforcement learning approach.

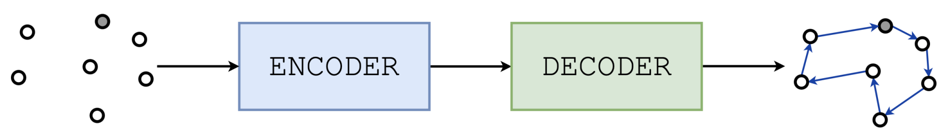

Figure 6.3: This figure provides a simplified illustration of the attention-based neural network. The input consists of coordinates, which are processed through the encoder and decoder. The decoder generates a solution route as the output, depicting the optimized route for the given input coordinates.

6.3.1 Encoder

The encoder is responsible for processing the input data and generating embeddings for all input nodes. These node embeddings serve as inputs for the decoder, which subsequently generates the output of the model. Figure 6.4 provides a visual representation of the entire encoding process [9].

The encoder takes the input features x_i, which usually represent the two coordinates of a node, and computes the initial node embeddings:

$$h_i^{(0)} = W^X x_i + b^X. \tag{6.9}$$

The encoder then updates these initial embeddings through N attention layers. At each layer $\ell \in 1, ..., N$, a new node embedding h_i^ℓ is generated by the encoder. After computing the node embeddings at the final attention layer, the encoder generates an aggregated embedding $\bar{h}^{(N)}$ that captures essential information from all node embeddings. This is achieved by computing the mean of the final node embeddings:

$$\bar{h}^{(N)} = \frac{1}{n} \sum_{i=1}^{n} h_i^{(N)}. \tag{6.10}$$

Consequently, the encoder outputs the node embeddings $h_i^{(N)}$ and the graph embedding $\bar{h}^{(N)}$, which together represent the essential information about the input graph for further processing by the decoder.

The Attention Layer in The Encoder
The attention layer in the encoder consists of two sublayers: a multi-head attention layer and a node-wise fully connected feed-forward layer. Both sub-layers have skip connections and batch normalization for improved training stability.

Sub-Layer 1: Multi-Head Attention
The first sub-layer of the attention layer in the encoder is the multi-head attention layer. This layer takes the input node embeddings $h_i^{(\ell-1)}$ from the previous layer and applies the multi-head attention mechanism to obtain an updated embedding:

$$\hat{h}_i = BN^\ell(h_i^{(\ell-1)} + MHA_i^\ell(h_1^{(\ell-1)}, ..., h_n^{(\ell-1)})). \tag{6.11}$$

The multi-head attention mechanism aggregates information from all nodes to enhance the representation of each node.

Sub-Layer 2: Fully Connected Feed-Forward Layer
The second sub-layer of the attention layer in the encoder is a fully connected feed-forward layer. This layer takes the aggregated node embedding from the first sub-layer as input, applies a fully connected feed-forward layer, and adds the result to the input. Finally, skip connections and batch normalization is applied to the output:

$$h_i^{(\ell)} = BN^\ell(\hat{h}_i + FF^{(\ell)}(\hat{h}_i)). \tag{6.12}$$

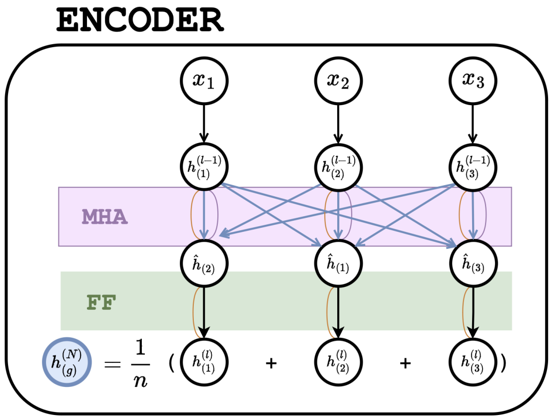

Figure 6.4: This figure illustrates the encoder process for three nodes, starting with the input node features, which are processed to generate initial embeddings, followed by the multi-head attention layer and the fully connected feed-forward layer. The output of the encoder consists of the node embeddings for all nodes and the graph embedding, which both serve as the input to the decoder. *The figure is inspired by [9].*

6.3.2 Decoder

The main function of the decoder is to generate the optimal route for a given routing problem by producing a sequence π of the input nodes. To accomplish this, the decoder takes the node embeddings $h_i^{(N)}$ and graph embedding $\bar{h}^{(N)}$ generated by the encoder as input.

To generate the solution route, the decoder computes context nodes. These context nodes represent the decoding context, which can contain information about the previously generated output nodes. To compute the context node, three elements are used as input: the graph embedding $\bar{h}^{(N)}$, the node embedding for the first node $h_{\pi_1}^{(N)}$, and the previous node embedding $h_{\pi_{t-1}}^{(N)}$.

When computing the context node for the first node, where there is no previous node, learned parameters v and v^f are used as input placeholders. The first context node embedding is then computed using a horizontal concatenation operation of these three elements:

$$h_{(c)}^{(N)} = \begin{cases} [\bar{h}^{(N)}, v^1, v^f] & t = 1 \\ [\bar{h}^{(N)}, h_{\pi_{t-1}}^{(N)}, h_{\pi_1}^{(N)}] & t > 1. \end{cases} \tag{6.13}$$

Attention Mechanism in The Decoder
The decoder uses the attention mechanism to update the context node embeddings. This is accomplished through multi-head attention. To compute the queries, keys, and values for the attention mechanism, the node embeddings h_i and the context node embedding $h_{(c)}$ are multiplied by learned weight matrices W^Q, W^K, and W^V, respectively:

$$q_{(c)} = W^Q h_{(c)}, \quad k_i = W^K h_i, \quad v_i = W^V h_i. \tag{6.14}$$

When queries, keys, and values are obtained, the next step is to compute the compatibility of the query with all nodes:

$$u_{(c)j} = \begin{cases} \frac{qk}{\sqrt{d_k}} & \text{if} j \neq \pi_{t'} \quad \forall t' < t \\ -\infty & \text{otherwise.} \end{cases} \tag{6.15}$$

This compatibility is used to compute attention weights, which determine the importance of each node for the current decoding step. The compatibility is computed using the scaled-dot product. During this process, nodes that have already been visited are masked out.

Probabilities
Finally, the probability of choosing the next node at each time step is obtained by adding a final decoder layer with a single attention head. This layer takes the computed compatibility and applies a tanh function to clip the results within a specific range:

$$u_{(c)j} = \begin{cases} C \cdot \tanh(\frac{qk}{\sqrt{d_k}}) & \text{if} j \neq \pi_{t'} \quad \forall t' < t \\ -\infty & \text{otherwise.} \end{cases} \tag{6.16}$$

In this thesis, the resulting scores are clipped to a range between -10 and 10 to stabilize the gradients during backpropagation and improve the overall training performance. If the scores are too large or too small, it can result in unstable gradients and hinder learning. Therefore, clipping helps to ensure that the gradients are within a reasonable range for efficient and effective training.

After clipping, the masked compatibilities are normalized using the softmax function to obtain probabilities for choosing the next node. These probabilities are used in the loss function to update the model parameters. The whole process of the decoder is shown in Figure 6.5, which illustrates this process for a TSP with three nodes.

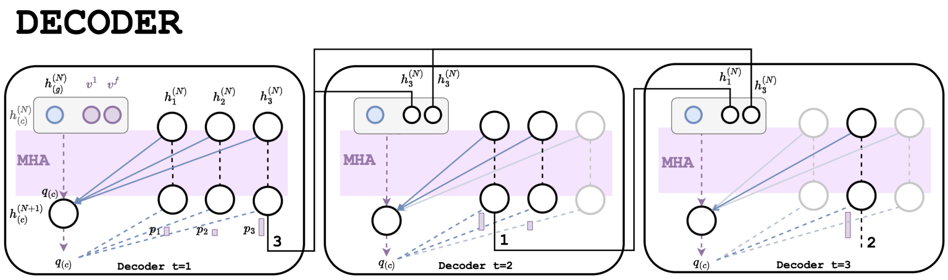

Figure 6.5: This figure demonstrates the computation process in the decoder for a Travelling Salesman Problem with three nodes. The decoder takes as input the graph embedding and the node embedding for the first node, both computed in the encoder. Additionally, the node embedding from the previous node is used as input for computing each context node. For the first context node, input placeholders are used instead, illustrated in purple. The decoder constructs a solution route $\boldsymbol{\pi} = (3, 1, 2)$, which can be seen in the output of each time step. *The figure is inspired by [9].*

6.3.3 Reinforcement learning

Our method utilizes reinforcement learning to train and improve the generated solution routes iteratively. Specifically, we employ the REINFORCE algorithm with a baseline, training the model through gradient descent. This approach enables the model to optimize its performance over time.

Reinforcement learning was chosen as the approach due to its demonstrated potential in addressing routing problems [9] [17] [89]. Its ability to learn from experience and adapt to changing conditions makes it well-suited for such problems. One advantage of reinforcement learning is that it eliminates the need for a training dataset, which can be challenging to obtain in traditional supervised learning methods. Unlike conventional approaches, the reinforcement learning methodology iteratively improves the solution through feedback. This iterative learning process allows our model to improve its solution route iteratively.

When applied to VRPs, the reinforcement learning technique learns through the process of making decisions regarding route selection and receiving rewards based on the performance of the chosen solution route. By leveraging this feedback, the model dynamically adjusts its behavior to optimize its performance over time. The efficacy of this approach has been demonstrated in various recent studies focused on solving VRPs [89] [9] [17].

Loss Function
A reinforcement learning approach is employed to train the attention-based neural network in this thesis, specifically utilizing a variant known as REINFORCE with a baseline. The loss function employed in this approach is defined as follows:

$$\nabla L(\theta|s) = \mathbb{E}p\theta(\boldsymbol{\pi}|s)\left[(L(\boldsymbol{\pi}) - b(s))\nabla\theta \log p_\theta(\boldsymbol{\pi}|s)\right]. \tag{6.17}$$

This loss function calculates the gradient of the loss concerning the model parameters θ, given a specific problem s. The expectation is taken over the policy distribution $p_\theta(\boldsymbol{\pi}|s)$, and the term $(L(\boldsymbol{\pi}) - b(s))$ gives the difference between the cost of the generated solution route $\boldsymbol{\pi}$ and a baseline value $b(s)$.

REINFORCE with Baseline
In this thesis, we employ the REINFORCE with a baseline algorithm. Using a baseline has been demonstrated to reduce gradient variance and enhance training speed in reinforcement learning scenarios [74]. We explore two different baselines: the greedy rollout baseline and the exponential moving average baseline. An analysis of their performance in the waste collection routing application is presented in Chapter 8.

The Exponential Moving Average Baseline
One of the baselines tested in this thesis is an exponential moving average. In this approach, the baseline value, M, is initially set to $L(\boldsymbol{\pi})$ in the first iteration. In subsequent iterations, it is updated using the formula:

$$M \leftarrow \beta M + (1 - \beta)L(\boldsymbol{\pi}), \tag{6.18}$$

where β represents the decay rate [9]. The decay rate β controls the influence of previous iterations on the updated baseline value. A higher decay rate puts more weight on recent iterations, allowing the baseline to adapt more quickly to changes in the solution quality. On the other hand, a lower decay rate gives more importance to previous iterations, allowing the baseline to capture longer-term trends.

The Greedy Rollout Baseline
The other baseline explored in this thesis is the greedy rollout baseline. It involves performing simulated rollouts from the current state and selecting actions greedily to maximize immediate rewards by choosing the node with the maximum probability at each decoding step without considering long-term consequences or exploring alternative paths [103]. While this approach offers computational efficiency, it can overlook better long-term solutions.

/7 Data

In the application of the proposed methodology to Remiks' waste collection operations, relevant data was gathered from various sources. This involved retrieving fuel consumption, and distance traveled data from Remiks' website, joining waste collection vehicles to gather real-time information on speed, load, and road gradient, and acquiring specific vehicle data and other relevant factors for precise fuel consumption calculations. The comprehensive analysis of this data will enable a more accurate analysis of the potential benefits associated with the proposed approach for Remiks.

This section begins by introducing the waste management company Remiks and examining their current operations. It then delves into the factors that influence fuel consumption, such as speed, load, and road gradient. Next, we provide more detailed information on vehicle-specific data and other relevant data. Finally, we narrow down the scope by selecting a specific focus area for conducting experiments.

7.1 Real-World Application: The Waste Management Company Remiks

The approach presented in this thesis introduces a novel combination of the GVRP model and the proposed attention-based neural network solution method. This integrated approach is applied to optimize waste collection operations for Remiks, a waste management company in Tromsø. This practical application serves to determine the applicability and effectiveness of the proposed method and contribute to the limited research on GVRP for real-world applications.

Remiks offers waste collection services in Tromsø and Karlsøy, focusing on household waste collection while ensuring efficient and sustainable waste management practices. Their main goals are to facilitate a well-functioning waste system and to be a driving force in reducing emissions through their values of environmental awareness and innovation [104].

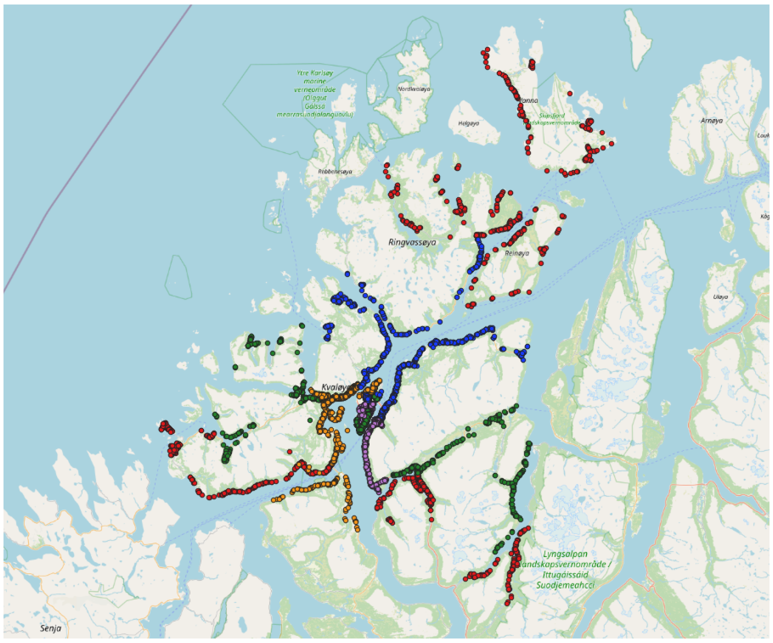

Figure 7.1: This figure shows a map of the Tromsø and Karlsøy area, highlighting the households that require waste collection services from Remiks. The waste collection trucks need to cover a significant area to collect waste from all households in the region. The current daily collection procedure generates significant emissions and poses challenges in terms of transportation costs.

While Remiks has integrated innovation and sustainability into their daily practices, including the implementation of waste-to-energy technology for district heating and the utilization of innovative waste sorting technologies such as optical sorting, the routine waste collection process contributes to significant emissions due to the reliance on diesel-powered vehicles. This poses a challenge for Remiks in their pursuit of emission reduction while striving to maintain efficient and cost-effective waste collection services. As shown in Figure 7.1, there are many households in Tromsø and Karlsøy that require waste collection services from Remiks, making it crucial for the company to address their emissions and achieve their goals of sustainable waste management practices.

Figure 7.2: The left image illustrates Remiks' vehicle fleet consisting of diesel-powered waste collection trucks. The right image showcases their innovative optical sorting technology, which plays a crucial role in their waste processing operations.

The proposed GVRP model, combined with the attention-based neural network solution method, provides Remiks with a promising approach to optimize their waste collection routes and work towards their emission reduction goals. The solution derived from the model has the potential to reduce emissions and transportation costs for their operations, providing Remiks with a competitive advantage.

Furthermore, implementing this approach gives Remiks the opportunity to have a deliberate approach to how they operate their waste collection routes. It gives a more proactive decision-making process that aligns with their goals rather than relying on the individual driver's subjective choices. Additionally, this approach enables the training of future generations of drivers without being solely dependent on the expertise of senior drivers who may no longer be available in the future.

To apply this method to the waste collection routes of Remiks, the next chapter presents data and insights on Remiks operations, along with other relevant information, to explore how the fuel-minimizing model can be utilized for route generation in this particular context.

7.2 Current Operations

Remiks has access to data on households requiring waste collection services, which can be geocoded and analyzed using QGIS. Figure 7.1 provides a geographic visualization of this data, presenting a full overview of the households that the company is obligated to collect waste from. Additional details about households on each collection day can be found in the Appendix Figures 9.1, 9.2, 9.3, 9.4, and 9.5.

Remiks operates five side loaders to service a designated area during weekdays. Side loaders are waste collection vehicles equipped with a mechanical arm, enabling efficient bin emptying by a single operator. The project thesis [19] investigated the different routes, uncovering variations in fuel consumption. Figure 7.3 emphasizes that Vehicle 4 consumed less fuel, predominantly servicing routes within the district area. In contrast, vehicles primarily operating in the city area, notably Vehicle 5, showed higher fuel consumption.

The existing route system was originally established several years ago and has undergone ad-hoc expansions to accommodate the city's growing developments. This presents an opportunity to identify more efficient routes that can better optimize waste collection operations. Furthermore, the current routes were not designed with a focus on minimizing emissions, making it worthwhile to explore the potential benefits of incorporating environmental considerations into the route planning process.

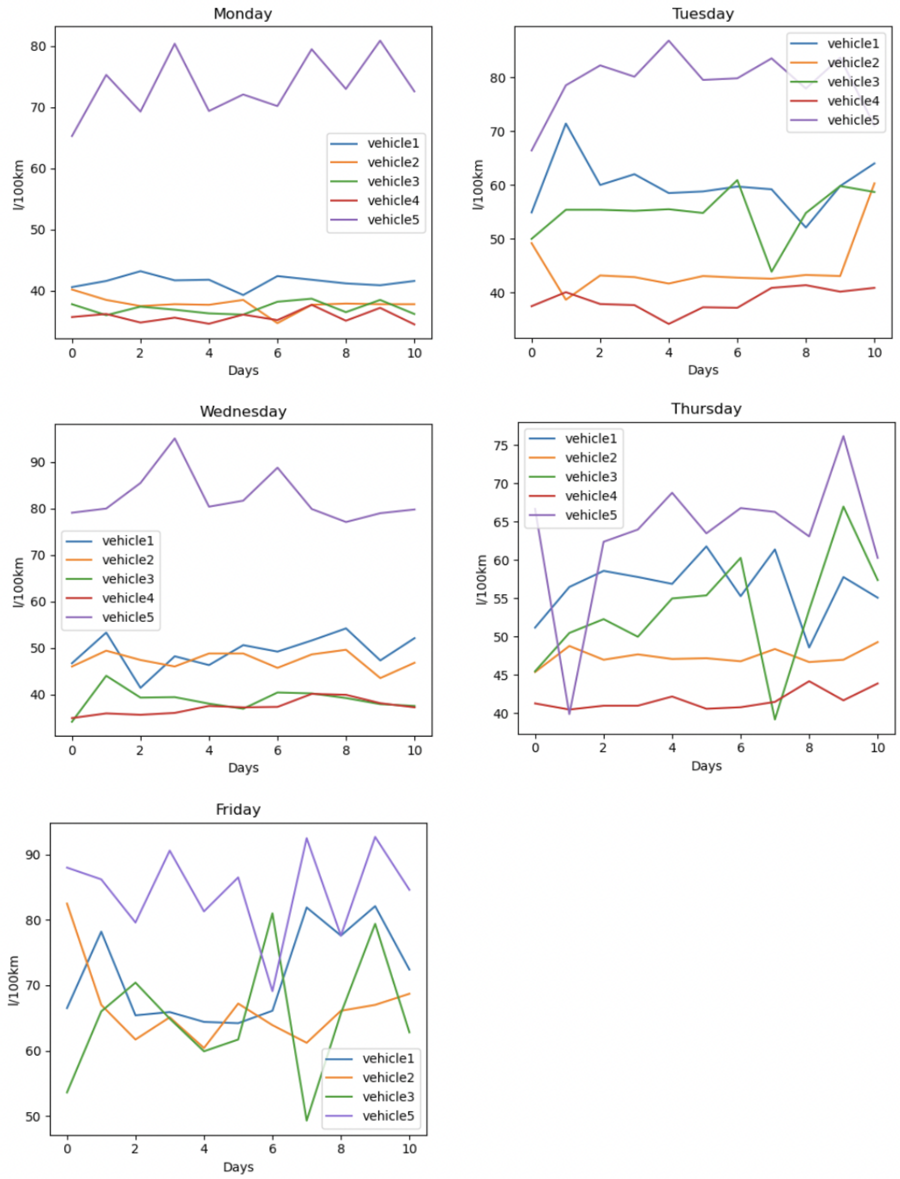

Figure 7.3: This figure illustrates the fuel consumption patterns of Remiks' waste collection operations carried out by five vehicles from Monday to Friday. The chart depicts the differences in fuel consumption across various routes within the current waste collection system. Vehicle 5 exhibits the highest overall fuel consumption, while Vehicle 4 demonstrates the lowest. The vehicle with the highest fuel consumption (purple) predominantly covers routes in the city area, whereas the vehicle with the lowest fuel consumption (red) primarily operates within the district area.

7.3 Speed, Load, and Road Gradient

When minimizing fuel consumption in waste collection routes, it is important to consider factors such as speed, load, and road gradient. These variables directly influence fuel consumption. For instance, the load carried by vehicles, which increases with the amount of waste collected, significantly impacts energy expenditure and fuel consumption. Furthermore, maintaining a consistent speed reduces the need for frequent acceleration and deceleration, optimizing fuel consumption. Additionally, road gradient plays a critical role in fuel consumption. Uphill slopes require more power and fuel, while downhill slopes offer opportunities for fuel savings. By analyzing and optimizing these variables, fuel savings can be achieved, resulting in cost reduction and reduced emissions from the waste collection operations.

Speed

To gather accurate speed data during waste collection operations, there are certain challenges that need to be addressed. The frequent stops and starts during the collection process make it difficult to obtain reliable speed data. However, by manually tracking one of the routes, it was possible to gather valuable insights into the average and maximum speeds during waste collection. Figure 7.4 demonstrates that the average speed during a route was 13.8 km/h, with the maximum speed recorded on this route being 54 km/h. While the traffic regulations may have some influence on the speed of the waste trucks, the frequent stops and starts significantly impact their overall speed.

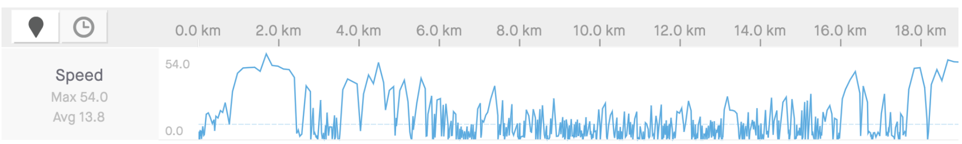

Figure 7.4: This figure illustrates the speed variations of waste collection operations along a selected route. The average recorded speed is 13.8 km/h, with the highest speed recorded at 54 km/h. The figure shows the impact of frequent stops and starts in waste collection, resulting in noticeable speed fluctuations.

Load

The side loaders have a curb weight of 13.5 tons, which is the weight of the truck without any waste collected or additional load. However, the weight measurement system used in each vehicle to monitor waste bin collection lacks the precision required to obtain accurate data for individual households. The weight information, expressed in tons, is displayed on the monitoring screen, with changes reflected for every 100 kilograms collected. However, the fully-loaded vehicle can reach a maximum weight of 19 tons, with 5 tons specifically allocated for waste collection.

Road Gradient
One of the notable characteristics of the city area in Tromsø is its significant variation in height or elevation. Steeper road gradients in this region result in increased fuel consumption for vehicles. For instance, the waste depot at the Remiks main office is situated at an altitude of 12 meters above sea level, while specific households are located as high as 100 meters above sea level. Figure 7.5 illustrates the altitude through an arbitrary waste collection route in the city area.

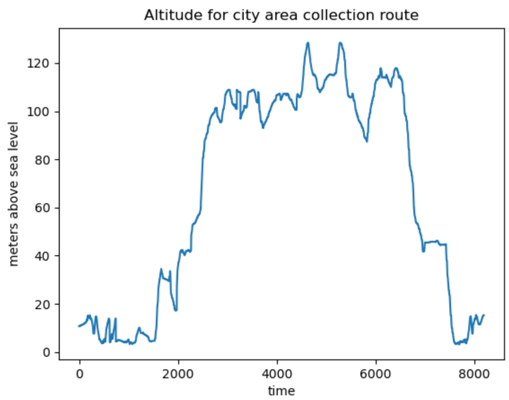

Figure 7.5: This figure illustrates the altitude variation observed along one of the existing waste collection routes. It shows the elevation changes encountered throughout the route, allowing for a better understanding of the topographical characteristics and potential challenges faced during waste collection operations.

An average road gradient matrix can be generated and utilized as a distance matrix to establish a connection between road gradient and route optimization for a specific problem. Equation 7.1 determines the road gradient by dividing the elevation difference between nodes by their distance, then multiplying the result by 100 [62]. This calculation gives the average road gradient percentage for each edge. Elevation data is obtained from the Google Maps API to acquire the necessary elevation information.

$$\text{Road Gradient} = \frac{\text{difference in elevation (masl)}}{\text{distance between nodes (m)}} \tag{7.1}$$

7.4 Vehicle Specific Data

Specific vehicle data is necessary to calculate the forces and estimate fuel consumption as described in the methodology. The required data is obtained from Remiks, emphasizing the Volvo FMX 420 2021 model. This model is a side loader commonly used in waste collection operations.

Drag Coefficient, C_d
The drag coefficient is a numerical measure representing how much an object generates drag force. For Volvo Trucks, an approximate drag coefficient of around 0.6 has been reported [105] [106] [107]. This coefficient provides insight into the vehicle's aerodynamic efficiency. This coefficient serves as an indicator of the vehicle's aerodynamic efficiency.

Frontal Area of Vehicle, A_f
A_f represents the frontal area of an object, which corresponds to the projected area perpendicular to the direction of motion. Based on measurements provided by Remiks [108] for the Volvo FMX 420 2021 model side loader, the width is 2.5 meters, the length is 2 meters, and the height is 4.1 meters. Calculating the frontal area as the product of width and height gives a frontal area of $10.25m^2$.

Mechanical Efficiency, η_m
Mechanical efficiency quantifies the effectiveness of power transmission from the engine to the wheels in a powertrain, considering the losses that arise during power transfer and conversion. The mechanical efficiency of heavy-duty trucks is typically estimated to be around 80% [109]. This means that the mechanical components of the truck, such as the drivetrain, gears, and bearings, can convert 80% of the input mechanical energy into helpful output mechanical energy. At the same time, the remainder is lost due to friction and other mechanical losses.

Low Heating Value, LHV
The low heating value (LHV) is a measure of the energy released when a fuel undergoes complete combustion, taking into account the latent heat of vaporization of the water vapor formed during the process. It signifies the maximum amount of heat energy obtainable from a fuel when completely burned, and the resulting water vapor is condensed and cooled. In the case of diesel engines, the approximate LHV ranges from 42 to 46 MJ/kg [110] [111]. A value of 44 MJ/kg will be used for the experiments.

Fuel Density, ρ_f
Remiks' waste collection vehicles are powered by diesel fuel, which has a density of 0.85 kilograms per liter [112].

Engine Thermal Efficiency, η_{th}
Thermal efficiency is a measure of how effectively a system converts thermal or heat energy into useful work. The thermal efficiency of modern heavy-duty trucks, including Volvo trucks, can vary between 40% and 45% [113] [114]. This means that approximately 40-45% of the input energy from the fuel is converted into useful work, while the rest is lost as waste heat.

7.4.1 Other Relevant Data

Density of air, ρ_a
The density of the fluid medium through which the vehicles move is represented by ρ, with air (a) being the specific medium in this case. The approximate air density is 1.23, kg/m^3 [115].

Coefficient of Rolling Resistance, C_r
The coefficient of rolling resistance, denoted as C_r, is a dimensionless constant that quantifies the resistance to rolling experienced by an object on a surface. In the case of a waste truck traveling on the roads in Tromsø, this coefficient can vary depending on several factors, including tire type, truckload, and road surface condition. As a general estimation, for a loaded waste truck on a typical paved road with dry conditions, the coefficient of rolling resistance can range from 0.008 to 0.014 [116] [117] [107]. However, it is essential to note that this range is subject to change if the road surface is wet or the tires are improperly inflated or worn out, as these conditions can significantly increase the coefficient of rolling resistance.

Focus Area for Experiments

A particular focus area within Remiks' data is selected to conduct the experiments. This focus area includes the city area of Tromsø. The area was chosen due to the high fuel consumption observed in this area, as indicated by the findings of the project thesis [19] and this thesis (Appendix). Additionally, this area provides a good representation of the overall situation due to its varying road gradients, which have been identified as contributing factors to increased fuel consumption.

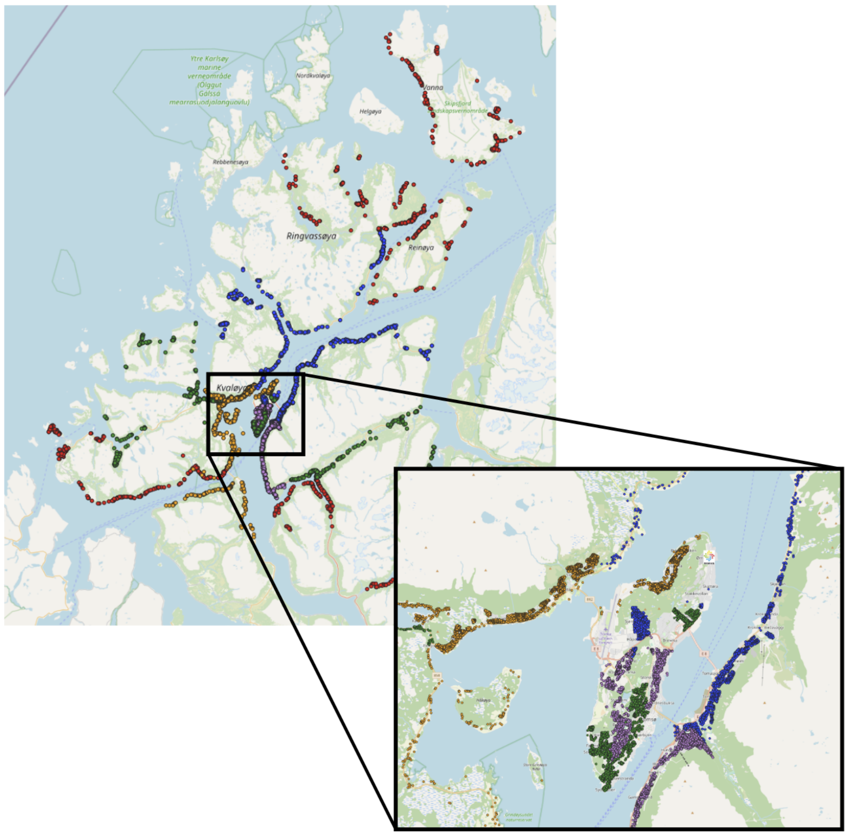

Figure 7.6: This figure shows a map of the focus area on Tromsøya, highlighting households needing waste collection services provided by Remiks.

Figure 7.6 visually represents the selected focus area. Minimizing fuel consumption is the primary objective, and road gradient is a crucial factor considered. The district areas, with a flat road gradient along the ocean, have limited potential for fuel reduction compared to the more varied gradients on Tromsøya. A prior study in Colombia successfully employed a similar approach, highlighting the significance of road gradient in reducing fuel consumption [100]. This serves as a promising foundation for the methodology employed in this thesis.

Part III

Results & Experiments

8
Results and Analysis

This chapter presents the results obtained from implementing the attention-based neural network, as described in the methodology section, to solve the Green Vehicle Routing Problem (GVRP) for waste collection operations in Tromsø for Remiks. The solution routes derived from our model are compared with the current waste collection practices observed in Tromsø. Additionally, an analysis is provided to explain the choices for obtaining the optimized route, including multi-objective optimization, the training process for the attention-based neural network, the choice of the distance matrix, and the scalability of the proposed approach to accommodate more nodes.

8.1 Sustainable Waste Collection Routes

This thesis presents a novel approach for the generation of sustainable waste-collection routes, utilizing an attention-based neural network. The proposed method is implemented and evaluated using real-world scenarios within the waste management operations of the Tromsø company. The obtained results reveal the potential to generate waste collection routes that significantly reduce fuel consumption compared to distance-minimized routes traditionally generated.

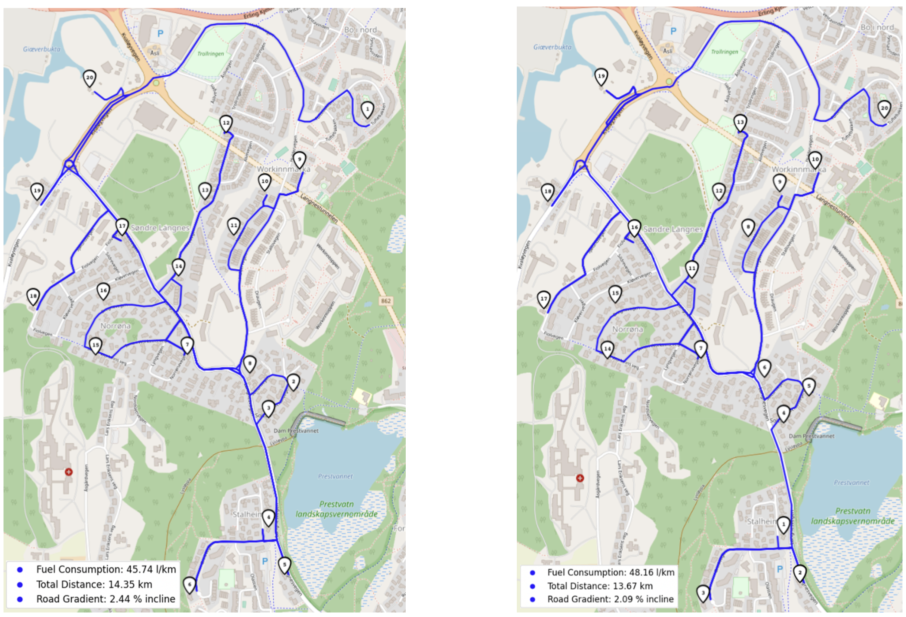

(a) Sustainable Waste Route

(b) Distance Minimized Waste Route

Figure 8.1: This figure illustrates the solution routes generated from the proposed method utilizing an attention-based neural network. In **(a)**, the route is generated with the objective of balancing road gradient and distance, tuned accordingly. **(b)** shows the route generated with the sole objective of distance minimization. The results highlight the capability of the method to generate fuel consumption minimizing routes, contributing to the reduction of emissions in the waste collection process.

Figure 8.1 provides a visual representation of the distance-minimized route and the fuel-minimized route. The fuel-minimized route is obtained using a multi-objective optimization approach, with 70% weighting assigned to road gradient and 30% weighting assigned to distance. This approach seeks to achieve a harmonious balance between minimizing overall distance and optimizing fuel efficiency while considering road gradient as a crucial factor influencing the fuel consumption of waste collection vehicles.

By integrating these two key parameters into the optimization process, the fuel-minimized route presents a more sustainable and environmentally conscious solution compared to the conventional distance-minimized route. The utilization of multi-objective optimization techniques enables the generation of waste collection routes with lower fuel consumption, thereby promoting sustainable waste management practices.

8.2 Analysis of Methodology and Results

This section presents an in-depth analysis of the choices for obtaining the optimized route using the proposed approach. The analysis includes a detailed analysis of the multi-objective optimization employed, the training process for the attention-based neural network, the selection of the distance matrix used, and the scalability of the proposed approach to handle more nodes. This comprehensive analysis will provide a better understanding of the methodology used to derive the optimized waste collection routes.

8.2.1 Distance Metrics

This section compares the results obtained using both Euclidean, road distance, and road gradient in waste collection route optimization. The strengths and weaknesses of each distance metric in the context of the optimization process are analyzed to gain a better understanding of their impact on the efficiency of waste collection operations.

Euclidean Distance
The Euclidean distance matrix is commonly used in various routing problems to measure distances between points in a two-dimensional space. It calculates distances based on straight-line paths using the Euclidean distance formula:

$$d(p, q) = \sqrt{(q_1 - p_1)^2 + (q_2 - p_2)^2}, \tag{8.1}$$

This formula enables easy computation of distances without the need for an API. However, it is important to recognize that relying solely on the Euclidean distance has limitations when applied to real-world scenarios.

Although the Euclidean distance provides a straightforward approach for validating route optimization techniques, it may not capture the complexities of road networks or generate realistic routes. Figure 8.2 illustrates how the resulting solution routes may include infeasible paths crossing oceans and mountains instead of adhering to road networks. Therefore, moving beyond the Euclidean distance is a significant advancement toward applying this method to real-world applications, as it allows for capturing the specifics of the road network and producing more realistic solution routes.

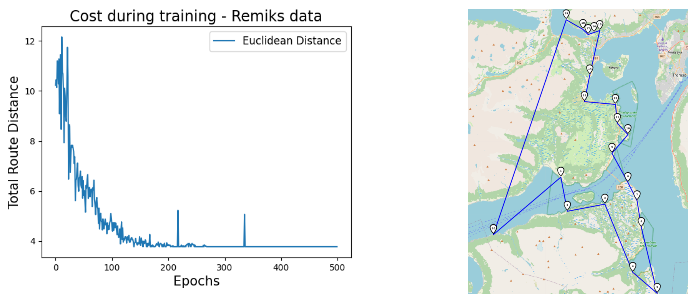

(a) Cost during Training **(b) Solution Route**

Figure 8.2: Plots of the cost during training (a) and the resulting solution route (b) when using the Euclidean distance as the distance matrix in the model. The cost during training shows the optimization progress over 500 epochs, and the solution route illustrates the generated route that includes unrealistic edges crossing oceans and mountains rather than strictly following roads.

Road Distance

In real-world scenarios, optimizing routes require accurate distance calculations. Our approach utilizes actual road distances as the distance matrix, enabling us to capture the intricacies of the road network. By doing so, we offer more precise and realistic solutions that align with real-world constraints, as illustrated in Figure 8.3. This method requires API integration to access up-to-date road network data, ensuring the accuracy and relevance of our results. Incorporating the complexities of road networks, our approach can help decision-makers with accurate and realistic solution routes.

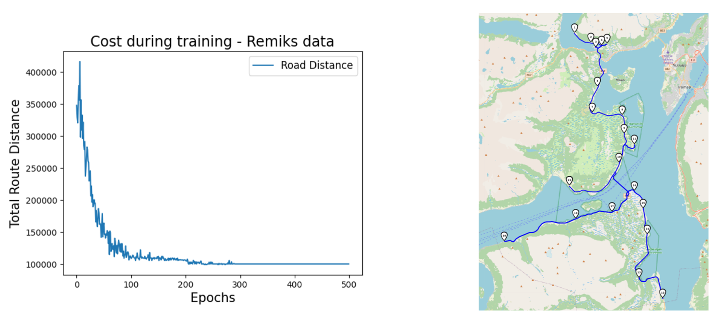

(a) Training Cost **(b) Solution Route**

Figure 8.3: Plots illustrating the training cost (a) and the resulting solution route (b) when using the road distance as the distance matrix in the model. The training cost shows the optimization progress over 500 epochs, while the solution route depicts the generated route that strictly follows feasible paths along the actual road network.

Road Gradient
In order to further enhance the sustainability of the waste collection solution routes, we have integrated the concept of fuel consumption minimization into our approach. Recognizing the significant impact of road gradient on fuel use, particularly in areas with varying elevations like Tromsø, we have incorporated road gradient as a key factor in our distance matrix.

Road gradient refers to the slope or incline of a road segment and plays a crucial role in influencing energy consumption, vehicle performance, and travel time across different routes. By considering road gradient as the distance matrix, we can effectively optimize routing decisions to reduce fuel consumption. By optimizing the road gradient in our solution routes, we can generate solutions routes that give Remiks the chance to actively contribute to emission reduction and align with their goals of promoting sustainable transportation. This optimization allows us to generate solution routes that align with Remiks' goals of promoting sustainable transportation and contributing to emission reduction.

To demonstrate the variations in generated routes resulting from road distance and road gradient, we consider an area on Tromsøya that exhibit noticeable road gradient characteristics, depicted in Figure 8.5.

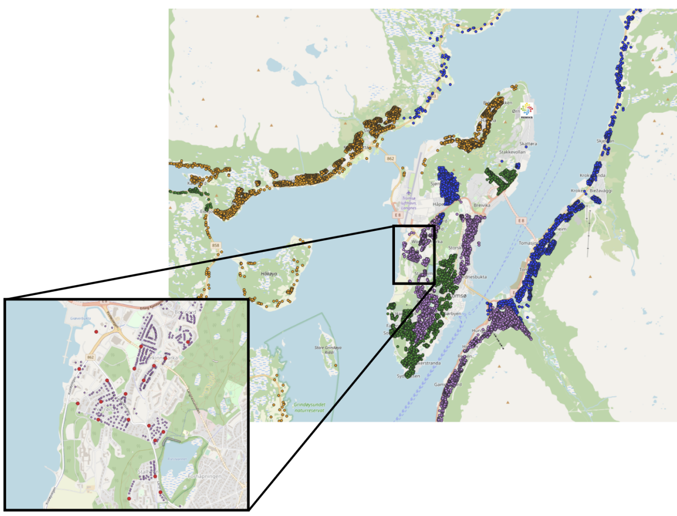

Figure 8.4: This figure depicts a terrain with diverse topography, specifically highlighting significant variations in road gradient. The purpose is to showcase the impact of road gradients on the distance matrix within the area. The terrain exhibits noticeable inclines, contributing to varying levels of elevation throughout the depicted region.

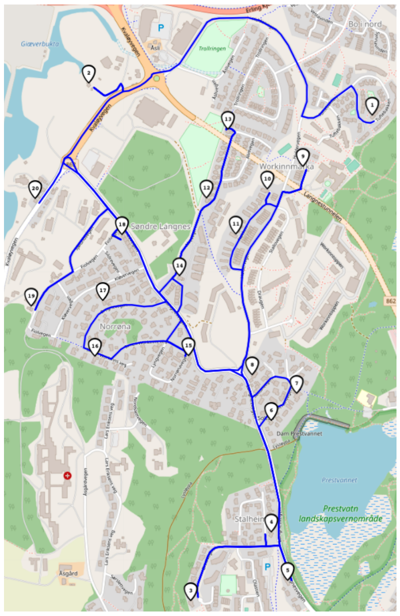

(a) **Solution route for road distance.**

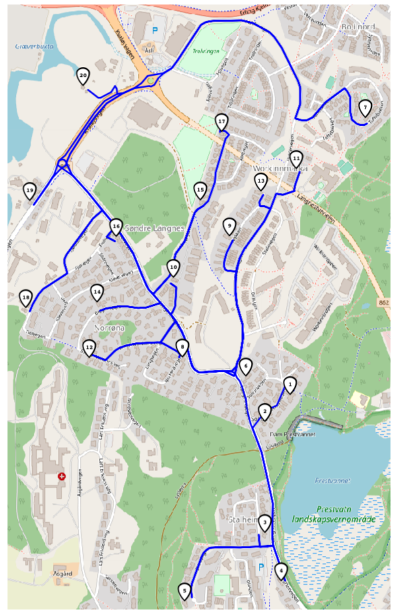

(b) **Solution route for road gradient.**

Figure 8.5: The figure illustrates the solution routes for an area characterized by noticeable road inclines using two different distance metrics: **(a)** road distance and **(b)** road gradient.

The same area was utilized to generate two solution routes: one optimized using road distance as the distance matrix and the other optimized using road gradient as the distance matrix. The resulting routes are illustrated in Figure 8.5. This example offers insights into the learning and optimization process of the model and its alignment with our objectives.

The route derived from using the road distance as the distance matrix, shown on the left, initially chooses the outline node and incorporates node 2 while ascending towards the peak of the hill. It then encounters nodes 3, 4, and 5 before descending downhill.

In contrast, the route generated based on the road gradient, displayed on the right, starts at the summit of the hill and covers nodes 1 to 5 while descending towards node 6. Interestingly, node 7, situated at a distance from the main route, is selected before the model backtracks to the hill for node 8. This behavior is influenced by the road gradient formula, which divides the elevation difference by the road distance. Despite the significant difference in road distance between nodes 7 and 8, the elevation remains relatively unchanged. Consequently, the model chooses the outlying node located between these two locations, which is suboptimal.

This observation highlights the importance of considering additional factors, such as fuel consumption or a balancing objective with distance, to achieve a better balance in route optimization. In this thesis, we explore a multi-objective approach to address this challenge.

8.2.2 Multi-Objective Optimization

In this section, we describe the process of setting up the multi-objective optimization problem to solve the GVRP for the waste collection operations of Remiks in Tromsø. We discuss how we weighed the two objectives of minimizing distance and minimizing fuel consumption against each other to arrive at the proposed solution. Additionally, we analyze the trade-offs between these objectives to understand better the solutions obtained.

As discussed in Chapter 6, we employ the following objective function for the multi-objective optimization:

$$min \quad \alpha \sum_{i,j \in A} E_{ij} x_{ij} + \beta \sum_{i,j \in A} c_{ij} x_{ij}, \tag{8.2}$$

This equation incorporates two scaling parameters, namely α and β, to effectively balance the influence of road gradient and distance within our model. The parameter α controls the weight assigned to the distance, while β determines the weight assigned to the road gradient. Importantly, these parameters are subject to the constraint that their sum must equal 1. To examine the impact of different parameter values, we conducted a series of tests, as illustrated in Figure 8.6.

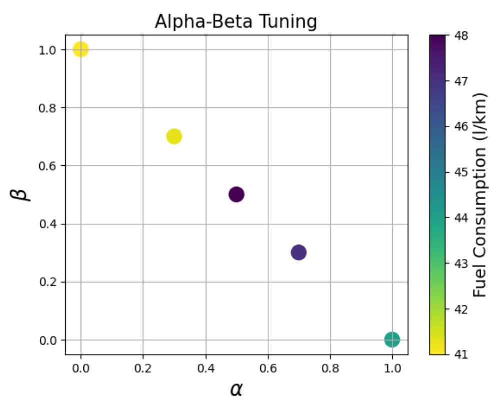

Figure 8.6: This figure illustrates the relationship between the scaling parameters α (distance) and β (road gradient) and their impact on fuel consumption. The figure shows different combinations of α and β values ranging from 0 to 1, representing various weightings of distance and road gradient in the optimization process.

This analysis reveals several key insights. Firstly, when the distance has zero contribution and the optimization is solely based on road gradient, the model achieves the lowest fuel consumption for a given route. On the other hand, a balanced weighting of 50-50 between distance and road gradient leads to suboptimal results, where neither factor dominates, resulting in longer routes and limited environmental benefits.

Optimizing solely based on distance yields lower fuel consumption. However, by incorporating a balanced contribution of both distance and road gradient, we can achieve a favorable compromise. Notably, when α is weighted at 0.7 and β at 0.3, we get a good balance, resulting in a route with reduced fuel consumption.

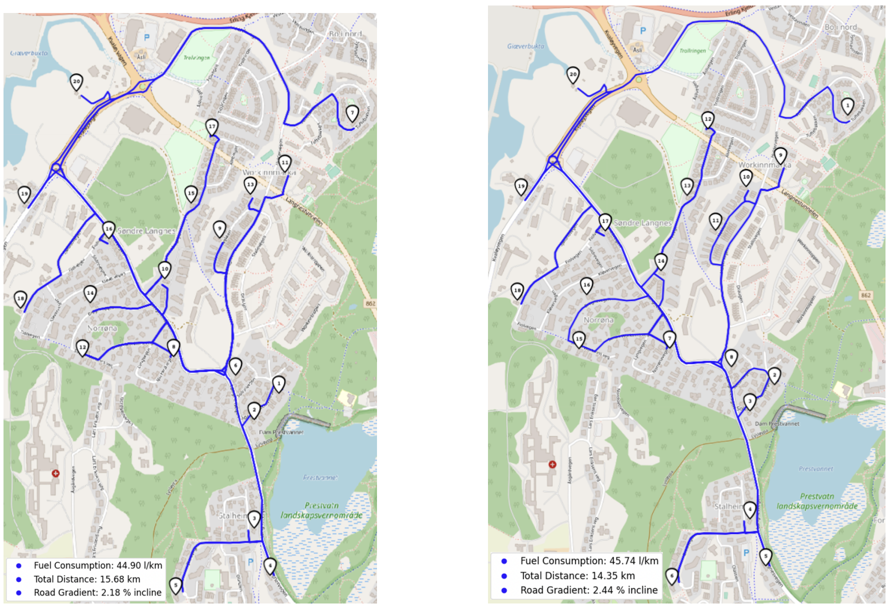

(a) Solution route: $\boldsymbol{\alpha} = 0$, $\boldsymbol{\beta} = 1$.

(b) Solution route: $\boldsymbol{\alpha} = 0.3$, $\boldsymbol{\beta} = 0.7$.

Figure 8.7: The figure showcases two solution routes with varying α-β tuning. The plots demonstrate that by effectively balancing the impact of distance and gradient, we can generate a route that exhibits enhanced visual attractiveness and simplicity for the driver. Notably, the balanced route excludes the outlying node 7, while the route generated solely based on gradient minimization includes it.

Figure 8.7 presents a visual representation of the solution routes, comparing those generated with and without the consideration of balanced road gradient and distance. The results demonstrate that the tuning process contributes to the generation of more realistic routes, incorporating aspects of simplicity. The inclusion of simplicity is crucial for ensuring the practicality and realism of the routes. Notably, this approach involves a slight compromise in fuel consumption, but it enables the generation of more holistic and sustainable routes.

8.2.3 Training the Attention-based Neural Network

This section delves into the specifics of the training process for the neural network employed in this study. It will cover the selection of hyperparameters and any adjustments made during training. By providing a detailed overview of the training process, this section aims to give insight into the factors that contribute to the proposed approach.

Learning Rate

In order to determine an appropriate learning rate, several values were tested and evaluated. Through a process of trial and error, it was observed that a learning rate of 10^{-4} gave the best performance in terms of training convergence speed and stability. This choice is supported by the results obtained, as illustrated in Figure 8.8. The plot showcases the training progress with different learning rates, demonstrating that a learning rate of 10^{-4} achieves the desired balance between rapid convergence and stable training.

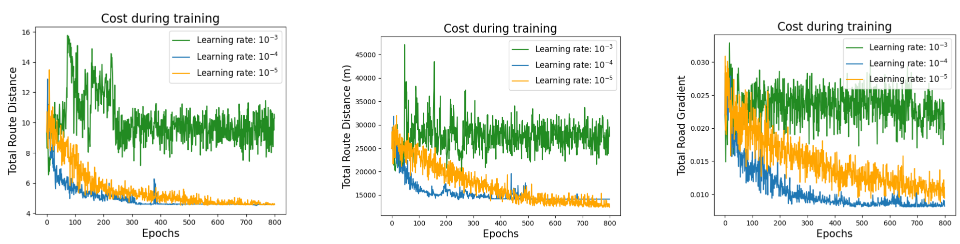

(a) Euclidean **(b)** Road distance **(c)** Road gradient

Figure 8.8: This figure presents the training progress with three different distance matrices: Euclidean distance **(a)**, road distance **(b)**, and road gradient distance **(c)**. It can be observed that the learning rate of 10^{-4} consistently outperforms the other learning rates across all three distance matrices.

Baseline

In this thesis, we conducted experiments using REINFORCE with two different baselines: exponential moving average and greedy rollout. Our results consistently showed that the exponential moving average baseline outperformed the greedy rollout baseline, as illustrated in Figure 8.9. These findings indicate that the exponential moving average baseline offers more reliable and stable estimates of expected rewards, thereby enhancing route optimization for the specific application under investigation.

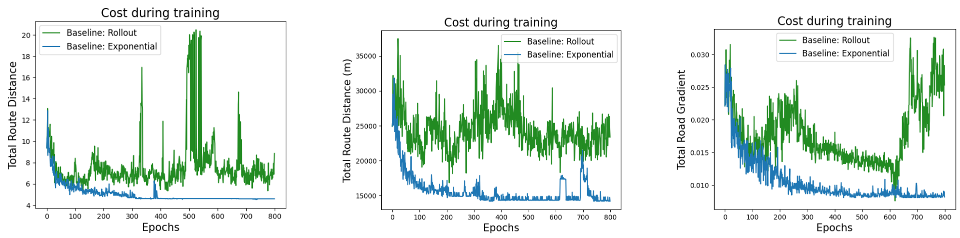

(a) Euclidean **(b)** Road distance **(c)** Road gradient

Figure 8.9: This figure illustrates the training progress using three different distance matrices: Euclidean distance **(a)**, road distance **(b)**, and road gradient distance **(c)**. It can be observed that the exponential baseline performs better than the rollout baseline in our application.

Exponential Baseline Decay Rate β

After identifying the exponential moving average baseline as the most effective for this application, we fine-tuned its performance by testing different β values. The parameter β represents the decay rate in the exponential moving average baseline. It determines the weight assigned to previous baseline values when updating the baseline in each iteration. A higher β value emphasizes the influence of historical baseline values, while a lower β value gives more weight to the current iteration's performance.

In the testing section, we conducted evaluations using different β values to assess their impact on the performance of the exponential moving average baseline. The results, illustrated in Figure 8.10, reveal how varying β influences the optimization outcomes. Notably, a β value of 0.3 yielded the most favorable results for Euclidean distance, whereas higher β values proved more effective for road distance and road gradient. It was observed that Euclidean distance remained relatively stable across all β values, while road distance and road gradient exhibited greater sensitivity to changes in the value.

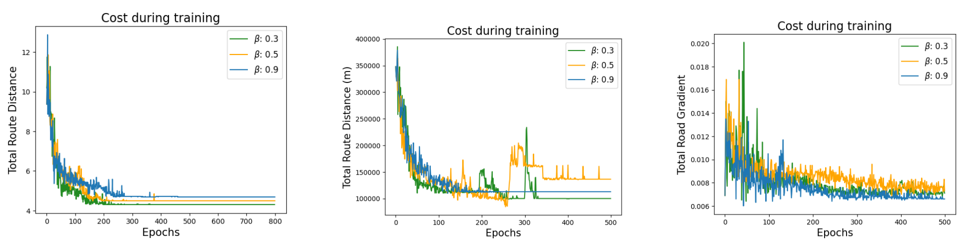

(a) Euclidean **(b)** Road distance **(c)** Road gradient

Figure 8.10: This figure demonstrates the impact of varying β values on the optimization outcomes for different distance metrics. Each subfigure represents a specific distance metric: (a) Euclidean distance, (b) road distance, and (c) road gradient.

8.2.4 Scaling

In this section, we explore the scalability of our method by increasing the number of nodes in the route optimization process. Specifically, we focus on one selected area with the highest fuel consumption (Appendix). This area is illustrated in Figure 8.11. By scaling up the chosen route from 20 nodes to 200 nodes, we aim to investigate the potential of this particular area and gain insights into the performance and effectiveness of our approach.

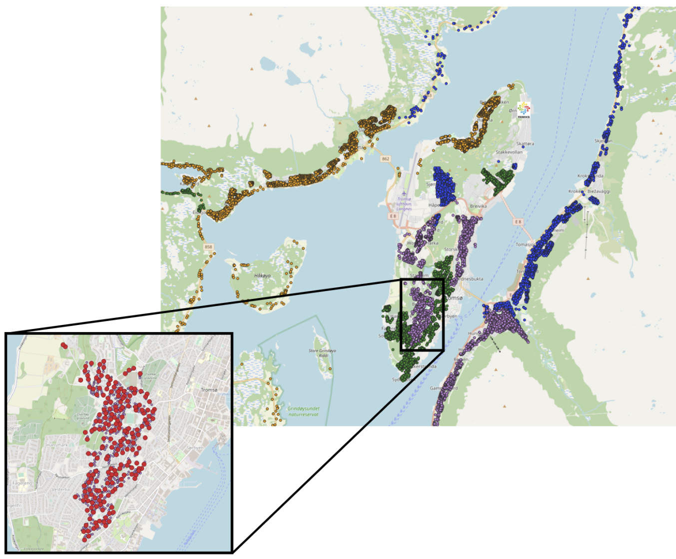

Figure 8.11: This figure illustrates the selected area for the scaling analysis. This particular area was found to have the highest fuel consumption in current waste collection operations. In the small image, the purple area represents the actual nodes present in the area. In contrast, the red nodes depict the 200 selected nodes that are strategically distributed to capture the node distribution in this area.

The scaling experiments shed light on the feasibility and performance of our approach when applied to larger route optimization problems. By investigating the potential of the area with the highest fuel consumption, we gain valuable insights into the adaptability and scalability of our method. These findings contribute to the broader understanding of our approach's applicability in real-world scenarios.

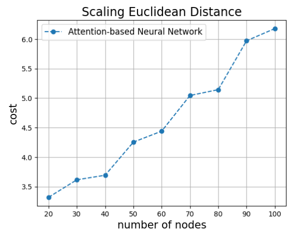

(a) Euclidean

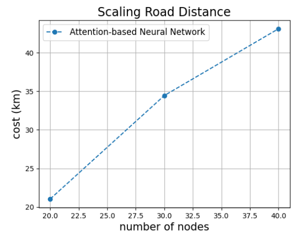

(b) Road distance

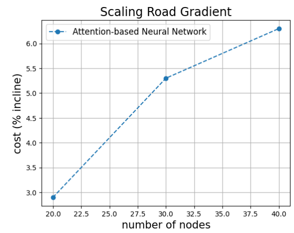

(c) Road gradient

Figure 8.12: This figure shows the results of the scaling experiments conducted. The experiments involved scaling up to 100 nodes for Euclidean distance and limited scaling to 20, 30, and 40 nodes for road distance and road gradient due to API restrictions. Further investigation into scaling for larger datasets is recommended for future research.

In Figure 8.12, we present the results of the scaling experiments. The scaling was conducted up to 100 nodes for the Euclidean distance metric, while for the road gradient and road distance, it was limited to 20, 30, and 40 nodes. The restricted API we used for retrieving road gradients and real road distances imposed these limitations on scaling. Therefore, we were unable to scale the road gradient and real road distance experiments beyond these node counts. It is worth emphasizing the importance of further exploration of scaling in future research.

The scaling results indicate that the method demonstrates promising performance even with larger problem sizes, as the scaling is not exponentially increasing and provides routes that appear visually appealing. However, to comprehensively assess the scalability of the proposed method, it should be compared to other approaches and subjected to more detailed testing.

Moving forward, it is crucial to consider scaling up the methodology for larger datasets and incorporating more complex constraints. This analysis will enable a thorough evaluation of the impact of scaling on the methodology's performance and the efficiency of the optimization process.

9

Concluding Remarks

The conducted experiments examined the effectiveness of the proposed method in addressing a real-world waste collection problem, and the results have highlighted its potential for sustainable route optimization. In the subsequent section, we delve into a discussion of the proposed method, exploring its advantages and limitations. Furthermore, this chapter presents the findings regarding the application of reinforcement learning and highlights the limitations concerning fuel consumption. Additionally, we explore the value that this method can offer to Remiks. Furthermore, we outline potential avenues for future research and development. Finally, a conclusion summarizes the key findings and contributions of this study.

9.1 Discussion

This thesis explores the application of an attention-based neural network methodology for optimizing waste collection routes in a real-world scenario. The results demonstrate the model's effectiveness in training and problem-solving for this purpose, indicating promising results.

The utilization of the attention model in waste collection operations for Remiks has shown its capability to minimize fuel consumption. One of the values this can give Remiks is consistency and a more targeted approach compared to their current operations, where they can use the method to work to achieve their goals of emission reduction. By utilizing the multi-objective optimization framework, where the parameters α and β control the weighting of distance and road gradient respectively, we can fine-tune these parameters to strike a balance between environmental considerations and distance minimization. This balance contributes to the realism of the generated routes, simplifying the task for the drivers. Studies have consistently highlighted the importance of simplicity in real-world systems, as it plays a vital role in gaining driver acceptance and ensuring the successful implementation of the routes [57] [118]. For example, nodes outlying from the main route would be counter-intuitive for the drivers and undermine the effectiveness of the generated routes. By taking into account a broader objective than only fuel consumption, this balance was achieved.

While the primary goal is to minimize environmental impact, it is essential to consider the associated costs. Taking the extra computational effort to strike a balance between environmental considerations, specifically fuel consumption in this case, can lead to significant benefits. The generated routes for Remiks are designed to be long-lasting, considering the pace of urban development and other changes. This ensures that the routes remain relevant for several years, making the investment of computational resources worthwhile. Thus, achieving sustainable waste collection routes is possible through a balanced approach that aligns with the principles of the Triple Bottom Line (TBL). TBL emphasizes a holistic, long-term perspective by considering economic, environmental, and social factors in decision-making processes.

Reinforcement Learning

Through the work of this thesis, reinforcement learning have shown promise in solving routing problems. While it may not yet be competitive with the established methods, it has demonstrated its capability to solve routing problems effectively.

One notable advantage of reinforcement learning is its ability to solve problems in real-time optimization without the need for a pre-existing training dataset. With a supervised approach pre-existing dataset is required to train a model. This dataset may consist of historical data or manually crafted rules based on domain expertise. However, creating and maintaining such datasets can be time-consuming, expensive, and impractical, especially in dynamic environments where conditions change frequently.

Reinforcement learning provides a data-driven approach to problem-solving, distinguishing it from methods that rely on pre-existing datasets. In the case of waste collection generation for Remiks, the utilization of reinforcement learning showed promising applicability, as it allows the learning agent to interact directly with the environment and learn from the feedback it receives. This characteristic of reinforcement learning holds great promise for real-world applications, as it enables the model to adapt and optimize waste collection routes based on real-time information and changing conditions.

However, it is important to acknowledge that reinforcement learning still faces challenges, particularly in achieving optimality. Routing problems, by their nature, are known to be NP-hard, making the search for an optimal solution computationally infeasible. Nevertheless, by continually pushing the boundaries and making incremental improvements in solution quality and computational time, reinforcement learning can make significant strides in addressing these challenges.

Fuel Consumption

While the attention-based method demonstrates promising results, it is important to acknowledge certain limitations and areas for improvement. One potential avenue for refinement lies in enhancing the estimation of fuel consumption. In this thesis, the focus was primarily on road gradients as a factor influencing fuel consumption. However, exploring the influence of load and speed in greater detail would be highly valuable. By incorporating these additional variables, a more comprehensive understanding of fuel consumption patterns and their impact on route optimization can be achieved. This would also enable a direct comparison between the actual fuel estimates obtained from Remiks' data and the generated estimates from the method used in this thesis.

A disparity was observed between the fuel estimates in this study and the actual fuel consumption gathered from Remiks. This disparity arises from the use of a simplified model that does not consider cumulative load, variating speed, or the effects of starts and stops. Despite these simplifications, the method still provides valuable insights, and the results generated are internally consistent and comparable.

Future research should consider incorporating the cumulative effect of load, accounting for varying speeds, and addressing the impact of starts and stops to improve the accuracy and comparability of the fuel consumption estimates. When considering speed, it is worth noting that many existing studies on route optimization assume an average speed of over 40 km/h [101] [63]. This assumption does not hold true for waste collection applications, as the average speed in this context was found to estimate to be 14 km/h.

Therefore, further exploration of speed is necessary to develop a more accurate fuel consumption model. Although this area of research is relatively underdeveloped for transportation scenarios involving lower speeds, it presents an interesting avenue that could yield valuable insights and results. By refining the model to better align with real-world scenarios, the generated estimates can be directly compared with actual fuel consumption, enhancing the overall applicability and validity of the research findings.

Value for Remiks

Training Purpose
Another benefit of the proposed method for waste collection routes proves to be highly advantageous for training purposes. The conventional training process often requires a significant amount of human time and has the potential for improvement in several aspects. By utilizing this method, drivers can immediately benefit from optimized routes, starting from day one, instead of having to manually optimize the routes themselves over many years of driving the waste collection route. This streamlined approach saves time and ensures optimal efficiency from the outset. It removes the dependency on senior drivers' knowledge and eliminates the requirement for new drivers to spend years driving routes before achieving an optimized waste collection route.

Goal-Oriented Strategies
The implementation of the proposed methodology for waste collection routes offers an additional advantage by enabling Remiks to enhance their focus on achieving their goals, specifically in terms of emission reduction. This approach mitigates the inherent randomness associated with route selection, as it eliminates the sole reliance on individual driver preferences. By employing this route optimization approach, Remiks can proactively implement a more targeted and goal-oriented operational strategy. Consequently, this positions Remiks as a competitive and sustainable player in the waste management industry, providing them with a competitive advantage.

Flexibility
Furthermore, the model can be extended to incorporate additional constraints or objectives based on specific requirements. For instance, the same model can be used to optimize vehicle range for electric vehicles by considering factors affecting fuel consumption and range. By including constraints such as battery capacity, charging station availability, and energy efficiency, the model can be adapted to find routes that maximize vehicle range while minimizing energy consumption. This flexibility allows for the application of the model to various optimization scenarios, making it a versatile tool for addressing different objectives for the company.

Sustainable Waste-Collection Routes
In real-world waste collection, the optimization of routes involves challenges that extend beyond efficiency and cost. Factors such as route simplicity and driver satisfaction also play a role but are challenging to objectively quantify [119]. It is crucial to acknowledge that what may appear optimal in theory does not always translate into practical effectiveness. Striking a balance between various factors, including driver satisfaction, environmental considerations, and distance, becomes essential for delivering successful solutions.

The proposed method aims to achieve this balance by considering both fuel consumption and distance through the $\alpha - \beta$ tuning process, leading to the generation of more sustainable and long-term routes. Taking a holistic approach that considers theoretical considerations alongside practical constraints is vital to ensure the viability and effectiveness of the proposed methods in real-world applications. This approach aligns with the principles of the Triple Bottom Line (TBL) framework, emphasizing the need to consider social, economic, and environmental aspects in addition to minimizing fuel consumption alone.

By adopting this holistic perspective, sustainable waste-collection routes can be achieved, taking into account the diverse range of factors that contribute to sustainability in waste management operations.

9.2 Future Directions

This section delves into several potential avenues for future research and development to enhance the effectiveness of the proposed model. In addition to improving fuel consumption estimations, as discussed earlier, there are other intriguing directions worth exploring. The following areas are highlighted as promising avenues: integrating real-time transportation information, incorporating a heterogeneous vehicle fleet, and considering waste bin clustering as an alternative approach for emission reduction.

Real-Time Transportation Information
Integrating real-time transportation information into the optimization model is a promising direction for future development. By leveraging dynamic real-world data, the model can effectively consider traffic conditions, construction work, and road closures, leading to more precise and adaptive route planning [120]. Furthermore, real-time information can offer valuable feedback to drivers, enhancing the training process and overall performance. It is worth noting that incorporating real-time data may impose computational-time challenges due to increased requirements. This research direction holds immense potential and can substantially benefit businesses leveraging the model.

Heterogeneous Vehicle Fleet
Another potential direction for further development of this model is to extend it to include a heterogeneous fleet of vehicles [103], encompassing both side loaders and back loaders and bins with underground tunnels used by Remiks for waste transportation. By incorporating the entire fleet, the optimization model can allocate vehicles across different types, leading to more efficient operations. This extension holds the potential to provide valuable insights and contributing to the ongoing research efforts in this domain.

Furthermore, it is worth noting that optimization naturally thrives in scenarios characterized by limited options and a constrained timeframe. Therefore, the inclusion of a wider range of choices and variations in the optimization models can lead to significant cost savings.

Clustering Waste Bins
While working on optimizing waste collection routes, an alternative approach emerged: exploring the best clustering of waste bins to optimize the overall objective. This method offers an additional avenue to reduce emissions in the waste collection by grouping waste bins together strategically. The process involves consolidating multiple bins into a single collection point, minimizing the number of starts and stops during collection, and consequently reducing fuel consumption. This approach not only saves time and effort in the collection process but also allows the optimization model to consider fewer nodes, resulting in faster computations.

Clustering can be optimized by identifying areas with high fuel consumption, detecting edges with significant energy usage, or employing waste sensors in the bins to determine which bins actually require collection [121]. Additionally, clustering can be achieved by identifying households that can be grouped together based on their waste generation patterns. By incorporating this strategy alongside the proposed model, waste management efficiency and emission reduction efforts can be further enhanced.

When discussing this topic with drivers, a key obstacle emerged: the social aspect. Clustering waste bins presents implications for the municipal tax system, as residents may resist sharing waste management costs with neighbors. Logistical and social dynamics both play a role in grouping bins together. A challenge lies in determining responsibility for bin grouping. Elderly individuals or those in hilly areas may lack the physical capability for this task. Implementing the clustering approach should consider these factors, potentially requiring additional support and resources.

Remiks is actively exploring future possibilities to address challenges, including the potential implementation of sensors in waste bins. This innovative idea involves utilizing real-time fill-level data to optimize the collection process based on actual waste volume. By integrating sensor data into the model, Remiks aims to enhance its effectiveness and leverage the potential of IoT technologies. Therefore, dynamic route planning based on sensor data could minimize unnecessary trips, optimize resource allocation, and presents an exciting avenue for the future development of the proposed method in this thesis.

9.3 Conclusion

In this thesis, an attention-based neural network is proposed for solving the Green Vehicle Routing Problem (GVRP) and optimizing waste collection routes. The primary objective of the proposed GVRP is to minimize fuel consumption, thereby reducing emissions in waste collection operations. This research aims to address the existing research gap in the field of GVRP and its real-world applications.

To validate the effectiveness of the proposed method, it is applied to a practical scenario involving the waste collection company Remiks in Tromsø. The results obtained from applying the proposed method demonstrate its success in generating waste collection routes that minimize fuel consumption and reduce emissions during the waste collection process. While the current approach primarily focused on accounting for road gradient as a factor influencing fuel consumption, further enhancements can be made by considering additional factors, such as variating speed and load. Nevertheless, the findings of this study highlight the potential of the proposed method to contribute to sustainable waste management practices and provide insights for real-world applications.

By combining the power of an attention-based neural network with the specific context of the GVRP in waste management, this research contributes to knowledge and practical solutions in route optimization and emission reduction in the waste collection domain. These findings pave the way for further exploration in this area.

Bibliography

[1] IEA. *Transport Improving the sustainability of passenger and freight transport*. 2023. URL: https://www.iea.org/topics/transport.

[2] European Commission. URL: https://climate.ec.europa.eu/eu-action/transport-emissions/road-transport-reducing-co2-emissions-vehicles/reducing-co2-emissions-heavy-duty-vehicles_en.

[3] Nada R. Sanders and John D. Wood. *Foundations of Sustainable Business: Theory, Function, and Strategy*. 2020. ISBN: 978-1-119-68855-6.

[4] Frederick S. Hillier and Gerald J. Lieberman. *Introduction to Operations Research*. Eleventh. New York, NY, USA: McGraw-Hill, 2021.

[5] Edmund K. Burke et al. "A multi-objective approach for robust airline scheduling." In: *Computers Operations Research* 37.5 (2010). Disruption Management, pp. 822–832. ISSN: 0305-0548. DOI: https://doi.org/10.1016/j.cor.2009.03.026. URL: https://www.sciencedirect.com/science/article/pii/S0305054809000896.

[6] Çağatay Iris and Jasmine Siu Lee Lam. "A review of energy efficiency in ports: Operational strategies, technologies and energy management systems." In: *Renewable and Sustainable Energy Reviews* 112 (2019), pp. 170–182. ISSN: 1364-0321. DOI: https://doi.org/10.1016/j.rser.2019.04.069. URL: https://www.sciencedirect.com/science/article/pii/S1364032119302904.

[7] Paolo Toth and Daniele Vigo. *Vehicle Routing: Problems, Methods, and Applications*. MOS-SIAM Series on Optimization Ser. 2014. ISBN: 978-1-61197-358-7.

[8] Ethem Alpaydin. *Introduction to machine learning*. third. Cambridge, Massachusetts London, England: MIT press, 2014.

[9] Wouter Kool, Herke van Hoof, and Max Welling. *Attention, Learn to Solve Routing Problems!* 2019. eprint: 1803.08475 (stat.ML).

[10] Wouter Kool et al. "Deep Policy Dynamic Programming for Vehicle Routing Problems." In: *CoRR* abs/2102.11756 (2021). arXiv: 2102.11756. URL: https://arxiv.org/abs/2102.11756.

[11] Sahil Manchanda et al. "Learning Heuristics over Large Graphs via Deep Reinforcement Learning." In: *CoRR* abs/1903.03332 (2019). arXiv: 1903.03332. URL: http://arxiv.org/abs/1903.03332.

[12] Chaitanya K. Joshi, Thomas Laurent, and Xavier Bresson. "An Efficient Graph Convolutional Network Technique for the Travelling Salesman Problem." In: *CoRR* abs/1906.01227 (2019). arXiv: 1906.01227. URL: http://arxiv.org/abs/1906.01227.

[13] J. J. Hopfield and D. W. Tank. ""Neural" computation of decisions in optimization problems." In: *Biological Cybernetics* 52.3 (1985), 141–152. DOI: 10.1007/bf00339943. URL: http://dx.doi.org/10.1007/bf00339943.

[14] Michel Deudon et al. "Learning Heuristics for the TSP by Policy Gradient." In: *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*. Ed. by Willem-Jan van Hoeve. Vol. 10848. Lecture Notes in Computer Science. Springer. Springer International Publishing, 2018, pp. 170–181. ISBN: 978-3-319-93030-5. DOI: 10.1007/978-3-319-93031-2_12.

[15] Wenyong Dong and MengChu Zhou. "A Supervised Learning and Control Method to Improve Particle Swarm Optimization Algorithms." In: *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 47.7 (2017), pp. 1135–1144. DOI: 10.1109/TSMC.2016.2557319.

[16] Elena Mocanu et al. "On-Line Building Energy Optimization Using Deep Reinforcement Learning." In: *IEEE Transactions on Smart Grid* 10.4 (2019), pp. 4148–4157.

[17] Mohammadreza Nazari et al. "Reinforcement Learning for Solving the Vehicle Routing Problem." In: *arXiv preprint arXiv:1802.04240* (2018). arXiv: 1802.04240 [cs.AI].

[18] Cecilio Mar Molinero. "Operational research: From war to community." In: *Socio-Economic Planning Sciences* 26.3 (1992), pp. 203–212. ISSN: 0038-0121. DOI: https://doi.org/10.1016/0038-0121(92)90011-S. URL: https://www.sciencedirect.com/science/article/pii/003801219290011S.

[19] Marit Utheim Breimo. *The Green Vehicle Routing Problem and the potential for Machine Learning in the field of Operational Research*. EOM-3010 Project Paper in Energy, Climate and Environment. 2022.

[20] Miltiadis Chalikias. *Applications of Operational Research and Mathematical Models in Management*. Basel, Switzerland: Mdpi AG, 2020. ISBN: 9783039433803, 9783039433810. DOI: 10.3390/books978-3-03943-381-0. URL: https://www.mdpi.com/books/pdfdownload/book/3080.

[21] Maurice Landry, Jean-Louis Malouin, and Muhittin Oral. "Model validation in operations research." In: *European Journal of Operational Research* 14.3 (1983), pp. 207–220. ISSN: 0377-2217. DOI: https://doi.org/10.1016/0377-2217(83)90257-6. URL: https://www.sciencedirect.com/science/article/pii/0377221783902576.

[22] Xin-She Yang. *Introduction to Mathematical Optimization: From Linear Programming to Metaheuristics*. 2008. ISBN: 978-1-907343-66-7.

[23] *Linear Programming - University of Kentucky*. URL: https://www.uky.edu/~dsianita/300/online/LP.pdf.

[24] Andrea Saltelli et al. *Global Sensitivity Analysis: The Primer*. Wiley-Interscience, 2008. ISBN: 978-0-470-05997-5. DOI: 10.1604/9780470059975.

[25] J. Adrian Bondy and U. S. Murty. *Graph Theory with Applications*. 1976. ISBN: 978-0-444-19451-0. DOI: 10.1604/9780444194510.

[26] H. Paul Williams. *Model Building in Mathematical Programming*. Fifth. Wiley-Blackwell, 2013. ISBN: 978-1-118-44333-0.

[27] M. R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. 1979. ISBN: 978-0-7167-1045-5. DOI: 10.1604/9780716710455.

[28] Tolga Bektas. “The multiple traveling salesman problem: an overview of formulations and solution procedures.” In: *Omega* 34.3 (2006), pp. 209–219. ISSN: 0305-0483. DOI: https://doi.org/10.1016/j.omega.2004.10.004. URL: https://www.sciencedirect.com/science/article/pii/S0305048304001550.

[29] Gilbert Laporte. “What you should know about the vehicle routing problem.” In: *Naval Research Logistics* 54.8 (2007), 811–819. DOI: 10.1002/nav.20261. URL: http://dx.doi.org/10.1002/nav.20261.

[30] G. B. Dantzig and J. H. Ramser. “The Truck Dispatching Problem.” In: *Management Science* 6 (1959), pp. 80–91.

[31] Nicos Christofides. “The vehicle routing problem.” In: *RAIRO - Operations Research - Recherche Opérationnelle* 10.V1 (1976), pp. 55–70. URL: http://eudml.org/doc/104635.

[32] Kris Braekers, Katrien Ramaekers, and Inneke Van Nieuwenhuyse. “The vehicle routing problem: State of the art classification and review.” In: *Computers Industrial Engineering* 99 (2016), 300–313. DOI: 10.1016/j.cie.2015.12.007. URL: http://dx.doi.org/10.1016/j.cie.2015.12.007.

[33] Gilbert Laporte. “Fifty Years of Vehicle Routing.” In: *Transportation Science* (2009). DOI: 10.1287/trsc.1090.0301.

[34] M.A. Hannan et al. “Capacitated vehicle-routing problem model for scheduled solid waste collection and route optimization using PSO algorithm.” In: *Waste Management* 71 (2018), pp. 31–41. ISSN: 0956-053X. DOI: https://doi.org/10.1016/j.wasman.2017.10.019. URL: https://www.sciencedirect.com/science/article/pii/S0956053X17307675.

[35] Andrew B. Kahng. T. C. Hu. *Linear and Integer Programming Made Easy*. Springer International Publishing, 2016. ISBN: 9783319240015. DOI: 10.1007/978-3-319-24001-5.

[36] C. Archetti, N. Bianchessi, and M. G. Speranza. “Branch-and-cut algorithms for the split delivery vehicle routing problem.” In: *European Journal of Operational Research* 238 (2014), pp. 685–698. DOI: 10.1016/j.ejor.2014.04.026.

[37] Timo Gschwind, Nicola Bianchessi, and Stefan Irnich. "Stabilized branch-price-and-cut for the commodity-constrained split delivery vehicle routing problem." In: *European Journal of Operational Research* 278 (2019), pp. 91–104. DOI: 10.1016/j.ejor.2019.04.008.

[38] Masoud Chitsaz, Jean-Francois Cordeau, and Raf Jans. "A branch-and-cut algorithm for an assembly routing problem." In: *European Journal of Operational Research* 282.3 (2020), pp. 896–910. DOI: 10.1016/j.ejor.2019.10.007.

[39] Luciano Costa, Claudio Contardo, and Guy Desaulniers. "Exact Branch-Price-and-Cut Algorithms for Vehicle Routing." In: *Transportation Science* 53.4 (2019), pp. 946–985. DOI: 10.1287/trsc.2018.0878.

[40] Nicos Christofides and Samuel Eilon. "An Algorithm for the Vehicle-dispatching Problem." In: *Journal of the Operational Research Society* 20 (1969), pp. 309–318.

[41] G. Laporte and Y. Nobert. "A Branch and Bound Algorithm for the Capacitated Vehicle Routing Problem." In: *OR Spectr.* 5.2 (1983), 77–85. ISSN: 0171-6468. DOI: 10.1007/BF01720015. URL: https://doi.org/10.1007/BF01720015.

[42] Edward A Silver. "An Overview of Heuristic Solution Methods." In: *Journal of the Operational Research Society* 55.9 (2004). Review, pp. 936–956. DOI: 10.1057/palgrave.jors.2601758.

[43] Franz Rothlauf. *Design of Modern Heuristics: Principles and Application*. 2011th. Natural Computing Series. Springer, 2011.

[44] Thibaut Vidal et al. "Heuristics for multi-attribute vehicle routing problems: A survey and synthesis." In: *Eur. J. Oper. Res.* 231 (2013), pp. 1–21. URL: http://dblp.uni-trier.de/db/journals/eor/eor231.html#VidalCGP13.

[45] Gilbert Laporte. "The vehicle routing problem: An overview of exact and approximate algorithms." In: *European Journal of Operational Research* 59.3 (1992), pp. 345–358. ISSN: 0377-2217. DOI: https://doi.org/10.1016/0377-2217(92)90192-C. URL: https://www.sciencedirect.com/science/article/pii/037722179290192C.

[46] Peng Si Ow and Thomas E. Morton. "Filtered beam search in scheduling." In: *International Journal of Production Research* 26.1 (1988), pp. 35–62.

[47] Michael L. Pinedo. *Scheduling: Theory, Algorithms, and Systems*. 5th. Springer, 2016.

[48] Shangce Gao et al. "Ant colony optimization with clustering for solving the dynamic location routing problem." In: *Applied Mathematics and Computation* 285 (2016), pp. 149–173.

[49] Sampo Kuutti et al. "A Survey of Deep Learning Applications to Autonomous Vehicle Control." In: *IEEE Transactions on Intelligent Transportation Systems* 22.2 (2021), pp. 712–733. DOI: 10.1109/TITS.2019.2962338.

[50] Gopal S. Tandel et al. "A Review on a Deep Learning Perspective in Brain Cancer Classification." In: *Cancers* 11.1 (2019), p. 111. DOI: 10.3390/cancers11010111.

[51] A. H. Bukhari et al. "Fractional Neuro-Sequential ARFIMA-LSTM for Financial Market Forecasting." In: *IEEE Access* 8 (2020), pp. 65110–65121. DOI: 10.1109/ACCESS.2020.2985763.

[52] John T. Houghton. *Global Warming: the complete briefing*. Cambridge University Press, 2015. ISBN: 13 978-0-521-88256-9.

[53] United Nations. "Review of Maritime Transport 2019." In: *United Nations Conference on Trade and Development* (2019). URL: https://unctad.org/system/files/official-document/rmt2019_en.pdf.

[54] V. Masson-Delmotte et al. *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Ed. by V. Masson-Delmotte et al. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press, 2021, p. 2391. DOI: 10.1017/9781009157896.

[55] United Nations Framework Convention on Climate Change (UNFCCC). *Paris Agreement*. https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement. 2015.

[56] Rodrigo Duarte Soliani. "The Concept of Sustainability in the Brazilian Road Freight Transportation Sector." In: *Sustainable Development Dimensions and Urban Agglomeration*. Ed. by Alessandra Battisti and Serena Baiani. Rijeka: IntechOpen, 2021. DOI: 10.5772/intechopen.97280. URL: https://doi.org/10.5772/intechopen.97280.

[57] Thibaut Vidal, Gilbert Laporte, and Piotr Matl. "A concise guide to existing and emerging vehicle routing problem variants." In: *European Journal of Operational Research* 286.2 (2020), 401–416. DOI: 10.1016/j.ejor.2019.10.010. URL: http://dx.doi.org/10.1016/j.ejor.2019.10.010.

[58] G Poonthalir and R Nadarajan. "A Fuel Efficient Green Vehicle Routing Problem with varying speed constraint (F-GVRP)." In: *Expert Systems With Applications* 100 (2018), pp. 131–144. DOI: https://doi.org/10.1016/j.eswa.2018.01.052.

[59] X.F. Yuan S.W. Zhou L.H. Wu Y.N. Wang. *Environmental/economic power dispatch problem using multi-objective differential evolution algorithm*. 2010. URL: https://www.sciencedirect.com/science/article/pii/S0378779610000696.

[60] S Ubeda, FJ Arcelus, and J Faulin. "Green logistics at Eroski: A case study." In: *International Journal of Production Economics* 131.1 (2011), pp. 44–51.

[61] Lijun Geng Jiangfeng Zhang Xueping Li Xiaoqiang Guo Liangce He Zhigang Lu. *Environmental economic dispatch of integrated regional energy system considering integrated demand response*. 2020. URL: https://www.sciencedirect.com/science/article/pii/S0142061519315352.

[62] Paulo Roberto de Oliveira da Costa et al. "A Genetic Algorithm for a Green Vehicle Routing Problem." In: *Electronic Notes in Discrete Mathematics* 64 (2018). 8th International Network Optimization Conference - INOC 2017, pp. 65–74. ISSN: 1571-0653. DOI: https://doi.org/10.1016/j.endm.2018.01.008. URL: https://www.sciencedirect.com/science/article/pii/S1571065318300088.

[63] Gilbert Laporte Tolga Bektaş. *The Pollution-Routing Problem*. 2011. DOI: https://doi.org/10.1016/j.trb.2011.02.004. URL: https://www.sciencedirect.com/science/article/pii/S019126151100018X.

[64] Canhong Lin et al. "Survey of Green Vehicle Routing Problem: Past and future trends." In: *Expert Systems with Applications* 41.4, Part 1 (2014), pp. 1118–1138. ISSN: 0957-4174. DOI: https://doi.org/10.1016/j.eswa.2013.07.107. URL: https://www.sciencedirect.com/science/article/pii/S095741741300609X.

[65] Eurostat. *Waste statistics*. 2023. URL: https://ec.europa.eu/eurostat/statistics-explained/index.php.

[66] Yizi Zhou et al. "Decision Support System for Green Real-Life Field Scheduling Problems." In: Springer. Loughborough, UK, 2017, pp. 355–369. DOI: 10.1007/978-3-319-71078-5_30.

[67] Angie Ramirez-Villamila, Anicia Jaegler, and Jairo R. Montoya-Torres. "Sustainable local pickup and delivery: The case of Paris." In: *Research in Transportation Business & Management* 45 (2022), p. 100692. DOI: 10.1016/j.rtbm.2022.100692. URL: https://www.sciencedirect.com/science/article/pii/S2210539522000342.

[68] Ákos Holló-Szabó and János Botzheim. "Bacterial Memetic Algorithm for Asymmetric Capacitated Vehicle-Routing Problem." In: *Electronics* 11.22 (2022), p. 3758. DOI: 10.3390/electronics11223758. URL: https://www.mdpi.com/journal/electronics.

[69] Konstantinos Koutroumbas and Sergios Theodoridis. *Pattern Recognition*. 4th. Academic Press, 2009.

[70] Asharul Islam Khan and Salim Al-Habsi. "Machine Learning in Computer Vision." In: *Procedia Computer Science* 167 (2020), pp. 1444–1451. DOI: 10.1016/j.procs.2020.03.355. URL: https://doi.org/10.1016/j.procs.2020.03.355.

[71] Artur Istvan Karoly, Robert Fuller, and Peter Galambos. "Unsupervised Clustering for Deep Learning: A tutorial survey." In: *ACTA POLYTECHNICA HUNGARICA* 15.8 (2018), pp. 29–53. ISSN: 1785-8860. DOI: 10.12700/APH.15.8.2018.8.2.

[72] Dong Gong et al. "Memorizing Normality to Detect Anomaly: Memory-augmented Deep Autoencoder for Unsupervised Anomaly Detection." In: IEEE International Conference on Computer Vision (2019), pp. 1705–1714. DOI: 10.1109/ICCV.2019.00179.

[73] Gabriel San Martin et al. "Deep variational auto-encoders: A promising tool for dimensionality reduction and ball bearing elements fault diag-

nosis." In: *Structural Health Monitoring - An International Journal* 18.4 (2019), pp. 1092–1128. ISSN: 1475-9217. DOI: 10.1177/1475921718788299.
[74] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning, Second Edition: An Introduction*. Adaptive Computation and Machine Learning Ser. Bradford Books, 2018. ISBN: 978-0-262-03924-6.
[75] Anne Håkansson and Ronald Lee Hartung. *Artificial Intelligence: concepts, areas, techniques and applications*. 1st. Studentlitteratur AB, 2020. ISBN: 9789144125992.
[76] David Silver et al. "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play." In: *SCIENCE* 362 (2018). ISSN: 0036-8075. DOI: 10.1126/science.aar6404.
[77] Jens Kober, J. Andrew Bagnell, and Jan Peters. "Reinforcement learning in robotics: A survey." In: *INTERNATIONAL JOURNAL OF ROBOTICS RESEARCH* 32.11 (2013), pp. 1238–1274. ISSN: 0278-3649. DOI: 10.1177/0278364913495721.
[78] Meixin Zhu, Xuesong Wang, and Yinhai Wang. "Human-like autonomous car-following model with deep reinforcement learning." In: *Transportation Research Part C: Emerging Technologies* 97 (2018), pp. 348–368. DOI: 10.1016/j.trc.2018.10.024.
[79] Simon Haykin. *Neural Networks: A Comprehensive Foundation*. 1999. ISBN: 978-0-7803-3494-6. DOI: 10.1604/9780780334946.
[80] Kevin L. Priddy and Paul E. Keller. *Artificial Neural Networks: An Introduction*. illustrated edition. SPIE Publications, 2005, p. 180. ISBN: 978-0819459879.
[81] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. Adaptive Computation and Machine Learning Ser. 2016. ISBN: 978-0-262-03561-3.
[82] Diederik P. Kingma and Jimmy Ba. "Adam: A Method for Stochastic Optimization." In: (2014). URL: https://arxiv.org/abs/1412.6980.
[83] Kaiming He et al. "Deep Residual Learning for Image Recognition." In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016), pp. 770–778. DOI: 10.1109/CVPR.2016.90.
[84] H. Ghaziri. "Supervision in the Self-Organizing Feature Map: Application to the Vehicle Routing Problem." In: *Meta-Heuristics*. Springer, 1996. ISBN: 978-1-4612-8587-8. DOI: 10.1007/978-1-4613-1361-8_39. URL: https://doi.org/10.1007/978-1-4613-1361-8_39.
[85] Matthias Schumann and Ralf Retzko. "Solving Vehicle Routing Problems with Self-Organizing Maps." In: 1995.
[86] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. "Pointer Networks." In: *Advances in neural information processing systems*. Vol. 28. 2015, pp. 2692–2700.
[87] Zhaoyang Niu, Guoqiang Zhong, and Hui Yu. "A review on the attention mechanism of deep learning." In: *Neurocomputing* 452 (2021), 48–62.

DOI: 10.1016/j.neucom.2021.03.091. URL: http://dx.doi.org/10.1016/j.neucom.2021.03.091.

[88] Irwan Bello et al. "Neural Combinatorial Optimization with Reinforcement Learning." In: *CoRR* abs/1611.09940 (2016). arXiv: 1611.09940. URL: http://arxiv.org/abs/1611.09940.

[89] Kaiwen Li et al. "Deep Reinforcement Learning for Combinatorial Optimization: Covering Salesman Problems." In: *CoRR* abs/2102.05875 (2021). arXiv: 2102.05875. URL: https://arxiv.org/abs/2102.05875.

[90] Yoav Kaempfer and Lior Wolf. "Learning the Multiple Traveling Salesmen Problem with Permutation Invariant Pooling Networks." In: (2019).

[91] Guillaume Bono et al. "Solving Multi-Agent Routing Problems Using Deep Attention Mechanisms." In: *IEEE Transactions on Intelligent Transportation Systems* 22.12 (2021), pp. 7804–7813. DOI: 10.1109/TITS.2020.3009289.

[92] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. *Neural Machine Translation by Jointly Learning to Align and Translate*. 2016. arXiv: 1409.0473 [cs.CL].

[93] Ming-Hsuan Guo et al. "Attention mechanisms in computer vision: A survey." In: *Computational Visual Media* 8 (2022), pp. 331–368. DOI: 10.1007/s41095-022-0271-y.

[94] Ashish Vaswani et al. "Attention Is All You Need." In: *CoRR* abs/1706.03762 (2017). arXiv: 1706.03762. URL: http://arxiv.org/abs/1706.03762.

[95] Petar Veličković et al. "Graph Attention Networks." In: *International Conference on Learning Representations*. Montreal Institute for Learning Algorithms. 2018. URL: https://openreview.net/forum?id=rJXMpikCZ.

[96] Yanni Dong et al. "Weighted Feature Fusion of Convolutional Neural Network and Graph Attention Network for Hyperspectral Image Classification." In: *IEEE Transactions on Image Processing* 31 (2022), pp. 1559–1572. ISSN: 1057-7149. DOI: 10.1109/TIP.2022.3144017.

[97] Ruiheng Wang et al. "User Identity Linkage Across Social Networks by Heterogeneous Graph Attention Network Modeling." In: *Applied Sciences* 10.16 (2020), p. 5478. ISSN: 2076-3417. DOI: 10.3390/app10165478. URL: https://www.mdpi.com/2076-3417/10/16/5478.

[98] Huan Zhao et al. "Predicting MiRNA-Disease Associations Based on Neighbor Selection Graph Attention Networks." In: *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 20.2 (2023), pp. 1298–1307. ISSN: 1545-5963. DOI: 10.1109/TCBB.2022.3204726. URL: https://ieeexplore.ieee.org/document/36067101.

[99] Cagatay Catal, Hakan Gunduz, and Alper Ozcan. "Malware Detection Based on Graph Attention Networks for Intelligent Transportation Systems." In: *Electronics* 10.20 (2021), p. 2534. DOI: 10.3390/electronics10202534. URL: https://www.mdpi.com/2079-9292/10/20/2534.

[100] José I. Huertas et al. “Real vehicle fuel consumption in logistic corridors.” In: *Applied Energy* 314 (2022), p. 118921. DOI: 10.1016/j.apenergy.2022.118921. URL: http://dx.doi.org/10.1016/j.apenergy.2022.118921.

[101] K Fagerholt, G Laporte, and I Norstad. “Reducing fuel emissions by optimizing speed on shipping routes.” In: *Journal of the Operational Research Society* 61.3 (2010), 523–529. DOI: 10.1057/jors.2009.77. URL: http://dx.doi.org/10.1057/jors.2009.77.

[102] John L. Sullivan, Geoffrey M. Lewis, and Gregory A. Keoleian. “Effect of mass on multimodal fuel consumption in moving people and freight in the U.S.” In: *Transportation Research Part D: Transport and Environment* 63 (2018), pp. 786–808. DOI: 10.1016/j.trd.2018.06.019.

[103] Jingwen Li et al. “Deep Reinforcement Learning for Solving the Heterogeneous Capacitated Vehicle Routing Problem.” In: *IEEE Transactions on Cybernetics* 52.12 (2022), pp. 13572–13585. DOI: 10.1109/TCYB.2021.3111082.

[104] Remiks. *Strategi 2021-2030*. 2021.

[105] *Drag Coefficient (CD)*. URL: https://vbi.truck.volvo.com/portal/perfman/020_terminology/drag_coefficient_(cd).htm.

[106] Xiuyu Liu et al. “Evaluation of Vehicle Braking Performance on Wet Pavement Surface using an Integrated Tire-Vehicle Modeling Approach.” In: *Transportation Research Record Journal of the Transportation Research Board* 2673.3 (2019), pp. 295–307. DOI: 10.1177/0361198119832886.

[107] Oscar Delgado, Felipe Rodriguez, and Rachel Muncrief. “Fuel efficiency technology in European heavy-duty vehicles: Baseline and potential for the 2020–2030 timeframe.” In: (July 2017).

[108] Remiks. *Personal communication*. Conversation with Remiks on 15.05.2023.

[109] Veerabhadra Basam. “IJESRT International Journal of Engineering Sciences Research Technology Performance and Emission Characteristics of Waste Plastic Pyrolysis Oil in D.I Single Cylinder Diesel Engine.” In: (Sept. 2015).

[110] *Heat Values of Various Fuels*. Website. Accessed on 18th May. World Nuclear Association. URL: https://world-nuclear.org/information-library/facts-and-figures/heat-values-of-various-fuels.aspx.

[111] *HEATING VALUES OF HYDROGEN AND FUELS*. Website. Accessed on 18th May. ESSOM CO., LTD. URL: https://chemeng.queensu.ca/courses/CHEE332/files/ethanol_heating-values.pdf.

[112] Benjamin G. Harvey, Walter W. Merriman, and Thomas A. Koontz. “High-Density Renewable Diesel and Jet Fuels Prepared from Multicyclic Sesquiterpanes and a 1-Hexene-Derived Synthetic Paraffinic Kerosene.” In: *Chemistry Division, Research Department, Naval Air Warfare Center Weapons Division (NAWCWD), United States Navy* (2023). URL: https://pubs.acs.org/EF.

[113] Press Information. *Fuel Use - How Low Can You Go?* Volvo - Construction Equipment Global. 2018. URL: https://www.volvoce.com/global/en/news-and-events/news-and-stories/2018/fuel-use-how-low-can-you-go/.

[114] Paul Breeze. *Power Generation Technologies*. English. 1st ed. Elsevier Science & Technology, 2005, p. 289. ISBN: 9780750663137.

[115] Scott Jenkins. *CRC Handbook of Chemistry and Physics*. 93rd. Vol. 119. 9. Chemical Engineering, 2012.

[116] Carsten Hoever and Wolfgang Kropp. "A model for investigating the influence of road surface texture and tyre tread pattern on rolling resistance." In: *Journal of Sound and Vibration* 351 (2015), pp. 161–176. DOI: 10.1016/j.jsv.2015.04.014. URL: http://www.elsevier.com/locate/jsvi.

[117] Jerzy A. Ejsmont et al. "Road texture influence on tyre rolling resistance." In: *Road Materials and Pavement Design* 18 (2017), pp. 181–198. ISSN: 1468-0629. DOI: 10.1080/14680629.2016.1160835.

[118] A. Poot, G. Kant, and A. P. M. Wagelmans. "A Savings Based Method for Real-Life Vehicle Routing Problems." In: *The Journal of the Operational Research Society* 53.1 (2002), pp. 57–68. URL: http://www.jstor.org/stable/822879.

[119] Maria Battarra, Günes Erdogan, and Daniele Vigo. "Exact Algorithms for the Clustered Vehicle Routing Problem." In: *Operations Research* 62 (Feb. 2014), pp. 58–71. DOI: 10.1287/opre.2013.1227.

[120] Javier Alonso-Mora et al. "On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment." In: *Proceedings of the National Academy of Sciences* 114.3 (2017), 462–467. DOI: 10.1073/pnas.1611675114. URL: http://dx.doi.org/10.1073/pnas.1611675114.

[121] Golman Rahmanifar et al. "Heuristic approaches to address vehicle routing problem in the Iot-based waste management system." In: *Expert Systems with Applications* 220 (2023), p. 119708. ISSN: 0957-4174. DOI: https://doi.org/10.1016/j.eswa.2023.119708. URL: https://www.sciencedirect.com/science/article/pii/S0957417423002099.

Appendix

Current Waste Collection Zones

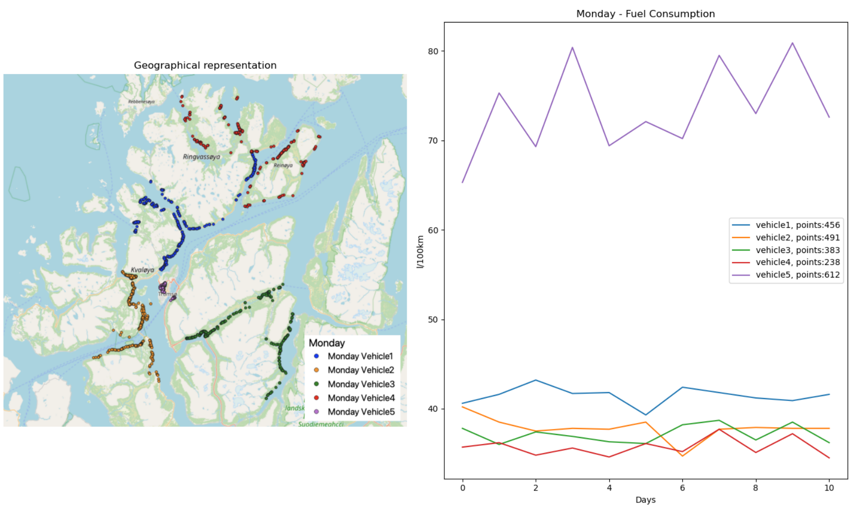

Figure 9.1: Monday waste collection routes. Geographical representation of the collection routes with their fuel consumption. Vehicle 5 has the highest fuel consumption and collects in the city area.

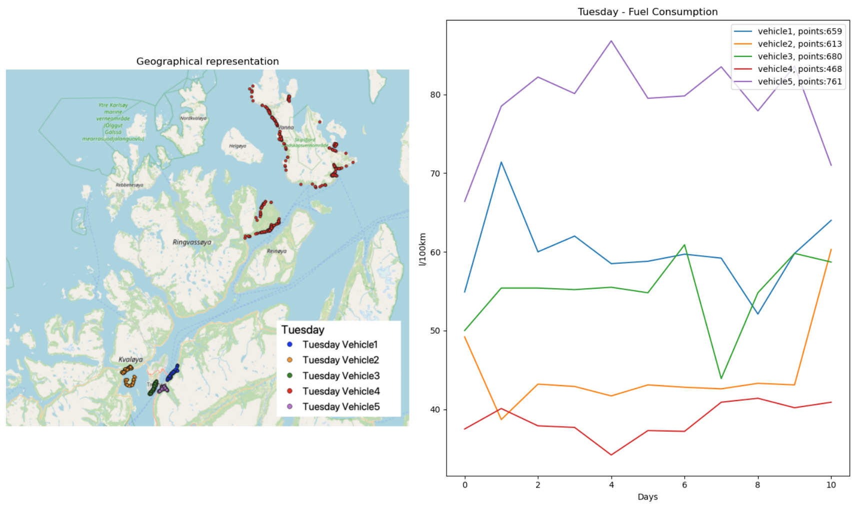

Figure 9.2: Tuesday waste collection routes. Geographical representation of the collection routes with their fuel consumption. Vehicle 4 has the lowest fuel consumption and collects in the district area.

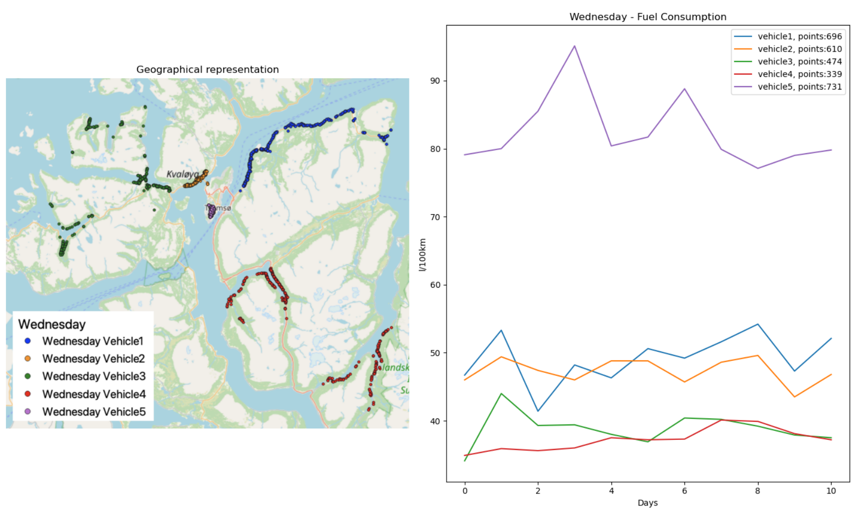

Figure 9.3: Wednesday waste collection routes. Geographical representation of the collection routes with their fuel consumption. Vehicle 5 has the highest fuel consumption and collects in the city area.

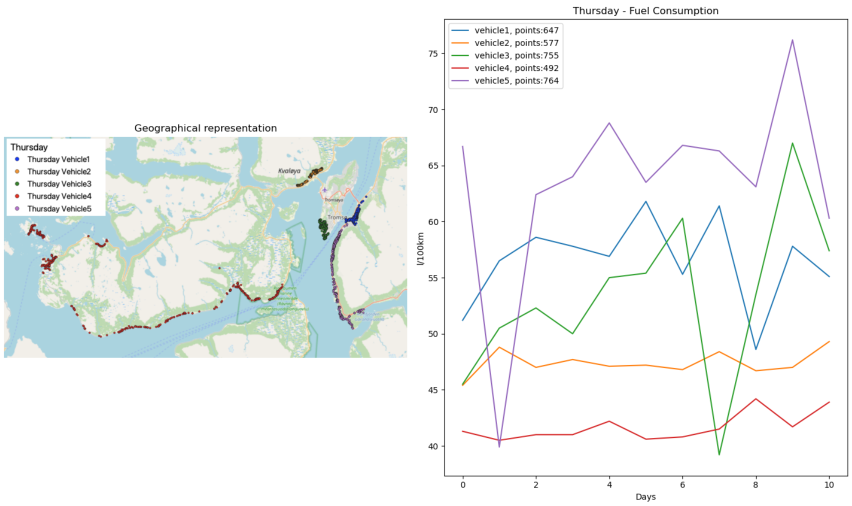

Figure 9.4: Thursday waste collection routes. Geographical representation of the collection routes with their fuel consumption. Vehicle 4 has the lowest fuel consumption and collects from the fewest households in the district areas.

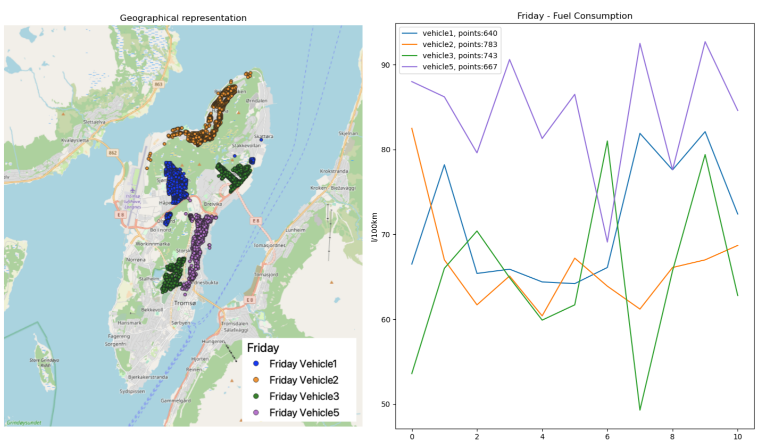

Figure 9.5: Friday waste collection routes. Geographical representation of the collection routes with their fuel consumption. All vehicles collect in the city area, and the overall fuel consumption is greater than on other days.