MSc Thesis
Industrial Engineering and Management

Improving the medium-term schedule in a home health care context with different demand frequencies

Janiek Smulders

Exam Committee:
Lead supervisor: Dr. D. Guericke
Second supervisor: Dr. S. Rachuba
Company lead supervisor: J.B.H. Vos

July, 2024

Department of Behavioural,
Management and Social Sciences

UNIVERSITY OF TWENTE.

ExpertCare
Verplegen met een hart

This thesis is written as part of the graduation assignment of the master programme of Industrial Engineering and Management at the University of Twente. The graduation assignment is conducted at ExpertCare.

University of Twente
Industrial Engineering and Management
Drienerlolaan 5
7522 NB Enschede
Tel.: 053 489911

ExpertCare
Injection Team
Perkinsbaan 14
3439 ND Nieuwegein
Tel.: 030 6008722

Improving the medium-term schedule in a home health care context with different demand frequencies

Date: July, 2024
Number of pages: 50
Number of appendices: 8

Author:
J.L. Smulders
Master student Industrial Engineering and Management, University of Twente

Supervisors University of Twente:
Dr. D. Guericke (1st supervisor)
Dr. S. Rachuba (2nd supervisor)

Supervisors ExpertCare:
J.B.H. Vos

Preface

This Master Thesis "Improving the medium-term schedule in a home health care context with different demand frequencies" concludes my Master Industrial Engineering and Management at the University of Twente. I would like to thank everybody who supported me to complete this thesis.

First and foremost, I would like to thank Bernard Vos for his warm reception at ExpertCare and for his great assistance throughout the process. The guidance you provided during this project were very helpful. I would also like to extend my gratitude to other staff at ExpertCare, for being willing to help me and provide me with data throughout the duration of my graduation assignment.

From the University of Twente, I would like to thank Daniela Guericke for giving me this opportunity and her enthusiasm and interest in the research. I enjoyed our meetings and valued your constructive feedback and the insights you gave. Also, I would like to thank Sebastian Rachuba for helping and providing feedback.

Lastly, I would like to thank my friends and family for their support.

I hope you enjoy reading this report.

Janiek Smulders

Utrecht, July 2024

Management summary

This research is executed at ExpertCare. ExpertCare is an organisation that provides home care services in several regions in the Netherlands. This study focuses on the Injection Team of ExpertCare. ExpertCare strives to improve its planning process of appointments in order to decrease the travel time. The Injection Team serves a large area with approximately 900 clients with a large variety in their visitation frequencies. We focus on the medium-term schedule of the Injection Team with the main aim to reduce the the travel time. For this we created an algorithm that schedules each client for the medium-term schedule. Additionally, we analyze the effects of having an extra nurse, the influence of the flexibility of planning clients' appointments and the impact of not using the option to move clients to other weeks.

Firstly, the current planning system of the Injection Team of ExpertCare is analyzed. At the moment the planning and scheduling is done manually by the planners of the team. This leads to an inefficient schedule, since the number of clients and the region in which they live is large. After analyzing the data of the current client pool, it is evident that there are three visitation frequencies with the largest number of clients, namely the monthly, quarterly and half-yearly frequency. Moreover, it is noticed that there are three places with a high density of clients, namely Utrecht, Amersfoort and Zeist. The average travel time per nurse per day is found to be 2.63 from the data of 2023. Additionally, the average number of appointments and the average number of nursing hours are 8.02 and 5.92, respectively.

In the literature review we investigated which solution approach fits best to the context of ExpertCare. It was established that the problem of the Injection Team can be described as the home health care routing and scheduling problem. It is clear from the literature that the visitation frequency of clients of the Injection Team is considerably lower than that in the literature. This would advocate for a planning horizon much longer than usually considered in the literature. Therefore, a heuristic approach is chosen to solve the problem, namely the adaptive large neighborhood search (ALNS). Firstly, a greedy heuristic assigns a scheme to each client to create an initial solution. Then the ALNS further improved this initial solution by destroying some scheme of clients and assigning new ones. The performance of the ALNS is compared with an exact model and it was found that it provides results close to the exact model. Additionally, the ALNS is able to find solutions on more complex instances in a reasonable time compared to the exact approach.

We investigated different scenarios, including freely scheduling clients, allowing complete flexibility in planning appointments, adding an additional daily route, giving clients the option to choose one or two time slot(s), restricting the ability to move clients to other weeks and increasing the number of clients to 1400. The results are shown in Table 1 on the next page. The third column gives the total travel hours over half a year in hours and the fourth column gives the run time in seconds. The fifth column gives the average number of traveling hours per nurse per day and the sixth column gives the average number of appointments per nurse per day. The last column gives the average number of nursing hours per nurse per day. The results show that when clients were freely scheduled, the total travel time over half a year was 573.67 hours, averaging 2.21 hours per day per nurse. Comparing this to the values found for the Injection Team in 2023 this gives a reduction in travel time of 16.14%. This increases the productivity of ExpertCare by 1.7% to 60.40%. However, no clear patterns on how to cluster clients effectively were identified. As a result, we are unable to offer guidelines for assigning clients to specific clusters based on their visitation frequency, location, or appointment week. The complex interplay between these parameters prevents clear categorization. Therefore, to improve the medium-term planning and the productivity, a more advanced planning system is necessary. Furthermore, other experiments have been conducted to see the influence of those instances. In experiment number 2 clients can be planned in any week and this gives a total travel time of 530.14 hours. This gives a reduction of 153.94 hours compared to the Injection Team in 2023 and this reduction in the travel time shows the effectiveness of using the possibility of moving clients from their starting weeks.

Adding an extra daily route increased the total travel time due to the algorithm spreading clients too widely across time slots. Another scenario allowed clients to choose their preferred time slots, resulting in a total travel time of 785.87 hours over half a year, a significant increase compared to free scheduling. This is expected as it reduces scheduling flexibility. A similar experiment was also conducted in which clients could choose two time slots and this gave a total travel time of 727.27 over half a year. Furthermore, fixing clients to their starting weeks also increased total travel time compared to freely planning clients. This is again due to the algorithm's inefficient planning clients spread out over all time slots. Lastly, increasing the number of clients to 1400 gives an increase of 244.28 compared to the Injection Team in 2023 which only served 879 clients.

Table 1: Results of all the experiments

Exp. num.	**Description of experiment**	**Solution of ALNS**	**Run time**	**Avg. travel.**	**Avg. num. app.**	**Avg. nurs.**
-	Injection Team	684.08	-	2.63	8.02	5.92
1	Flexible clients	573.67	3824.38	2.21	9.26	6.94
2[1]	Complete flexibility	530.14	10931.24	2.04	9.25	6.94
3	Number of daily routes	610.68	4650.58	1.88	7.39	5.54
4	Clients choose one slot	785.87	2039.16	2.02	6.18	4.64
5	Clients choose two slots	727.27	2817.76	1.89	6.27	4.70
6	No movement of clients	607.88	3777.97	1.71	6.86	5.14
7	More clients	928.36	5045.50	2.38	9.74	7.30

[1] Infeasible in practise.

Recommendations given to ExpertCare are considering using a different way of planning clients in the medium-term. Given that only four clients have a weekly frequency, relying on a weekly planning method may be inefficient. Instead, considering a monthly planning method or a half-yearly planning system, would be recommended. Another strategy is employing a medium-term planning, where each client is assigned specific days for all future appointments. When new clients join, an algorithm can identify the best time slots based on the current schedule, prioritizing days with other nearby clients. This would reduce time window and working time violations by recognizing fully booked time slots. Additionally, allowing clients to book their appointments online could save time for planners and nurses, despite increasing total travel time. As the number of clients grows and more daily routes are added, this option may become more viable. So it is also recommend for the Injection Team to establish for themselves what they want to offer clients, whether they would like the most efficient planning in terms of traveling time or let the clients choose their own appointment moment.

This thesis offers several practical contributions. First, it conducted a comprehensive context analysis of the Injection Teams' current situation, providing valuable insights into their client pool. It clarified for ExpertCare which data is easily accessible and which is challenging or unavailable. Additionally, a prototype algorithm was developed to create an efficient medium-term plan for ExpertCare's current clients. The research concluded that due to the complex interplay of client parameters, clear patterns for clustering clients were not found. Therefore, it is recommended that ExpertCare adopt a more advanced planning system to enhance productivity. Additionally, it is recommended to have three daily nurse instead of two, to decrease the workload of the nurses and due to the consistent increase in number of clients over time.

Contents

1 Introduction

This thesis aims to increase the productivity of the Injection Team of ExpertCare. The following chapter introduces the research conducted at ExpertCare. Section 1.1 gives a background of ExpertCare and the research motivation. Then in Section 1.2 the problem identification is given and in Section 1.3 the research questions are described. Lastly, in Section 1.4 the research design is discussed.

1.1 Background

1.1.1 The company

ExpertCare is a specialized home care organization that provides high-quality skilled medical and nursing home care services. These services include hospital-at-home care for adults and children in the comfort of their own homes or in Villa ExpertCare centres, ensuring personalized and comprehensive care. Villa ExpertCare provides specialised nursing care for children between 0 and 18 years old, offering support for those with chronic or severe illnesses. These are located in Vleuten, Rijswijk, Waalre and Wezep. The main office of ExpertCare is in Nieuwegein. ExpertCare employs over 300 staff members across its locations.

This study centres on the Injection Team, comprising six nurses delivering injections to more than 800 clients at their homes with frequencies such as weekly, monthly, quarterly, or semi-annually. This is done on behalf of six hospitals. The majority of these clients are diagnosed with cancer. The Injection Team serves clients residing in a broad area spanning from Gouda to Barneveld, encompassing both rural and (sub)urban areas. The Injection Team at ExpertCare was set up in 2022. With new clients joining and existing ones leaving on a weekly basis, the route planning is never consistent. The routes on a daily basis are made manually such that the travel time is minimized given all the client appointments of that day. However, doing this manually requires a lot of time and it is more prone to errors which increases the travel time. In addition, there is potential to further optimize the clustering of the clients on a weekly basis to reduce overall travel time.

1.1.2 Research motivation

The goal of this research is to minimize the nurse travel time. This would enable the Injection Team to assist more individuals and spend less time on traveling to the clients and on planning. Also, client preferences should be considered, such as enabling stability in visitation time slots and enabling continuity of care, meaning that the client is helped by the same nurse as much as possible. ExpertCare is especially interested in the clustering of its clients, since it pursues to transition from 2 daily routes to 3 daily routes.

1.2 Problem identification

Figure 1 illustrates the problem cluster demonstrating the interconnectedness of various problems and their underlying causes. Some of these aspects will be elaborated on further below.

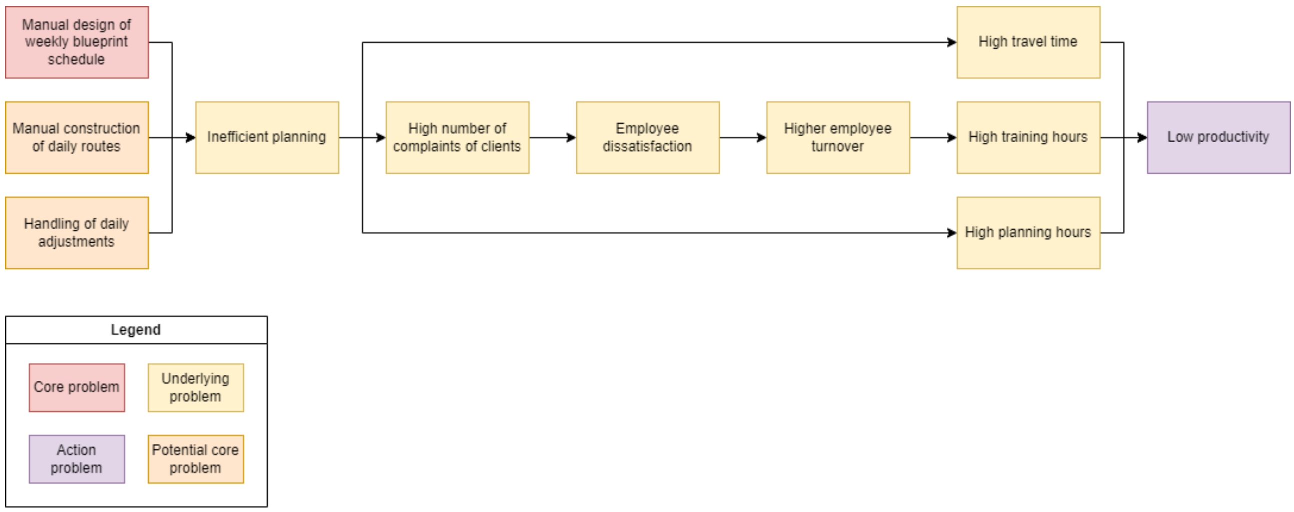

Figure 1: Problem cluster

Low productivity

The low productivity of the Injection Team is experienced by the management. The target productivity score is set at 62%. However, in 2023 an average productivity score of 58.7% was reported from internal reports which is below the target. This gives a difference of 3.3% which amounts to 226.57 hours. When considering that each year has 260 working days, this would amount to 0.87 hours per day. Although the Injection Team was established in 2022, it has consistently failed to achieve the target score. Therefore, it is an important issue to tackle for the Injection Team. The productivity score is calculated by dividing the direct working hours by the total amount of working hours. The total amount of hours consists of the direct work hours and the indirect work hours. The indirect work hours comprise of administration, yearly performance review meetings, oncology meetings, training hours, office hours (other), travel time, planning hours, educational time, team discussions, on-call service, work supervision hours, holiday hours and sick leave. The five indirect work hours categories with the highest impact on the productivity in 2023 are the travel time (17.67%), the holiday hours (11.56%), the training hours (3.50%), the office hours (other) (2.46%) and the planning hours (2.11%). The office hours (other) refer to the time during which the nurse retrieves medication from the storage. The holiday hours and the office hours (other) are left out of the problem cluster since this cannot be influenced.

High travel time

The high travel time is caused by inefficient scheduling of appointments over a week and the daily constructed routings. This would mean the nurses spend more time on the road and have less time for helping clients.

High training hours

The increased client complaints can contribute to employee dissatisfaction, resulting in higher employee turnover rates. Receiving many complaints during a working day negatively affects the nurses and may therefore lead to a higher employee turnover. This means new employees will have to be hired which leads to higher training hours since they must follow a training first.

High planning hours

The planners are the ones who make the daily routings for the nurses. Since this is done manually, this requires a lot of time of the planners. Moreover, many daily adjustments in the schedule are necessary and since the planning is done manually these disruptions in the schedule, takes up a lot of time of the planners. The planners themselves are also nurses of the Injection Team. Therefore, if less time is spent on making the planning, more time can be used to serve clients.

Manual planning and scheduling
The inefficiency in the planning is caused since the planning and scheduling of appointments and routes are done manually. The planning and scheduling consist of two parts. Firstly, a blueprint for the weekly schedule of appointments is used. This gives an indication at what moment of the week (morning or afternoon) appointments can be planned in a certain region (cluster). The routing is not included in the blueprint and we will elaborate more on the blueprint in Section 1.4. This blueprint is based on the planners' knowledge and intuition. It does not take into account how the clients are distributed over the region. Secondly, the planning of the daily routes of the nurses must be made. This is done by the planners manually in a planning system called "Nedap ONS" given the appointments of a day. This simply works by dragging the appointments to change the order and then examine what a good order could be.

Handling of daily adjustments
Many daily adjustments occur due to for example, new clients or discharge of a current client or a change in the frequency of medication. These disruptions in the schedule have to be manually adapted in the current schedule, which takes up a lot of time for the planners. This is considered as a potential core problem, since a more efficient way of handling these changes in the schedule could be constructed.

The core problem
As indicated in the problem cluster depicted in Figure 1 the core problem of this research will be formulated formally in this section. We selected this core problem over the others, because we believe that this core problem will have the greatest impact effect. The core problem revolves around the manual creation of the blueprint relying solely on the planners' knowledge and intuition. This method lacks efficiency and leads to prolonged travel times for the nurses. Figure 2 exemplifies the design of a blueprint. The table indicates in which region the nurse will attend clients in the morning and in the afternoon. The coloured regions in the map give an indication of what it covers and these are called clusters. The clusters are based on the geographic location of clients and the planners' knowledge and intuition. Thus, to design a blueprint for the daily routes, an efficient method to cluster the clients is necessary. As can be seen in Figure 2 clusters can be overlapping. This means that a client living in two clusters has more appointment options available to them.

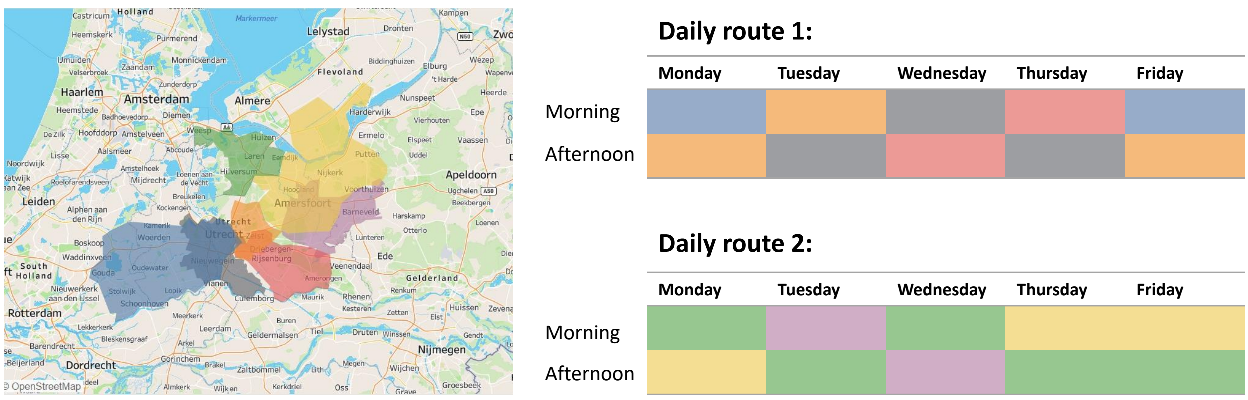

FIGURE 2: Blueprint of daily routes

1.3 Research questions

In order to solve the core problem, we formulate and answer the main research question:

How can we improve the medium-term planning in order to enhance the productivity of the Injection Team of ExpertCare?

To answer the main research question, several sub-questions are formulated to provide a framework for the research.

1) Analysis of the current situation
First, we will analyse the current scheduling process for the client appointments within the system and what current clusters are used. In addition, we will evaluate its impact on the overall performance. To comprehend the existing process, we engage in interviews with the company's planning staff. Through these interviews, we gather insights into the current design of the process, its requirements, perceived constraints and its performance. Thus, we will answer the following questions:

1. How is the current planning of the Injection Team organized?
   (a) How does the Injection Team currently plan appointments?
   (b) What is the current blueprint schedule used to plan appointments?
   (c) What is the distribution of visitation frequency of the clients?
   (d) How are the stakeholders involved and what do they want?
   (e) What is the current performance of the planning system?

2) Literature review
After assessing the current situation, our next step involves exploring suitable methods and solution approaches that align with our requirements. Subsequently, we will compare these options to determine the most suitable one tailored to our specific needs. Through a comprehensive literature review and analysis of its findings, we aim to establish a robust foundation for our forthcoming solution approach.

2. Which solution approach fits best to the context of ExpertCare?
   (a) Which aspects of the home health care routing and scheduling problem have been addressed in previous research?
   (b) What other related problems exist in the literature?
   (c) What are the solution approaches employed in the literature?

3) Solution design
Once the conducted literature research is done, we will develop a model used as a solution approach. The found method from the literature research should still be tailored to our problem context.

3. How should the solution approach be designed?
   (a) What are the characteristics of the problem that needs to be solved?
   (b) What are the assumptions of the model used as a solution approach?
   (c) How can we adapt the model used as a solution approach to our problem context?

4) Experiments
Once we have developed the solution approach, we can experiment with different settings to see how the solution approach performs. Things to consider are for example, the flexibility of clients. It is important to actively involve the planners in this part since they are the end-user of the product.

4. How can we tune the solution approach to ensure high-quality outcomes?
   (a) How should the parameters of the model be set to generate high-quality solutions?
   (b) How should the experimental settings be set?

5) Results analysis
After performing all the different experiments, we can analyze the influence of the different instances on the performance of the model.

5. What is the influence of the different instances on the performance of the model?
   (a) What is the performance of the model?
   (b) What is the effect of freely planning clients?
   (c) What is the influence of allowing complete flexibility in planning appointments?
   (d) What is the influence of adding an additional daily route?
   (e) What is the effect of clients choosing their own appointment moment?
   (f) What is the influence of clients choosing two appointment moments?
   (g) What is the influence of not allowing clients to be moved to a different week?
   (h) What is the effect of having more clients?

6) Evaluation
Based on the results obtained in the previous phase, we can draw the conclusions and give recommendations to ExpertCare. The deliverables to ExpertCare are guidelines on which criteria the clustering of patients can be done and recommendations on how to improve their productivity.

6. What are the conclusions we can draw and the recommendations we can give to ExpertCare based on the results of this research?
   (a) What conclusions can be drawn based on the results compared to the current performance of ExpertCare?
   (b) What recommendations can we give to ExpertCare?
   (c) What are other possibilities for future research?

1.4 Research design

Based on the research questions in Section 1.3, we divide our research in six phases in order to answer the main research question. The six phases are shown in Figure 26 in Appendix A and each phase represents a chapter in this report. First, we analyze the current situation of the Injection Team using both qualitative and quantitative methods which is covered in Chapter 2. Following this, we conduct a literature review to analyze, compare, and select established solution methodologies from existing literature. This is discussed in Chapter 3. Next, we design a solution by adjusting the chosen approach to fit the specific context of the study which is covered in Chapter 4. Subsequently, we conduct experiments to fine-tune the parameters of the solution approach and establish the experimental settings and this is covered in Chapter 5. Then, we analyze the performance of the solution under different experimental settings which is covered in Chapter 6. Finally, we evaluate the research findings and their implications for ExpertCare and this is covered in Chapter 7.

2 Context analysis

This chapter analyses the current planning system used by the Injection Team of ExpertCare in order to answer the first research question: *How is the current planning of the Injection Team organized?*. Section 2.1 discusses the current planning system used by the Injection Team of ExpertCare. Additionally, in Section 2.2 the clusters of the Injection Team are analyzed and in Section 2.3 the different visitation frequencies of the clients are analyzed. Moreover, Section 2.4 contains the stakeholder analysis to identify all the demands of the stakeholders. In addition, Section 2.5 discusses the performance indicators of the planning system. Lastly, we will conclude this chapter in Section 2.6.

2.1 Current planning system

This section discusses the current planning system employed by the Injection Team of ExpertCare. Firstly, the nurses and clients are introduced. Then the planning of appointments and the routing is discussed.

Nurses
At the beginning of this project the Injection Team comprised six nurses, including one designated as the team coach. While the coach primarily does not serve clients, she occasionally assists when needed. From the remaining 5 nurses, three work one day a week, while two work three days a week. In case of insufficient number of nurses, the coach or nurses from other teams step in, a situation currently occurring weekly. For the remainder of this paper, we will refer to nurses who are not part of the Injection Team, excluding the coach, and who step in to assist clients of the Injection Team, as substitute nurses.

Clients
The clients of the Injection Team of ExpertCare are patients that require medication via injections. The hospital will refer them to ExpertCare such that the medication can be given at the client's home. Once the client has been forwarded to ExpertCare, an intake appointment will be made. Since new clients register and current clients leave the Injection Team on a daily basis the number of clients is never stable. In Figure 3 the evolution of the number of clients over time indicates that there is a clear growth in the number of clients over time.

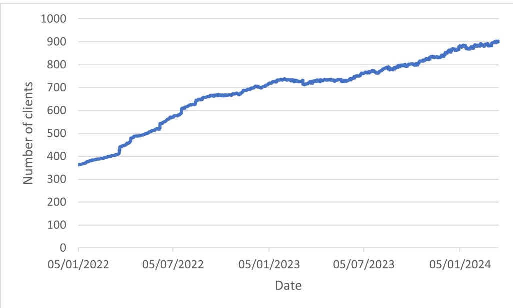

FIGURE 3: Number of clients of Injection Team over time

Appointment
There are two types of appointments, namely an intake appointment and a normal appointment. An intake appointment, which is the initial appointment for new clients, requires more time. During this appointment, the nurse educates the client about their medication and ExpertCare, and gathers necessary information about the client. These appointments last for 75 minutes. Normal appointments encompass all subsequent visits and have a duration of 45 minutes. Both types of appointments include 10-15 minutes of pre-appointment document review and approximately 5 minutes of post-appointment reporting. Consequently, nurses spend approximately 25-30 minutes for normal appointments and 55-60 minutes for intake appointments at the client's residence. The pre-appointment document review

is often done the day before and the post-appointment reporting is often done immediately after helping the client. So the time that should be considered for planning the appointments within a day is approximately 30 minutes.

The nurse always calls the client approximately 20-30 minutes before the appointment to let them know they are coming. When the nurse has arrived, the nurse assesses the client's health and administers medication. Since the clients are referred by the hospital the assessment cannot be negative. Subsequently, the next appointment is scheduled, considering the client's care frequency. Care frequencies include weekly, bi-weekly, tri-weekly, four-weekly, six-weekly, eight-weekly, quarterly or semi-annually. A new appointment is set in the week in which medication is needed by considering the blueprint schedule as depicted in Figure 2. Based on the region where the client lives, several options for an appointment are available. Then the client can choose their preferred option or choose an alternative if needed. The blueprint is based on a specific clustering of the clients which will be elaborated on in Section 2.2.

Daily routes construction
A few days prior to the start of the week, a planner utilizes the "Nedap ONS" software to manually create daily routes based on scheduled appointments. The software presents a list of appointments for each day in a particular order, along with corresponding routing information. Planners can easily adjust the order of appointments by dragging them within the software interface, visually observing how routing changes occur. Planners strive to arrange routes logically within the software. Client-specific preferences for time slots are inserted as notes in the system, but these are frequently overlooked by planners, resulting in these preferences not being considered.

2.2 Clustering analysis

In September of 2023 the clustering of clients by region was introduced by one of the planners of the Injection Team. The clients' location of the Injection Team are shown in a map in Figure 4. The percentages mentioned are from the total number of visits or total number of clients. The number of clients is not equal to the number of visits in half a year due to the different frequencies of the clients. In total there are 7 clusters:

- **Cluster 1:** covers the largest area of all the clusters. It stretches from Gouda to Utrecht. The majority of their clients live in Utrecht. It is visited twice in the morning. This cluster contains approximately 126 (14%) clients that in total need care 361 (15%) times over half a year.
- **Cluster 2:** lies centrally and covers Bunnik and Zeist among other places. The cluster is visited once in the morning and twice in the afternoon. This cluster has around 99 (11%) clients. Over half a year time they require about 265 (11%) visits in total.
- **Cluster 3:** is located in the south and also covers Bunnik and Zeist among other places. It is visited once in the morning and once in the afternoon and it contains approximately 150 (17%) clients. In total they require 407 (16%) visits in half a year.
- **Cluster 4:** covers the area around the office and Utrecht. It is visited once in the morning and twice in the afternoon and it has around 106 (12%) clients. In total 307 (12%) visits are made over half a year time to this cluster.
- **Cluster 5:** is located in the north and covers Hilversum and Huizen among other places. It is visited twice in the morning and thrice in the afternoon and it contains approximately 103 (12%) clients. They need to be visited 410 (17%) times in half a year in total.
- **Cluster 6:** covers Amersfoort and the northern region of Amersfoort. It is visited twice in the morning and once in the afternoon and it contains around 193 (22%) clients. They need about 491 (20%) of visits over half a year time in total.
- **Cluster 7:** covers Amersfoort as well and the southern region around Amersfoort. It is visited once in the morning and once in the afternoon and it has approximately 102 (12%) clients. Approximately 226 (9%) visits need to be made to all the clients over half a year.

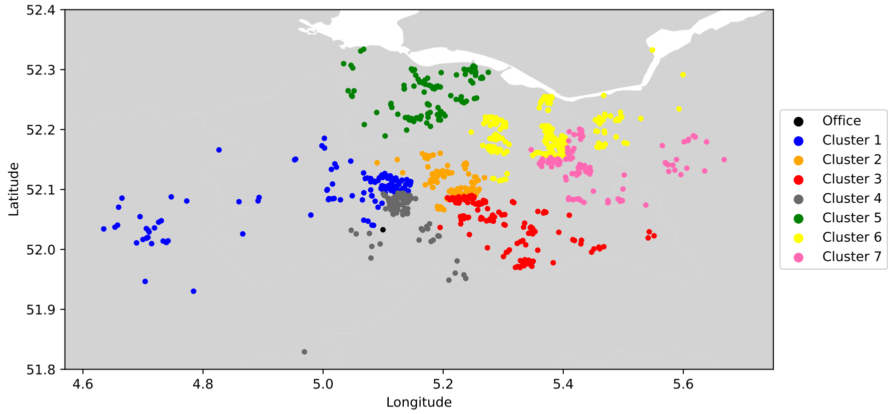

FIGURE 4: Map of client locations

Some clusters overlap slightly, namely cluster 1 & 4, cluster 2 & 3 and cluster 6 & 7. The clients in the overlapping clusters are divided evenly among those clusters for our analysis. Moreover, the data on the number of clients is based on the data as of 19/03/24.

To actually see which clients are in which cluster, it is assumed that the day of the last appointment of each client is repeated for every number of weeks in which they need an injection. So the blueprint shown in Figure 2, is translated to the following clustering, shown in Figure 5. This shows for each client when they are scheduled and since two nurses work on a day, there are two options. So Monday Morning 1 and Monday Afternoon 1 is done by one nurse and Monday Morning 2 and Monday Afternoon 2 is done by another nurse. This is the same for the other days of the week. In Table 10 in Appendix C the number of clients and the average visits in half a year per cluster is displayed.

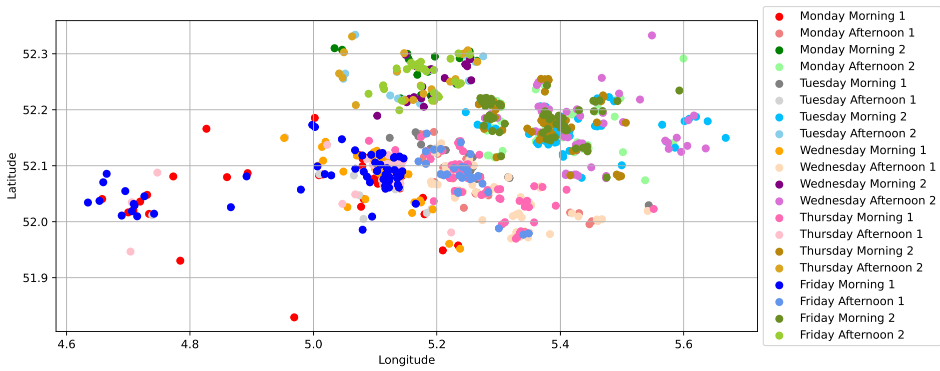

FIGURE 5: Clustering of clients of current schedule of the Injection Team

2.3 Frequency analysis

As mentioned before there are eight different appointment frequencies that clients can have. In Table 2, the number of clients and number of visits in half a year per frequency is displayed. Furthermore, in Figure 6 the map with the clusters are depicted with the according care frequency of each client. From the analysis, it can be noticed that there are 3 main groups of frequencies that dominate. The frequencies of 4-, 12- and 26-weekly amount to 96% of all the clients and 87% of the total number of visitations over a year. The 12-weekly frequency contains the most clients (44%) and the monthly frequency amounts to the most visitations (39%). The 4-weekly frequency is scattered over the whole map, whereas the 12-weekly frequency has about 70% of their clients in clusters 5, 6 and 7 (the north-east on the map in Figure 6). Additionally, the 26-weekly frequency has almost half of their clients in cluster 2 and 3 (centre and south-east on the map in Figure 6). All the other frequencies are scattered over the map. Additionally, Figure 7 illustrates the density of all the clients highlighting the most frequented places. The darker colour indicates a higher density of visitations. We can notice three places in Figure 7 that are the most visited. These places are Utrecht, Amersfoort, and Zeist, with 157, 104, and 84 clients situated in each respective location and total average amount of visits over half a year 440, 254 and 247, respectively. These 3 locations are the 3 overlapping parts of the clusters.

TABLE 2: Number of clients and visitations per frequency

Frequency	Number of clients	Approximate number of visits in half year
Weekly	4	104.00
2-weekly	9	117.00
3-weekly	8	69.54
4-weekly	148	967.69
6-weekly	3	13.15
8-weekly	7	23.15
12-weekly	384	856.62
26-weekly	316	316.00
Total	**879**	**2467.15**

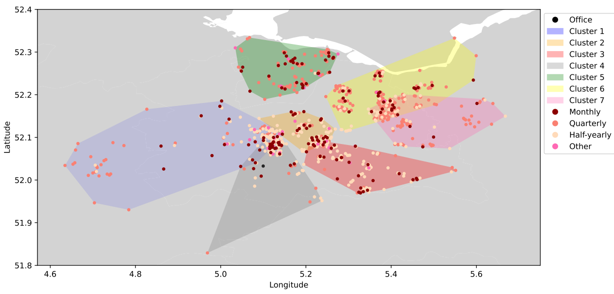

FIGURE 6: Map of client locations and frequency

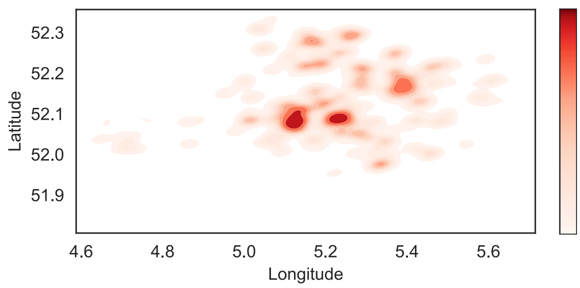

FIGURE 7: Density of the clients

2.4 Stakeholder analysis

The stakeholders in the planning process of the Injection Team of ExpertCare are the clients, the nurses of the Injection Team and the board of ExpertCare. All three stakeholder groups have demands of the planning system that they wish to be fulfilled. After some discussions with members of the Injection Team, the wishes per stakeholder group are listed here below:

The clients:

- The clients want to be scheduled during their preferred time window. Currently, the nurses can either arrive at the client in the morning from 9:00 to 13:00 or in the afternoon from 12:00 to 16:00. The overlap in the time slots allows for more flexibility for the nurses of the Injection Team. For example, if the afternoon is busier than the morning they can already start helping afternoon clients at 12:00 and vice versa.
- Many clients also prefer that they are helped by the same nurse for each appointment as much as possible.

The nurses:

- The nurses prefer to have balanced workdays, ensuring a consistent workload of approximately 8 ± 1 hours on a day. This includes the time slots as mentioned before (from 9:00 till 16:00) in which they can arrive at clients homes. However, this also includes the time from the office to the first client and from the last client to the office if this is not done between 9:00 and 16:00, because nurses can leave the office already at 8:30. Additionally, it also includes the time that is needed for the pre-appointment document review for each client which is usually done the day before by the nurse.

The board of ExpertCare:

- The board requires that all clients have to be served.
- They also want that clients and nurses are satisfied such that they stay with the company.
- A nurse can work at most 12 hours on a day following the collective labor agreement.
- They aim to have a cost efficient solution.

2.5 Performance indicators

This section discusses the performance indicators and how the Injection Team performed in 2023. Performance indicators are quantifiable metrics that indicate the performance of the schedules and routes. It allows us to compare different schedules and routes. The data available over 2023 does however not include substitute nurses, since it only takes the data from the nurses within the Injection Team.

2.5.1 Travel time

The travel time is an important performance indicator since the less time is spent on traveling, the more clients can be helped. The reported travel time is based on the system "Nedap ONS" which estimates the travel time from one client to another. It does this by determining the Euclidean distance between two clients and dividing this distance by a standard travel speed which is set to 50 km/h. However, in reality the actual travel time is higher than reported by Nedap ONS, since they cannot drive in a straight line to the clients and not everywhere a speed of 50 km/h is possible. Therefore, the difference between the time in Nedap ONS and the actual time spend traveled is compared. Over a period of one month the nurses measured their actual travel time and compared this with the time reported in Nedap ONS. Figure 8 shows a histogram plot of the extra travel time measured by the nurses. It is evident that, on average, each trip has an extra travel time of more than 5 minutes. In addition, in Appendix B Figure 27 displays the measurements, the trendline of the measurements and the line if the measurements are equal to the time in Nedap ONS. It can be noticed that there is a slight trend that when routes are longer the extra travel time is likely to be also longer. All the measurements can be found in Table 9 in Appendix B. To the reported data we add 5 minutes to each trip from one customer to another in order to make the travel time more accurate.

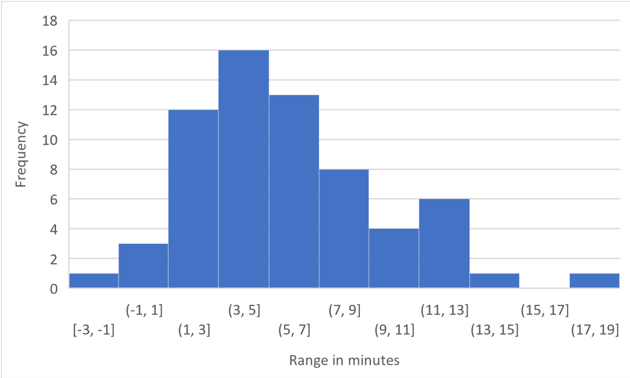

FIGURE 8: Minutes of extra travel time per route

We analyzed the data to determine the average travel time per nurse per day. To the reported data we add 5 minutes to each trip from one customer to another in order to make the travel time more accurate. In Figure 9 the boxplot of the average number of traveling hours per nurse is shown. It is clear that the distribution is symmetric. In Figure 10, the boxplot of the average travel hours per day of the week per nurse is displayed. There are no big differences between the different days of the week. It is determined that on average the travel time per day per nurse is around 2.63 hours.

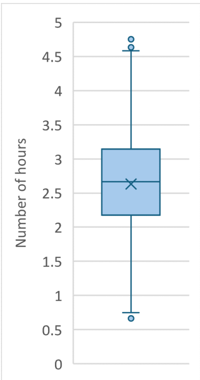

FIGURE 9: Average number of traveling hours per nurse on a day

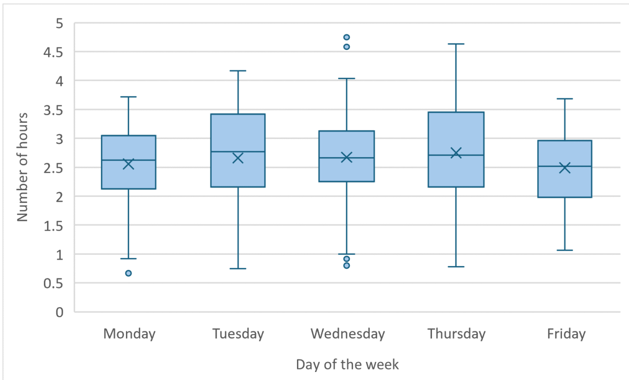

FIGURE 10: Average number of traveling hours per nurse over a week

2.5.2 Continuity of care

Continuity of care is described as having steady and reliable healthcare services over time. It ensures continuity and predictability in the treatment and support for the patient. Maintaining a consistent nurse facilitates continuity of care. Having the same nurse for a client brings several advantages with it. Firstly, it is preferred by the clients since it gives more trust and stability for them. Furthermore, when a nurse attends to the same patient on each visit, they develop familiarity with the client's location, optimizing navigation, parking, and overall efficiency during subsequent visits.

Wirnitzer et al. [38] present several performance metrics for evaluating the continuity of care (coc). One of their objectives is to minimize the relative number of different nurses per patient. This is calculated by dividing the total number of different nurses that have visited a certain client by the total number of visits to that client. Different from Wirnitzer et al. [38] is that we subtract this value from 1 such that a higher value of coc corresponds with a better continuity of care. Also, from the total number of different nurses 1 is substracted, since a client will always be helped by at least one nurse. This shown in Equation 1. Wirnitzer et al. [38] do not mention what a good value for the coc is. A coc of 0 would indicate that a client has had a different nurse for each visit and a coc of 1 would indicate that a client has had the same nurse for each visit. Thus, a higher coc is preferred.

$$\text{coc} = 1 - \frac{\text{Total number of different nurses} - 1}{\text{Total number of visits}} \tag{1}$$

The continuity of care over the year 2023 is calculated by calculating for each client the coc according to Equation 1 and then taking the average and the weighted average on the total number of visits. When taking the average a value of 0.65 is found and when taking the weighted average a value of 0.67 is found. Note that only the clients that have had more than one visit in the year 2023 are taken into account and that substitute nurses are not taken into account as we do not have this data. This would mean that in reality the values found would be lower.

Another way of calculating the coc mentioned in Wirnitzer et al. [38] is the ratio of nurse switches during subsequent appointments. This is done by computing the total number of nurse switches by the total number of visits of a client. This is shown in Equation 2. From the denominator 1 is subtracted since at least two visits are needed to have a nurse switch. A nurse switch means that a different nurse visits a client than in the previous visit of that client. Wirnitzer et al. [38] do not mention what a good value for the ratio of switches is. A ratio of switches of 0 would mean that a client has only had the same nurse, and a ratio of switches of 1 would mean that the client has a different nurse compared to the previous visit every time. Thus, a lower ratio of switches indicates that fewer switches are made between nurses in subsequent switches, so a lower ratio is preferred.

$$\text{Ratio of switches} = \frac{\text{Total number of nurse switches}}{\text{Total number of visits} - 1} \tag{2}$$

The ratio of nurse switches over the year 2023 is calculated by calculating for each client the ratio of switches according to Equation 2 and then taking the average and the weighted average on the total number of visits. When taking the average a value of 0.63 is found and when taking the weighted average a value of 0.57 is found. Similarly, note that only the clients that have had more than one visit in the year 2023 are taken into account and that substitute nurses are not taken into account as we do not have this data. This would mean that in reality the values found would be higher. Additionally, when a nurse leaves the team and a new nurse joins the team, this would lead to more nurse switches for clients, because the nurse leaving cannot help them anymore and they will get a new nurse visiting them.

2.5.3 Balanced working days for the nurses

Another performance indicator is the workload distribution for the nurses. The nurses prefer to have consistency in the length of their working days. They prefer to work around 8 ± 1 hours, which includes nursing and traveling hours.

In Figure 11 the boxplot of the average number of appointments per nurse over a week is displayed for 2023. The average number of appointments on a day per nurse is around 8.02. Moreover, Figure 12 shows the boxplot of the average number of nursing hours per nurse over a week for 2023. The average number of nursing hours on a day per nurse is around 5.92 hours and in general the distribution is symmetric. Note that this excludes the travel time.

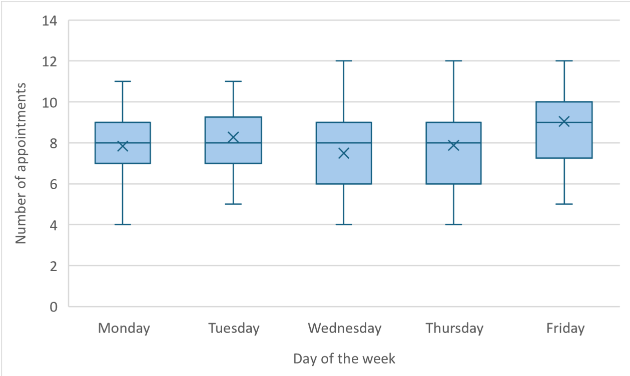

FIGURE 11: Average number of appointments per nurse over a week

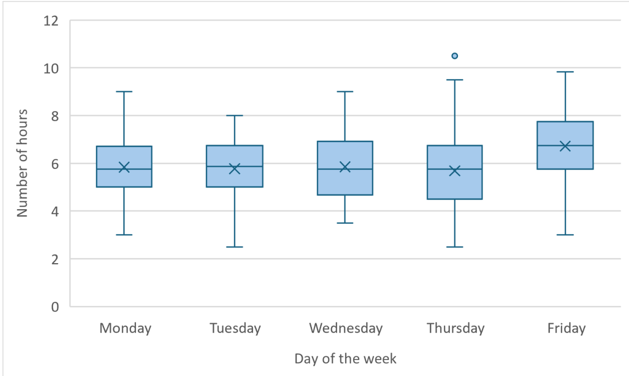

FIGURE 12: Average number of nursing hours per nurse over a week

2.6 Conclusion

This chapter answered the following research question: *How is the current planning of the Injection Team organized?*. First, we present the current planning system employed by the Injection Team. This is done by first introducing the nurses and the clients and then discussing the planning of appointments and the construction of the routes. Additionally, we analyzed the clusters and the frequency of the clients. The analysis revealed the presence of 3 dominant frequencies that occur substantially more relative to the remaining frequencies. Moreover, a stakeholder analysis is done to determine all the wishes and requirements. Lastly, the chapter concluded with several key performance indicators to see how the Injection Team is currently performing. In the next chapter, a literature review is done to determine a suitable solution approach.

3 Literature review

This chapter addresses the second research question: *Which solution approach fits best to the context of ExpertCare?*. Firstly, Section 3.1 describes the home health care routing and scheduling problem (HHCRSP) and examines various aspects of it. Secondly, Section 3.2 discusses a related problem, namely the Vehicle Routing Problem and several variants of it. In addition, Section 3.3 describes a variety of solution approaches applicable for solving this problem. Lastly, Section 3.4 concludes the chapter by proposing a solution approach and addressing the existing literature gap.

3.1 Aspects of the home health care routing and scheduling problem

The planning of the nurses and the design of the routes to visit their clients greatly impact the quality of the service as well as the travel time and associated costs. The problem of scheduling and designing these routes is known in the literature as the HHCRSP. The problem has gained much more attention over the recent years due to the increase in number of home care businesses and the reduction it can have on the costs and client dissatisfaction [11]. The HHCRSP consists of two components, namely the routing and the scheduling of the nurses. These two components can either be solved simultaneously or separately depending on the nature of the problem. The HHCRSP is related to blueprints in terms of planning and designing frameworks that guide the efficient delivery of home health care services. The HHCRSP aims to minimize or maximize a criterion under certain constraints. This criterion and these constraints are determined by the characteristics of the problem and the demands of the stakeholders. In the context of home health care the three main stakeholder groups are the home care organization, the caregivers and the patients [23]. Several literature reviews have been written about the HHCRSP, see for example the papers of Chabouh et al. [8], Cissé et al. [9], Euchi et al. [11] and Goodarzian et al. [14]. Since there are numerous papers about HHCRSP, the focus of this study lies on papers considering the HHCRSP with a long planning horizon. Since not many papers have planning horizons longer than a week, also planning horizons of a week are included. In the following subsections various aspects concerning the HHCRSP are discussed, namely the planning horizon, objectives, constraints, pattern-based and stochasticity. Then in the last subsection, we will discuss the problem of the Injection Team considering these aspects and make a comparison to the literature.

3.1.1 Planning horizon

The planning horizon in the HHCRSP describes the time frame during which the scheduling and routing is made. Its duration varies based on the availability and quality of information and the planning horizon of interest [9]. In the HHCRSP model, the planning horizon commonly spans one day or week. Several papers use a weekly rolling horizon (e.g. Wirnitzer et al. [38]). This means that decisions are made for a fixed period and then updated as time progresses. The time horizon shifts as each period passes and it allows for updates based on newly available information or changed circumstances. There are only a few papers that use time horizons longer than a week. Nickel et al. [29] construct a medium-term master schedule that plans over multiple weeks. Moreover, Hewitt et al. [18] use a planning horizon of 2 to 3 months to demonstrate the advantages of using longer planning horizons over shorter ones. This is done by comparing the longer horizon planning strategy to a planning strategy using a weekly rolling horizon. The paper concludes that medium-term planning leads to a significant decrease in travel time than short-term planning.

3.1.2 Objectives

There are various different objectives in the HHCRSP. We have classified them into three groups, namely cost minimization, patient satisfaction maximization and nurse satisfaction maximization. The most commonly used objective for the HHCRSP is cost minimization [9]. The cost could be described as the traveling costs, the traveling time, the distance or the number of staff. Another objective used in the context of HHCRSP is to maximize the satisfaction of the clients. For the clients this could be the continuity of care and whether their time or nurse preference is fulfilled. Furthermore, objectives of the nurse satisfaction maximization could be workload balancing or minimizing overtime.

These objectives can either be used as a single objective or as multiple objective functions. Bard et al. [2] have as a single objective that optimizes the total travel cost. For multi-objective optimization frequently a weighted objective function is used. Nickel et al. [29] use a weighted objective function for the overtime costs, the travel distance, the number of unscheduled tasks and the patient-nurse loyalty. Moreover, Maya Duque et al. [25] perform a lexicographic optimization. Their most important objective is maximizing the service level and their second objective is minimizing the travel distance. In this case, a large decrease in travel time is allowed for a certain small drop in the service level.

3.1.3 Constraints

There are numerous different constraints used in the literature of HHCRSP. In this section the most frequently used ones are discussed. A very common constraint is the time window constraint. These can either be a hard time window or a soft time window. In the case of a hard time window the service must start and end within the time window. In the case of soft time windows, it is allowed that the time window is violated but this could induce a penalty cost. For example, Lin et al. [24] use soft time windows. Another frequently used constraint is the qualification constraint. It ensures that the nurse with the right qualifications helps a client. It is modelled as a hard constraint, since this condition cannot be violated. Additionally, working time regulation of the nurses are sometimes considered. This includes for example, the work time, the break and holidays of the nurses. Guericke and Suhl [17] consider many working regulations in their model to ensure applicability in practise. In addition, the continuity of care is a common constraint in the context of HHC. This could either ensure in consistency of the client's time slot or the consistency of the nurse over a period of time.

3.1.4 Pattern-based

In a few HHCRSP the flexibility of client's patterns is taken into account. A client's pattern describes the days on which they are visited. For example, a client needs to be visited 2 times a week, then possible patterns are Monday and Wednesday or Tuesday and Thursday. The introduction of multiple possible patterns for clients makes the problem more complex [35]. This is due to the fact that there are more decision variables introduced to the problem. In the papers of Bard et al. [1], Cappanera and Scutellà [7], Maya Duque et al. [25], Shao et al. [35] and Yalçındağ et al. [39] all deal with the assignment of a pattern to a client. Often first, the patient is assigned to a nurse, taking into account qualifications and workload balancing. Then, a pattern is assigned to a patient and lastly the routes are determined. Nickel et al. [29] also work with a type of pattern, but they call it all possible shift combinations for each job. However, they only use it for their weekly optimal plan and not for their medium-term master schedule. In their master schedule they say that each job has only one possible shift combination.

3.1.5 Stochasticity

Stochasticity is another aspect that can be considered in the HHCRSP. Stochasticity in the HHCRSP introduces variability, which can manifest in different facets such as service time, nurse availability, and client demand. Ignoring this uncertainty assumes a static environment where parameters like client numbers and nurse availability remain constant. Conversely, accounting for stochasticity allows for a dynamic model where these parameters fluctuate over time, yielding a more realistic scenario. Notably, research on uncertainty in HHCRSP is limited compared to deterministic approaches. It has recently been identified as a promising future research direction, especially in strategic planning [8]. Fathollahi-Fard et al. [12] take into account the uncertainty in travel and service time as well as patients' satisfaction. Additionally, Hewitt et al. [18] take into account the uncertainty in when and where new patients will request care in the future. They do this by taking into account a set of dummy patients for the future which is based on a point estimate which is the expected number of patients per week. Furthermore, Nikzad et al. [30] take into account the stochasticity of travel and service times.

3.1.6 Classification of ExpertCare's case

In Appendix D an overview of the HHCRSP literature is given in Table 11. We will examine the five aspects discussed above in the context of the Injection Team. Firstly, the planning horizon seen in the literature is often a week and at most 2-3 months. Since numerous clients of the Injection Team have a visitation frequency of 26 weeks, a planning horizon of at least 26 is reasonable. This is a lot longer than seen in the literature. Secondly, the objective of the Injection Team is to minimize the total travel time. This objective is commonly used in the literature. Thirdly, the constraints the Injection Team has are time windows of the arrival time at the clients and a maximum working time for the nurses on a day. The skill linking constraint is not relevant for the Injection Team since all nurses are able to give an injection to all clients. Fourthly, the problem of the Injection Team is a pattern-based problem. This is because the timing of most injections is flexible. This means that the days on which clients are visited is not fixed and the time windows can also be more flexible and larger. Therefore, the papers using a pattern-based approach are the most relevant to our problem. These are Bard et al. [1], Bowers et al. [4], Cappanera et al. [7], Maya Duque et al. [25], Nickel et al. [29], Shao et al. [35] and Yalçındağ et al. [39], as can be seen in Table 11. Lastly, the stochasticity considered in the found literature is the arrival of new clients, uncertainty in time windows, travel time, service time and clients' preferences. Since, we are focusing on the medium-term planning, we have chosen not to consider the changes in appointments and the arrival of new clients. Additionally, the Injection Team works with deterministic values for the service time and the travel time determined in Nedap ONS. Thus, the HHCRSP literature has several relevant papers for the case of ExpertCare, but there are also several differences between the case of ExpertCare and the analyzed literature.

3.2 Variants of Vehicle Routing Problems

Since the scheduling problem of the Injection Team has several differences from the HHCRSP, methods from other industries are also relevant. For example, skill linking is no requirement for the Injection Team and the days on which clients are visited are not fixed. This is because all nurses of the Injection Team are qualified to give injections and with most injections they give the timing is not very strict. There are numerous different sectors in which vehicle routing occurs, such as delivery of packages, waste collection and repairs or installations at people's home. We will discuss different variants of vehicle routing problems (VRPs) which are relevant for the HHCRSP in the following subsections.

3.2.1 Vehicle Routing Problem with time windows

The classical Vehicle Routing Problem (VRP) deals with finding optimal routes to visit each customer exactly once within a single period. The VRP with time windows (VRPTW) generalizes the VRP by adding time windows in which customers have to be visited [5]. This could either be modelled as hard time windows or as soft time windows. However, the VRPTW typically does not take into account a planning horizon longer than a day. Therefore, it does not take into account consistency or working regulations for nurses over longer periods.

3.2.2 Periodic Vehicle Routing Problem

The Periodic Vehicle Routing Problem (PVRP) extends the VRP concept to multiple periods, requiring the determination of visit schedules and routes across a planning horizon to minimize total transportation costs. The PVRP was first introduced in 1974 by Beltrami and Bodin [3] in a paper discussing garbage collection. It aimed to optimize routing plans for visiting a set of customers over multiple planning horizons. Addressing the PVRP involves tackling two interconnected challenge: determining the optimal visiting schedule for each customer within the planning horizon, which corresponds to an assignment problem, and subsequently planning the vehicle routes for each day, which corresponds to the VRP [6]. This dual approach is essential for optimizing the overall transportation efficiency.

3.2.3 Consistent Vehicle Routing Problem

The Consistent Vehicle Routing Problem (ConVRP) was first introduced by Groër et al. [15]. The time horizon for the ConVRP is multiple days. The ConVRP adds consistency to the VRP in two ways. Firstly, it is ensured that each customer is always helped by the same person. Secondly, the visits of a customer should occur roughly at the same time. The ConVRP assumes that the visit schedules of the customers are known in advance. Kovacs et al. [22] generalize the ConVRP by allowing more than one driver to visit a customer. They put a limit on the number of drivers per customer. Moreover, they penalize the variation in arrival time in the objective function instead of having a constraint for this.

3.2.4 Periodic Vehicle Routing Problem with Driver Consistency

The Periodic Vehicle Routing Problem with Driver Consistency is introduced by Rodríguez-Martín et al. [32] and it extends the PVRP by ensuring that each customer is served by the same driver in each visit. In contrast to the ConVRP, the visit schedules of the customers are unknown and need to be determined. The model does not take into account time consistency in visitations. In the paper they formulate the problem as an integer linear program and solve it with an exact brand-and-cut algorithm.

3.3 Solution approaches

In this section different solution approaches are discussed. Konstantakopoulos et al. [21] divided the solution approaches into three classes namely, exact approaches, heuristics and metaheuristics. We will discuss those three approaches in this section.

3.3.1 Exact approach

Exact algorithms were first suggested to solve VRPs because of their ability to generate high-quality solutions. However, for extensive problems involving more than 100 customers, the computational time required for exact algorithms to find the best solution becomes impractical [21]. Consequently, heuristic and metaheuristic algorithms were devised, as they provide a more favorable trade-off between solution quality and computational time [21]. Méndez-Fernández et al. [28] solve their problem by both an exact approach and a heuristic approach, namely simulated annealing. The heuristic approach was discovered to effectively solve a greater number of instances compared to the exact approach. Additionally, it yielded solutions of comparable quality to those generated by the exact approach. The comparison of an exact approach with a heuristic approach is also done in Guericke [16]. The considered heuristics are large neighborhood search, adaptive large neighborhood search and reduced variable neighborhood search. All the heuristics outperformed the exact approach in most cases. In addition, the heuristics are computationally much faster. Since our problem is rather complex we will not go much in the details of the exact approaches. However, they are important for us to have a reference to evaluate and compare the performance of a non-exact method. As a non-exact method does not guarantee to find the optimal solution, it is significant to determine their effectiveness by benchmarking them against an exact solution.

3.3.2 Heuristics

Konstantakopoulos et al. [21] further classify the heuristics into several classes. We will discuss in this section the construction and the 2-phase heuristics.

Construction heuristics

Construction heuristics are relatively simple heuristics and are often able to quickly generate feasible solutions. This makes them very suitable for large-scale complex problems. However, some construction heuristics may perform well for specific problem instances, but poorly for others, limiting their generalizability. Well-known construction heuristics for VRPs are the nearest neighbor heuristic and the Clarke and Wright heuristic [10]. The nearest neighbor simply takes a node and then adds the next node that is the nearest. Additionally, the Clarke and Wright heuristic is based on a cost savings of 2 routes. At first all nodes are initialized as a single route. Then based on the highest cost saving of combining 2 routes, these 2 routes are combined. This is done until no costs cannot be saved anymore or no more routes can be merged. The Clarke and Wright heurisitc has several advantages, such as its simplicity and its ability to produce quality solutions in a reasonable time [37].

2-phase heuristic

2-phase heuristics are problem-solving approaches that involve two distinct phases to find approximate solutions efficiently. In the first phase, initial solutions are generated quickly using heuristic methods. Then, in the second phase, these initial solutions are improved iteratively to enhance their quality. This iterative improvement can involve various techniques like local search or optimization algorithms.

The cluster-first route-second method is a 2-phase heuristic approach that aims to solve the VRP. It does this by first clustering the locations into smaller sets and then the vehicle routes are determined in each set. Clustering is the grouping of similar objects based on similar features. For clustering in transportation it can help to reduce the travel time and to improve the route efficiency. There are numerous criteria on which clustering can be done, such as geographic location, similar visitation frequency, traffic patterns, accessibility or time windows. Recently also a study has been done in the context of HHC. Pahlevani et al. [31] cluster the region of customers into smaller regions. They do this based on the similarity between clients in location and required skill of the nurse and on the dissimilarity in time windows. In this way the problem can be solved in a reasonable computation time.

A related concept is territory planning. It is not a heuristic or algorithm, but more a strategic concept. It is described as the process of dividing an area into smaller regions in order to efficiently allocate resources to the smaller regions. This is often used in the residential waste collection problem [19]. This segmentation into smaller areas facilitates solving sub-problems more efficiently. Moreover, if certain drivers are assigned to each area, it makes sure that the driver knows the route well and also ensures in the case of HHC better continuity of care. Hurkmans et al. [19] state that there are 3 objectives for territory planning, namely minimum overlap, minimum travel time and balanced workload. They use a K-means algorithm to construct an initial solution and perform an Adaptive Large Neighborhood Search (ALNS) algorithm to improve the solution. Zhen et al. [40] perform the territory planning on a longer horizon with customer frequency demand. In their work they formulate a partitioning model and solve the problem with a column generation based algorithm.

3.3.3 Metaheuristics

Metaheuristics conduct a thorough exploration of the solution space, leading to superior solutions compared to classical approaches [26]. However, this enhanced performance comes at the cost of computational efficiency and simplicity. There are numerous metaheuristics in the literature. In this section we will discuss several that have been employed often in the HHCRSP and VRP context.

Greedy Randomized Adaptive Search Procedure (GRASP)

GRASP was first proposed by Feo et al. [13] in 1994. GRASP consists of two phases. First a construction phase and then a local search phase. In the first phase GRASP constructs initial solutions greedily while incorporating randomness to diversify the search space. This randomness helps prevent the algorithm from being trapped in local optima. Then in the second phase GRASP applies a local search to improve the solution quality further. This phase involves iteratively exploring neighboring solutions and moving towards better solutions, often guided by a heuristic. This heuristic used could be for example Simulated Annealing or a Variable Neighborhood Search which will be explained next.

Simulated annealing (SA)
Goodarzian et al. [14] found that SA was one of the three most popular solution approaches used for the HHCRSP in the period from 2019 to 2023. SA, pioneered by Kirkpatrick et al. in 1983 [20], is particularly effective for solving combinatorial optimization challenges with a vast array of potential solutions. It excels at navigating away from local optima, facilitating convergence towards nearly optimal solutions. SA achieves this by permitting the acceptance of suboptimal solutions, known as hill-climbing moves. The overarching strategy involves initially accepting all solutions to diversify exploration, gradually transitioning to solely accepting improvements to intensify the search.

Variable Neighborhood Search (VNS)
VNS was first introduced by Mladenović and Hansen in 1997 [27]. VNS involves systematically changing the neighborhood structure during the search process, allowing for exploration of different solution spaces and escape from local optima. VNS iteratively explores solutions by moving between various neighborhoods, each defining a different notion of proximity to the current solution. A local search is done to improve the current solution within each neighborhood. Then based on an acceptance criteria, either the solution is updated or not.

Large Neighborhood Search (LNS)
LNS was initially proposed by Shaw [36] to solve VRPs in 1998. It focuses on exploring large solution neighborhoods rather than local search methods, which typically explore smaller neighborhoods. LNS has a destroy phase and a repair phase. In the destroy phase they remove a subset of the solutions components and in the repair phase the solution is reconstructed again. Then it is evaluated if the current solution is updated to the new solution. The destroy, repair and evaluation phase are done iteratively until a stopping criteria is met.

Adaptive Large Neighborhood Search (ALNS)
ALNS was first mentioned in 2006 by Ropke and Pisinger [33] as a solution approach to the pickup and delivery problem with time windows. ALNS operates by iteratively exploring a solution space and dynamically adapting its search strategy based on the performance of previously applied operators. Instead of using a single destroy and a single repair operator as in LNS, ALNS uses a set of destroy and repair operators. Each operator is assigned a weight, which determines the probability to be selected in the next iteration. These weights undergo updates after each iteration or set of iterations, prioritizing the selection of effective operators to enhance solution quality.

3.4 Conclusion

This chapter focuses on addressing the research question: *Which solution approach fits best to the context of ExpertCare?*. This question has been answered and the final conclusions and a choice of model are given in this section. In the domain of HHC, while operational planning is broadly analysed, mid-term planning is also gaining more attention. The frequency of care intervals experienced by the clients of the ExpertCare's Injection Team often exceeds those documented in the literature. Literature typically cites weekly, daily, or multiple daily care intervals, whereas the Injection Team's clients require care less frequently, ranging from weekly to semi-annually. Consequently, while literature often advocates for weekly master plans, ExpertCare's approach should be for longer time horizons. Additionally, in the literature the days or moments in which the patients require care is often given whereas for the Injection Team this is more flexible. This flexibility of planning patients can be explored in constructing a master schedule to reduce the travel time. There are a few papers that include this flexibility, but often only for a time horizon of a week or only explored to mainly maximize client's preferences. In our paper we will apply this over longer time horizon and the main goal will be to minimize the travel time.

Thus, our work will investigate the HHCRSP with a significantly lower visitation frequency of clients, an aspect not previously addressed in the literature. Moreover, we will consider a longer time horizon, which has also not been explored in existing research. A two-phase approach is chosen, because Shao et al. [35] showed its success to a similar problem. Additionally, for the second phase we have selected an ALNS as this has been seen in Nickel et al. [29] and Guericke and Suhl [17] as a successful method. In the next chapter the solution design is introduced.

4 Solution design

This chapter answers the third research question: *How should the solution approach be designed?*. Firstly, in Section 4.1 we will give the problem description and in Section 4.2 we will give the underlying assumptions of the model. Furthermore, in Section 4.3 we will describe the solution approach. Lastly, in Section 4.4 we will conclude this chapter.

4.1 Problem description

The model has the task to generate a master schedule with a minimized travel time such that all clients are helped. For this it has to assign a scheme $s \in \mathcal{S}_i$ to each client $i \in \mathcal{C}$. The concept of schemes ($\mathcal{S}_i$) is explained below. The schedule is planned over a time horizon of $\mathcal{W} = \{1, 2, ..., W\}$ weeks, where W is the total number of weeks of the planning horizon. The set $\mathcal{T}$ denotes the set of time slots and $\mathcal{T}_w$ denotes the set of time slots for week w. A time slot refers to a time frame in which nurses can arrive at clients' homes, so there could be multiple time slots on a day. For every time slot t the routing also needs to be determined for the clients assigned to this time slot. The variable $y_{t,i,j}$ is equal to 1 if in time slot t a nurse goes from client i to client j and 0 otherwise. The nurse always starts and ends at the office (denoted by O). The travel time between location i and location j is denoted as $tt_{i,j}$. The travel time is symmetric, because it is based on the euclidean distance between two locations with a constant added to it. This means that the travel time from location i to location j is the same as going from location j to location i or simply $tt_{i,j} = tt_{j,i}$. The clients are defined as $\mathcal{C} = \{1, 2, ...C\}$, where C is the total number of clients. A client i has a visitation frequency f_i. The visitation frequency denotes the number of weeks that must be between subsequent appointments. Additionally, each client has a week number parameter b_i which indicates the week in the time horizon in which their first appointment is.

Generating schemes

Thus, a scheme needs to be assigned to each client. A scheme refers to a pattern which we discussed in Section 3.1.4. Recall that a client's pattern describes the days on which they are visited. For the remainder of this paper, we will refer to patterns as schemes. This description of schemes is inspired by Bard et al. [1] and Shao et al. [35]. Each client i has a set of schemes $\mathcal{S}_i$, which consists of all the possible schemes client i can have. This set of schemes depends on the frequency f_i of the client, the week number of the first visit in the time horizon b_i, the set of time slots in each week $\mathcal{T}_w$ and the planning horizon. When assuming the first appointment must occur in week b_i, then the scheme set of each client consists of $|\mathcal{T}_w|$ schemes. For example, we have a planning horizon of 4 weeks and $\mathcal{T}_1 = \{\text{Mon}_1, \text{Tue}_1\}$, $\mathcal{T}_2 = \{\text{Mon}_2, \text{Tue}_2\}$, $\mathcal{T}_3 = \{\text{Mon}_3, \text{Tue}_3\}$ and $\mathcal{T}_4 = \{\text{Mon}_4, \text{Tue}_4\}$, which means that each week consists of only two time slots. This is also displayed in Figure 13 for clarification. The subscript for each day indicates the week number. Then if a client has a frequency of 1 visit per 2 weeks and the week of their first appointment in the planning horizon is 1, then $\mathcal{S}_i = \{\{\text{Mon}_1, \text{Mon}_3\}, \{\text{Tue}_1, \text{Tue}_3\}\}$. The first scheme in this set, $\{\text{Mon}_1, \text{Mon}_3\}$, means that the client will be helped in week 1 on Monday and in week 3 on Monday. Additionally, if a client i has a frequency of 1 visit every week then $\mathcal{S}_i = \{\{\text{Mon}_1, \text{Mon}_2, \text{Mon}_3, \text{Mon}_4\}, \{\text{Tue}_1, \text{Tue}_2, \text{Tue}_3, \text{Tue}_4\}\}$. Note that when a client has a weekly visitation frequency, then b_i is always 1, which is the first week of the planning horizon. The generation of schemes can easily be adapted if clients can for example also be helped after their starting week. This allows for more schemes. In our example, the client with a frequency of 1 visit per 2 weeks, then $\mathcal{S}_i$ becomes $\{\{\text{Mon}_1, \text{Mon}_3\}, \{\text{Tue}_1, \text{Tue}_3\}\{\text{Mon}_2, \text{Mon}_4\}, \{\text{Tue}_2, \text{Tue}_4\}\}$.

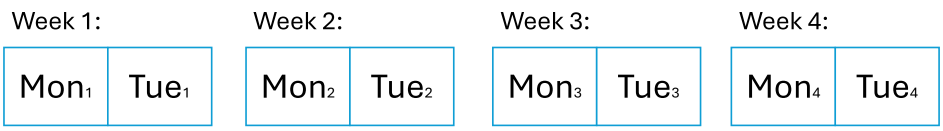

Figure 13: Time slots visualization

The model formulation

In Table 3 the descriptions of all the sets, parameters and variables are shown. Furthermore, the model formulation is given below. Next, the objective function and constraints are explained.

TABLE 3: Symbols of sets, parameters and variables and their meaning

Sets	
$\mathcal{C}$	Set of clients
$\mathcal{S}_i$	Set of schemes for client i
$\mathcal{T}$	Set of time slots
$\mathcal{T}_w$	Set of time slots for week w
$\mathcal{W}$	Set of weeks
Parameters	
$tt_{i,j}$	Travel time between location i and location j
$a_{t,s}$	Binary variable that indicates whether time slot t is in scheme s
b_i	Week of first visit to client i in the planning horizon
c_i	Earliest time at which the appointment of client i can start
d_i	Latest time at which the appointment of client i can start
f_i	Visitation frequency of client i
h	Duration of an appointment in minutes
D_t	Maximal duration of a time slot t in minutes
M	Big number
O	Office
Variables	
$x_{i,s}$	Binary variable that indicates whether client i is served according to scheme s
$y_{t,i,j}$	Binary variable that indicates whether in slot t a nurse goes from location i to location j
$z_{t,i}$	Positive integer that indicates the number of clients visited after visiting location i in time slot t
$z'_{t,i}$	Non-negative value that indicates the time at which the nurse enters location i in time slot t

$$
\begin{aligned}
\min \quad & \sum_{t\in\mathcal{T}}\sum_{i,j\in\mathcal{C}} tt_{i,j}y_{t,i,j} && && (3)\\
\text{s.t.} \quad & \sum_{s\in\mathcal{S}_i} x_{i,s} = 1 && \forall i \in \mathcal{C}, && (3a)\\
& \sum_{i\in\mathcal{C}\cup O} y_{t,i,j} = a_{t,s}x_{j,s} && \forall j \in \mathcal{C}, \forall s \in \mathcal{S}_j, \forall t \in \mathcal{T}, && (3b)\\
& (\sum_{i,j\in\mathcal{C}\cup O} (tt_{i,j}+h)y_{t,i,j}) - h \le D_t && \forall t \in \mathcal{T}, && (3c)\\
& \sum_{i\in\mathcal{C}\cup O} y_{t,i,j} = \sum_{i\in\mathcal{C}\cup O} y_{t,j,i} && \forall t \in \mathcal{T}, \forall j \in \mathcal{C}\cup O, && (3d)\\
& z_{t,O} = 0 && \forall t \in \mathcal{T}, && (3e)\\
& z_{t,j} \ge z_{t,i} + 1 - M(1-y_{t,i,j}) && \forall i \in \mathcal{C}\cup O, \forall j \in \mathcal{C}, \forall t \in \mathcal{T}, && (3f)\\
& z_{t,i} \le \sum_{i\in\mathcal{C}}\sum_{s\in\mathcal{S}_i} x_{i,s}a_{t,s} && \forall i \in \mathcal{C}, \forall t \in \mathcal{T}, && (3g)\\
& x_{i,s} \in \{0,1\} && \forall i \in \mathcal{C}, \forall s \in \mathcal{S}_i, && (3h)\\
& y_{t,i,j} \in \{0,1\} && \forall t \in \mathcal{T}, \forall i,j \in \mathcal{C}\cup O, && (3i)\\
& z_{t,i} \in \mathbb{Z}^+ && \forall t \in \mathcal{T}, \forall i \in \mathcal{C}\cup O && (3j)
\end{aligned}
$$

The objective in (3) is to minimize the total travel time over the time horizon. The first constraints (3a) ensure that each client gets assigned one scheme. Constraints (3b) make sure that a client can only be visited in a time slot if the time slot is in their assigned scheme. Constraints (3c) guarantee that the maximum duration of a time slot is not violated. The h is subtracted since the office does not have a service time. Constraints (3d) are the flow conservation constraints. Constraints (3e) assure that the number of visited clients at the office is 0 for all time slots. Constraints (3f) make certain that the number of visited clients will be one higher after visiting a client. This is needed to ensure there are no subtours, meaning that each time slot has one route that starts and ends at the office. Constraints (3g) enforce an upper bound on the number of clients visited in a time slot to be the total number of clients scheduled in that time slot. Constraints (3h) and (3i) are the binary constraints and constraints (3j) ensure that the number of visited clients are positive integers.

Model with time windows

As discussed in Section 2.4 the clients of the Injection Team have certain time windows to be served. There are two time windows in which an appointment can start, either between 9:00 and 13:00 or between 12:00 and 16:00. To also take into account the time windows in which they must be served, the model can be adapted easily, shown below:

$\min \sum_{t\in\mathcal{T}}\sum_{i,j\in\mathcal{C}} tt_{i,j}y_{t,i,j}$

s.t. Constraints (3a), (3b), (3c), (3d), (3e), (3h) and (3i)

$$
\begin{aligned}
& z'_{t,j} \ge (z'_{t,j} + tt_{i,j})y_{t,i,j} + z'_{t,j}(1-y_{t,i,j}) && \forall i \in \mathcal{C}\cup O, \forall j \in \mathcal{C}, \forall t \in \mathcal{T} && (4)\\
& z'_{t,i}a_{t,s} \ge c_i a_{t,s}x_{i,s} && \forall i \in \mathcal{C}, \forall s \in \mathcal{S}_i, \forall t \in \mathcal{T} && (5)\\
& z'_{t,i}a_{t,s} \le d_i a_{t,s}x_{i,s} && \forall i \in \mathcal{C}, \forall s \in \mathcal{S}_i, \forall t \in \mathcal{T} && (6)\\
& z'_{t,i} \ge 0 && \forall t \in \mathcal{T}, \forall i \in \mathcal{C}\cup O && (7)
\end{aligned}
$$

The variable $z'_{t,i}$ is introduced and it represents the time at which the nurse enters location i in time slot t. This is defined for all locations, but for the office it only represents the time at the beginning which is set to 0. Constraints (4) update the time at which the following clients can be entered. Moreover, constraints (5) and (6) are the time window constraints. Lastly, constraints (7) ensure that z is a positive continuous variable. Constraints (3f), (3g) and (3j) can be removed from the previous model, as the constraints (4), (5), (6) and (7) inherently eliminate subtours.

4.2 Assumptions

Several assumptions are made for the model and these are listed in this section.

- It is assumed that a client will be served in the same time slot every week. For example, if a client is helped on Monday morning, it will be helped in all the other appointments also on Monday morning. This is done for consistency for the planners and the clients. At the moment this is also done in practise.
- Another assumption that has been made is that each client has a given week in which their first appointment in the time horizon takes place. This can be moved forward or backward in time by a few days depending on their frequency. Then based on their frequency of care, the weeks in which their following appointments will take place are determined. This assumption is made, since clients who need an appointment every half year or quarter cannot randomly plan their appointment in any week.
- In addition, to determine the length of a planning horizon to work on a rolling horizon basis, the least common multiple of all the frequencies is taken. Since the least common multiple of 1, 2, 3, 4, 6, 8, 12 and 26 is 312, we would ideally need to extend our planning horizon to 312 weeks. Unfortunately, this extended timeframe is impractical for our model as it would significantly increase computation time. Therefore, a planning horizon of 26 weeks is taken. However, for this to work on a rolling horizon basis, some frequencies have to be rounded. This has the implication on the results that the model may provide solutions that are optimal for the 26-week horizon but not for the actual longer-term needs, leading to suboptimal performance over the full period.
- Furthermore, it is assumed that each week has the same number of time slots.
- Moreover, the stochasticity in service time is not taken into account. In their current planning process deterministic values for the service time is also adopted. The appointment service time is half an hour.
- Additionally, the travel times are taken the same as in their current planning process which is calculated in Nedap ONS. However, we add 5 minutes to each trip from one customer to another. This is based on the data demonstrated in Section 2.5.1.

4.3 Solution approach

In Section 3 several works were discussed with similar problems and most of them use a heuristic approach to solve the problem, due to the complexity of the problem and the computation time needed. Additionally, the time horizon of our problem is much larger compared to the literature which makes a heuristic approach suitable. This paper adopts an Adaptive Large Neighborhood Search (ALNS) as a solution approach which is shown in Figure 14. In the first phase, an initial solution is constructed in a greedy manner. In the second phase, we implement an ALNS to improve the solution. In our model we have two possibilities for the time slots, namely a time slot as an entire day and a time slot as half a day. The latter represents the current morning slot and an afternoon slot as currently employed by the Injection Team. The latter is more complex than the former, as the time windows of clients are smaller.

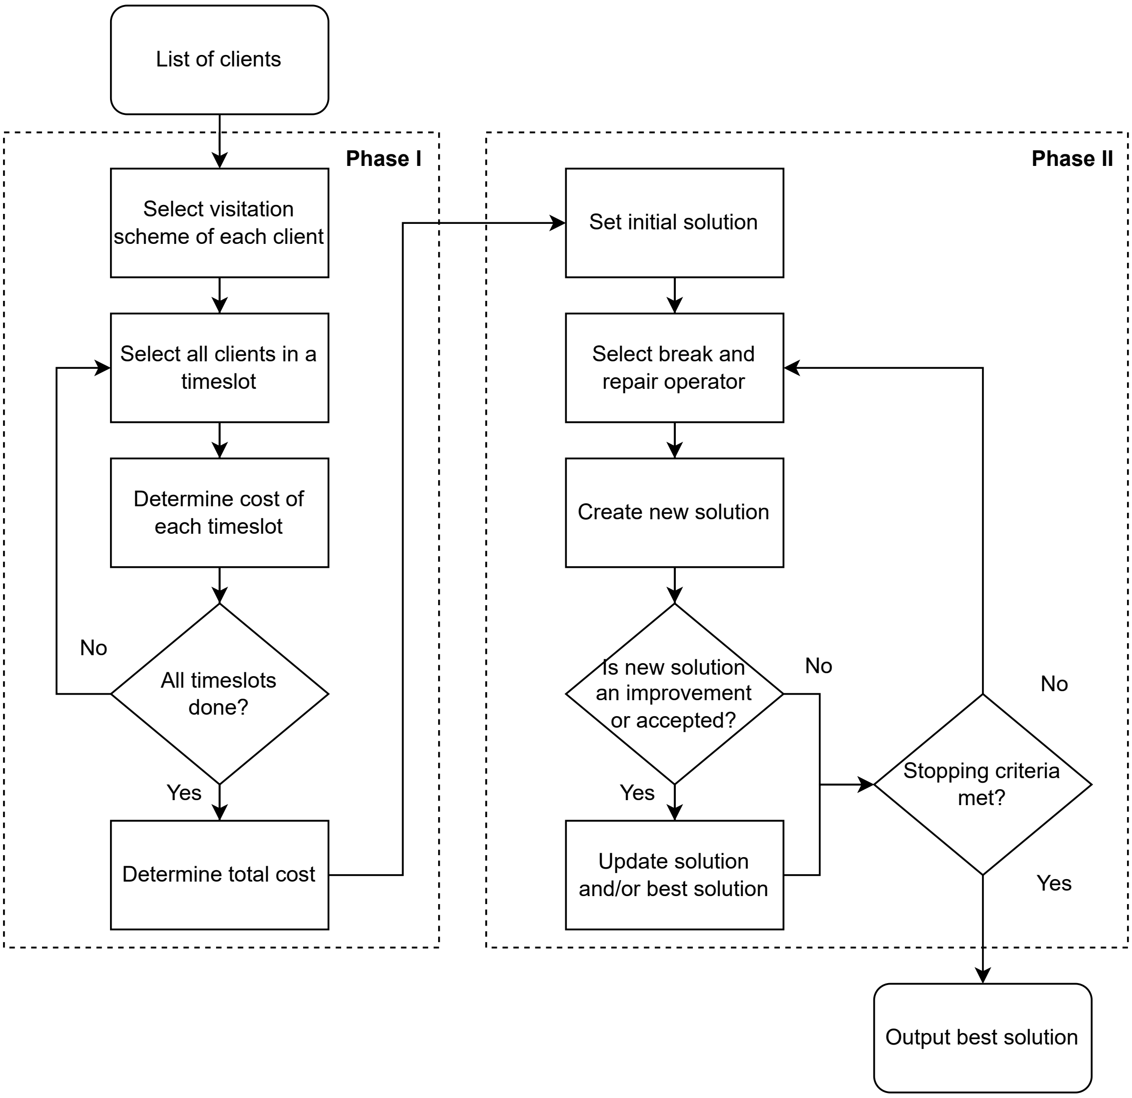

FIGURE 14: Solution approach

4.3.1 Phase I - Initial solution

At first an initial solution needs to be constructed such that it is known on which day which clients will be visited. For this, we need to assign a scheme for each client. We have constructed two greedy heuristics that generate an initial solution. The greedy heuristic will iteratively assign a feasible scheme from the set $\mathcal{S}_i$ to each client i. Only feasible schemes of clients are evaluated, which are the schemes that satisfy the time windows of customers and the working hours of the nurses. The clients are sorted based on their frequency. The ones which require the most visits during the time horizon get assigned a scheme first. This is done because they pose the greatest obstacles for developing feasible solutions [6]. We have designed two greedy heuristics that assign schemes in the following way:

- Assignment can be based on the number of already scheduled appointments on certain days. If there are already many appointments assigned to a day, it might be less preferable to assign a client here. So for each time slot in a scheme it is evaluated how many other clients are already scheduled. Then a scheme among the least busy schemes is randomly selected and assigned to the client. When running this greedy heuristic with the given input of the Injection Team, it is not able to generate a feasible solution. This can be explained by the possibly inefficient schedules created as it does not take into account the travel time. The pseudocode can be seen in Algorithm 3a in Appendix E.

- Another way the assignment can be done is based on the distance. For each feasible scheme of a client, each time slot is evaluated. A flowchart for the decision making is given in Figure 15. Firstly, it is considered whether there are other clients in the time slot, as seen in the first question in the flowchart in Figure 15. If this is the case then the average distance to the other clients in that time slot is added to the score. If there is no one in the time slot scheduled yet then it is considered if there are clients in the time slot on the same day of the same route, as seen in the second question in the flowchart in Figure 15. If this is the case the minimum distance to one of the clients in this other time slot is added to the score. However, if this is not the case we will go to the third question in the flowchart in Figure 15. If the average distance to their 10 or 5 closest neighbors is determined. If this is below $\frac{1}{3}$ then a score of 0.2 is added to the score and otherwise 0.5 is added to the score. The score of each time slot is then summed for each scheme and normalized by their number of appointments. Then a scheme with the lowest score is selected. If there are multiple with the lowest score then one is randomly assigned from those. The pseudocode for this is shown in Algorithm 1a and Algorithm 1b for time slots of days and time slots of half a days, respectively. When a time slot lasts a whole day, then the second question in the flowchart of Figure 15 is skipped. This greedy heuristic is able to generate feasible solutions and will therefore be used for our ALNS. Below Figure 15 we will give an explanation of the choices made in the flowchart.

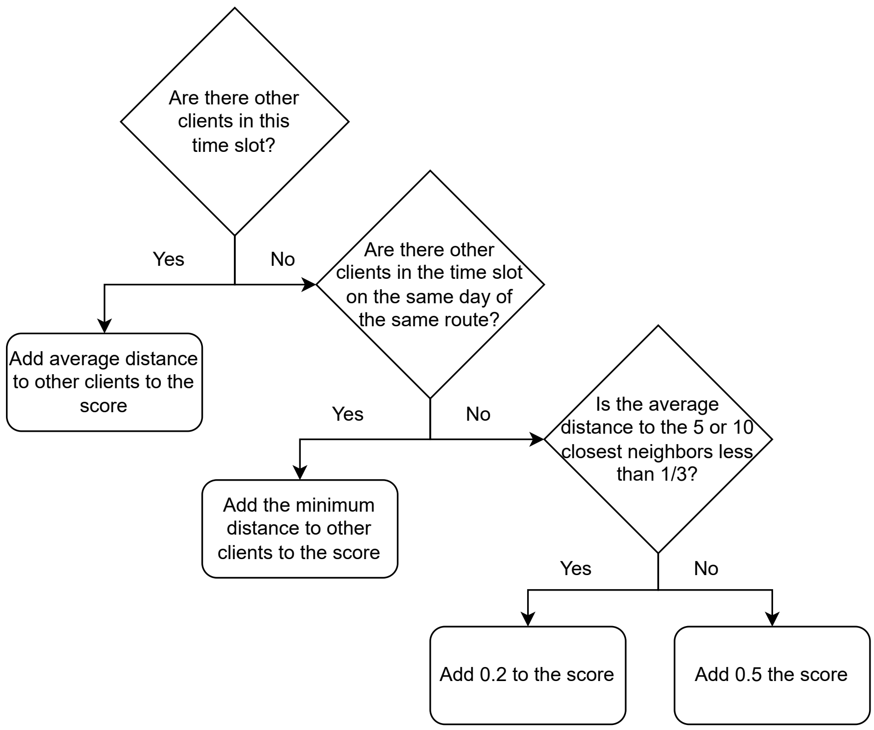

FIGURE 15: Flowchart of the decision making in the greedy heuristic

Firstly, when there are other clients in the time slot, the average distance to these clients is added to the score. This is done to ensure that the client is assigned to a time slot where it is close to the other clients. Secondly, the minimum distance to other clients is added to the score when there are clients in the time slot on the same day of the same route. This other time slot would then either be the morning or the afternoon route done by the same nurse. In this case, we take the minimum distance to the other clients, because the time slot considered should be close to the other time slot. Lastly, a constant is added to the score. If we do not add a constant to the score, then at first all the empty time slots will be filled. However, it could be that there are other time slots that have clients that are close to them. Thus, we add this constant in order to not assign clients to empty slots when there are time slots with clients close to them. The constant is either 0.2 or 0.5. This depends on the average distance to their 10 or 5 closest neighbors. When the time slot lasts a day 10 neighbors are chosen and when a time slot lasts half a day 5 neighbors are chosen. The values 5 and 10 are chosen, because this is the rounded average number of clients in a time slot. Over all clients the average distance to their 5 or 10 closest neighbors is equal to 0.33 hours and 0.31 hours, respectively. Therefore, if the average distance to their closest neighbors is below $\frac{1}{3}$ then this client has many close neighbors meaning that it can be more easily assigned to an empty slot and therefore gets a lower score. When it is above $\frac{1}{3}$ then this client does not have many close neighbors which means that is should not be more easily assigned to an empty slot. By adding a higher score, the client is more likely to be assigned to a time slot with a client close to it.

Algorithm 1a Greedy heuristic: minimize distance (1 time slot is a day)

```
Sort clients by frequency
for each client i do
    for each feasible scheme s in S_i do
        Score = 0
        for each time slot t in s do
            if no client in t then
                Score = Score + value based on their closeness to their neighbors
            else
                Score = Score + average distance from i to all other client already in t
            end if
        end for
        Normalize Score with number of appointments in s
        Save Score to scheme s
    end for
    Select a scheme with the lowest Score
end for
```

Algorithm 1b Greedy heuristic: minimize distance (1 time slot is half a day)

```
Sort clients by frequency
for each client i do
    for each feasible scheme s in S_i do
        Score = 0
        for each time slot t in s do
            if no client in t and there are clients in the other time slot on the same day then
                Score = Score + minimal distance from i to client in the other time slot
            else if no client in t and no clients in the other time slot on the same day then
                Score = Score + value based on their closeness to their neighbors
            else
                Score = Score + average distance from i to all other client already in t
            end if
        end for
        Normalize Score with number of appointments in s
        Save Score to scheme s
    end for
    Select a scheme with the lowest Score
end for
```

4.3.2 Phase II - Adaptive Large Neighborhood Search

Neighborhood

Neighborhood solutions to the current solution in the ALNS are explored by generating destroy and repair operators. The performance of each operator is tracked and operators are adaptively selected based on their performance. The destroy operators select a number of clients, based on the degree of destruction (dod), and removes their assigned scheme. Then the repair operator will go through the unassigned clients and assign a new scheme to them. In case the repair operator is not able to assign a scheme to a client, then the neighborhood solution will be set equal to the current solution.

The following destroy operators are used for our ALNS:

1. Random destroy: randomly selects a number of clients whose assigned scheme is removed.
2. Worst destroy: removes the schedule of a number of clients that are the "farthest". For this each assigned scheme of a client is evaluated. For each time slot the average distance to the other clients already in that time slot is summed. In case, the time slot is empty and time slots last half a day, then the other time slot on that day is considered. The minimum distance from the considered client to any client in the other time slot is taken. Finally, the number of clients with the highest value are removed, normalized by their number of appointments in their scheme.
3. Frequency destroy: randomly selects a frequency and then randomly removes schemes of clients with that frequency. When the degree of destruction is higher than the number of clients with that frequency, then randomly clients are selected to be removed.
4. Every time slot destroy: randomly selects a time slot and then randomly removes one client. Time slots cannot be selected again unless all time slots have been selected already.
5. One cluster destroy: randomly selects a cluster and then randomly destroys clients within that cluster. When the degree of destruction is higher than the number of clients in that cluster, then another cluster will be randomly selected.
6. Worst cluster destroy: randomly selects a cluster and then removes the clients with worst destroy. This is done up to half of the clients in that cluster (round down). When the degree of destruction is higher than half of the clients in the cluster, a new cluster is randomly selected.
7. Area destroy: generates a point on the map based on the density of the clients. This means that a point near a big city is more likely to be generated than in a rural area. Then clients are

randomly removed if they are within 0.1 hour within this point. If the degree of destruction is higher than the number of clients within this area, then the radius is increased by 0.01 hour until enough clients are present in the area.

8. Worst normalized destroy: is similar to worst destroy, but now the values of each client are normalized by the average of their distances to their 5 or 10 closest neighbors depending on the length of a time slot. This is done to decrease the preference of removing the clients that have neighbors living far away from them.

The following repair operators are used for our ALNS:

1. Random repair: loops over the clients based on their frequency and randomly assigns a feasible scheme to a client that has no scheme.
2. Greedy repair sorted: loops over the clients based on their frequency and assigns a feasible scheme that is closest to a client. For this each feasible scheme for a client is evaluated. For each time slot the average distance to another client already in that time slot is summed. In case, the time slot is empty and time slots last half a day, then the other time slot on that day is considered. The minimum distance from the considered client to any client in the other time slot is taken, normalized by their number of appointments in their scheme. The scheme with the lowest value is assigned.
3. Greedy repair unsorted: does the same as greedy repair sorted, but now it loops over the clients randomly and not based on their frequency.

Evaluation step

In this step the newly constructed solution from the neighborhood is evaluated, such that we are able to see if the new solution is an improvement. This evaluation needs to determine the quality over the whole time horizon of the routing. Often in the literature approach the routing for a day is constructed based on the time constraints of the clients. These are often smaller than for the Injection Team clients and therefore determine the broad outline of their routes. In our case we have more flexibility in the order of the clients and the time horizon is much longer than a week, which means that the routing needs to be done frequently for one evaluation. Therefore, we can employ a heuristic approach that is computationally fast, but still gives reasonable solutions. For this the Clarke and Wright algorithm is a suitable option. This algorithm is also known as the savings algorithm [10]. We have chosen for the Clarke and Wright heurisitc in our evaluation step due to the several advantages it has, such as its simplicity and its ability to produce quality solutions in a reasonable time (as mentioned in Section 3.3.2).

Thus, over the whole time horizon the routing is done for each day and the travel time is summed resulting in the total travel time. Some schemes for clients have less appointments than other schemes, so the final output could be biased towards schemes with less appointments as this client is less visited, which makes it likely to have less travel time. Thus, to avoid this bias in the final output, we have introduced a penalty for having less appointments than the initial solution. For each appointment the final output has less than the initial number of appointments a value of 1.53 hours is added to the total travel time. This 1.53 hours is chosen, since it is the maximum distance between any two clients.

Additionally, after every L iterations several parameters need to be updated, where L is the Markov chain length. The weights (w_j) of the operators need to be updated based on the roulette wheel parameter (RW). This is done according to the following formula: $w_j = w_j(1-RW) + RW(\frac{\pi_j}{\text{usage}_j})$ for j being all destroy operators and all repair operators, π_j is the success of operator j in the previous set of iterations and usage_j is the usage of operator j in the previous set of iterations. In this way each operator is assigned a weight that depends on its past behavior. Below all the steps of the ALNS are given for more clarity. Additionally, in Step 1 several parameters need to be initialized and this will be explained in detail in Section 5.2. For more clarity r represents repair operators, d represents destroy operators and j represents repair and destroy operators.

ALNS: steps taken
Step 1: Initialize data
Initialize the current solution x_0 and the best solution so far x_{best} and initialize the starting temperature T_{start}, the stopping temperature T_{stop}, the current temperature T, the decrease factor α, the Markov chain length L, the degree of destruction dod, the roulette wheel parameter RW, the success values σ_1, σ_2, σ_3 and σ_4 and the counter i. Additionally, set the starting weight w_j and the success rate π_j of each operator j to 1 and 0, respectively.

Step 2: Generate a neighbor solution
A neighbor solution is created by performing a destroy and repair operator on the current solution. Firstly, select a destroy operator d where each destroy operator has the probability of being selected of $\frac{w_d}{\sum_{d=1}^{8} w_d}$. Additionally, select a repair operator r where each repair operator has the probability of being selected of $\frac{w_r}{\sum_{r=1}^{3} w_r}$. Perform the destroy and repair operator to find the new solution x.

Step 3: Evaluate the new solution
There are 4 options:

- If $f(x) < f(x_{best})$, then update the best solution so far. Then x_{best} and x_0 becomes x. Also, add σ_1 to π_j, where j is the used destroy and repair operator.
- If $f(x) < f(x_0)$ then accept the new solution. Then x_0 becomes x. Also, add σ_2 to π_j, where j is the used destroy and repair operator.

When the new solution performs worse than the current solution, we generate a random number between 0 and 1 (z).

- If $f(x) > f(x_0)$ and $z \leq e^{\frac{x_0 - x}{T}}$ (commonly used acceptance criteria [34]), then the new solution is not better than the current solution, but still accepted by some probability. Then x_0 becomes x. Also, add σ_3 to π_j, where j is the used destroy and repair operator.
- If $f(x) > f(x_0)$ and $z \geq e^{\frac{x_0 - x}{T}}$, then the new solution is not accepted. Also, add σ_4 to π_j, where j is the used destroy and repair operator.

Step 4: Update counter
If $i < L$, then $i = i + 1$ and go to **Step 2**, otherwise go to **Step 5**.

Step 5: Update weights and temperature
Update the weights, $w_j = w_j(1 - RW) + RW(\frac{\pi_j}{\text{usage}_j})$ for j being all destroy operators and all repair operators, where usage_j is the usage of operator j and π_j is the success of operator j in the previous set of iterations. If $T > T_{stop}$, then $T = T \cdot \alpha$ and go to **Step 2**. Otherwise, the algorithm is terminated and we return x_{best}.

Pseudocode ALNS
The pseudocode of the ALNS is given below.

Algorithm 2a Adaptive Large Neighborhood Search (ALNS)

```
Require: Initial solution, StartTemp, StopTemp, α, Mlength
Initialize Solution and CurrentBest
Initialize parameters and termination criteria
while Temp > StopTemp do
    for i in range(Mlength) do
        Select destroy operator d based on probabilities and update usage of d
        Select repair operator r based on probabilities and update usage of r
        NewSolution := r(d(Solution)
        if NewSolution < Solution then
            if NewSolution < CurrentBest then
                CurrentBest := NewSolution
            end if
            Solution := NewSolution
        end if
        if AcceptStrategy = TRUE then
            Solution := NewSolution
        end if
        Update weights based on solution quality and operator selection
    end for
    Update weights based on usage and success of operators
    Update probabilities based on weights
    Decrease Temp
end while

Return: CurrentBest
```

4.4 Conclusion

This section concludes the solution design in which we focused on the research question: *How should the solution approach be designed?*. First, we discussed the problem description and introduced a model for this. Then two greedy heuristics are discussed which construct initial solutions. Additionally, the ALNS framework is introduced which will further improve the initial solutions. In the following chapter the parameter tuning of our approach will be done as well as the execution of several experiments. The quality of the solutions will be evaluated in Section 6.1.

5 Experiments

This chapter discusses the experimental evaluation of our model in order to answer the research question: *How can we tune the solution approach to ensure high-quality outcomes?*. In Section 5.1 the problem specific data input is discussed. By this we mean data input that are specific to the problem of ExpertCare. This includes among others the starting weeks of clients. In Section 5.2 several parameters required as input for our ALNS are discussed. Then, in Section 5.3 the different scenarios for our experiments are given. Lastly, Section 5.4 concludes this chapter. All experiments are performed on a laptop with a central processing unit of AMD Ryzen 5 PRO 7530U with a base speed of 2.0 GHz and a random-access memory of 32 GB with a transfer rate of 3200 MT/s.

5.1 Data input from the Injection Team

5.1.1 Client data

In the context analysis in Section 2 the clients of the Injection Team were already thoroughly analyzed. The number of clients in the Injection Team reported are 879. The clients have several parameters, namely their visitation frequency, their home location and their starting week. Firstly, in Table 2 in Section 2.3 the frequencies of the clients are shown. Secondly, in Figure 6 a map with the home location of the clients are shown and their according frequency. Thirdly, the starting weeks are chosen based on the week in which clients had their last appointment. However, since some frequencies do not divide 26 to an integer, the number of appointments a client has differs per starting week. For example, a client with an 8 weekly frequency and starting week 1 or 2, will give 4 appointments whereas if the starting week is higher than 1 or 2, it will result in 3 appointments in the time horizon. Thus, the starting week of each clients gives 26 scenarios of starting weeks of clients. We have chosen a random one and kept it the same during all experiments for consistency.

The clients do not actually have to be scheduled necessarily in their starting week, as they can also be moved to the week before or the week after the starting week depending on their frequency. The amount of days that a client can move is obtained from a member of the Injection Team. Clients that require visits every 2 weeks cannot be moved. Additionally, clients with a 3, 4, 6, and 8 weekly frequency can be moved up to 3 days into the week after their starting week. The clients with a quarterly and half yearly frequency can be moved up to 5 days into the week after their starting week. Furthermore, the clients with the half yearly frequency can be moved up to 5 days in the week before their starting week. The weekly frequency is trivial since, it is already scheduled in each week.

5.1.2 Other input data

In Table 4 an overview is given of the other input data used for our model, including the number of clients. Our model will utilize a 26-week planning horizon. Typically, two nurses are scheduled daily, unless specified otherwise. Additionally, two time slots on a day is used, unless indicated otherwise. Two time slots on a day mean a morning slot and an afternoon slot. The nurse can arrive at the clients in the morning slot between 9:00 and 13:00 and in the afternoon between 12:00 and 16:00. Note that the nurse can already leave the office at 8:30 as mentioned in Section 2. Lastly, the total number of time slots depends on the number of time slots on a day and the number of nurses working. The number of time slots on a day can for our model either be one or two. Then using the following formula the number of total time slots can be determined:

$$\text{Number of total time slots} = \text{number of time slots on a day} \cdot 5 \cdot 26 \cdot \text{number of nurses working},$$

where 5 is the number of working days in a week and 26 denotes the number of weeks of the planning horizon. Thus, this leads for our problem to a total of 520 time slots over the planning horizon.

Table 4: Overview of input data

Number of clients	879
Planning horizon in weeks	26
Number of daily nurses	2
Number of time slots on a day	2
Total number of time slots	520

5.2 Data input for parameters of adaptive large neighborhood search

This Section considers the data input parameters needed for our ALNS. The parameters that are needed for the ALNS were mentioned in Step 1 in Section 4.3.2. Firstly, we discuss the initial solution and the best solution so far in Section 5.2.1. Moreover, we have chosen to tune a number of parameters, namely T_{start}, T_{stop}, α, L, the degree of destruction (dod) and the roulette wheel parameter (RW). This is discussed in Section 5.2.2. Properly tuning parameters is important to help the algorithm find higher-quality solutions and to let it operate more efficiently, balancing exploration and exploitation to avoid unnecessary computations. However, we have chosen to not tune some parameters, namely σ_1, σ_2, σ_3 and σ_4, because tuning all parameters is resource-intensive and the other parameters are deemed more influential. In subsection 5.2.3 the parameters we did not tune are discussed.

5.2.1 Initial solution and best solution so far

The initial solution could either be constructed by the greedy heuristic that aims to distribute the workload or the other greedy heuristic that aims to minimize the overall distance. The first one is with our given input, not able to generate an initial solution. This can be explained by the fact that the schedule is rather full and by not taking into account travel time, inefficient schedules are created which require too much time. However, the second greedy heuristic is able to construct feasible initial solutions. Thus, the initial solutions generated are constructed by the greedy heuristic that aims to minimize the overall distance. In the ALNS the current solution and the best solution so far is initialized to this solution. The best solution so far is set to the current solution since this is the best solution we have found so far.

5.2.2 Fine tuning of parameters

To determine suitable values for the parameters of the ALNS, we perform several experiments. The parameters that need to be tuned are T_{start}, T_{stop}, α, L, the degree of destruction (dod) and the roulette wheel parameter (RW). Table 5 shows the result of the conducted experiments. The first column is the experiment number and the second column shows the total number of iterations performed during the experiment, which depends on T_{start}, T_{stop}, α and L. The run time of each experiment is shown in seconds and the column called outp. stands for the total travel time of the final solution including the penalty as discussed in Section 4.3.2. The column called Impr. shows the improvement made from the initial solution by the ALNS. The last column gives the standard deviation of the output in percentages. Each experiment is done with 3 repetitions, because it was noticed that for this number of repetitions the standard deviation was quite low already. This is done to ensure reliability of the results. The seeds are chosen randomly for each repetition and are the same across different experiments. It is aimed for each experiment to last 15 minutes in order to speed up the fine tuning process and to be able to perform several different experiments. In order to achieve this run time, the T_{start}, T_{stop}, α and L are tuned. For clarification, this section is about the experiments for the fine tuning of the parameters of the ALNS. In Section 6.2 we will perform the experiments that will provide a final output and we call these the final experiments. When conducting the final experiments, the values of the T_{start}, T_{stop}, α and L will differ since the runtime can be extended. During fine-tuning, the goal is to perform numerous experiments quickly to explore various parameters. However, in the final experiments, ensuring robustness in the results is prioritized over the speed of the process. Therefore, these experiments are to get an idea of the proportion between these parameters. It can be seen that the largest improvement is achieved in experiment number 9. Below we will discuss each parameter of the ALNS in more detail.

Table 5: Experiments for fine tuning of the parameters of the ALNS

<table>
<tr><th>Exp.</th><th>Iter.</th><th>T_{start}</th><th>T_{stop}</th><th>α</th><th>L</th><th>RW</th><th>Dod</th><th>Run time</th><th>Outp.</th><th>Impr.</th><th>St. dev.</th></tr>
<tr><td>1</td><td>5400</td><td rowspan="19">10</td><td>0.00001</td><td rowspan="20">0.05</td><td rowspan="21">20</td><td rowspan="10">0.1</td><td>1</td><td>1087.81</td><td>594.13</td><td>33.15</td><td>0.76</td></tr>
<tr><td>2</td><td>3060</td><td>0.004</td><td>5</td><td>912.84</td><td>581.02</td><td>46.26</td><td>0.65</td></tr>
<tr><td>3</td><td>3600</td><td>0.001</td><td>10</td><td>1209.10</td><td>583.24</td><td>44.05</td><td>0.65</td></tr>
<tr><td>4</td><td>1800</td><td>0.1</td><td>20</td><td>841.06</td><td>591.26</td><td>36.03</td><td>1.63</td></tr>
<tr><td>5</td><td>1800</td><td>0.1</td><td>30</td><td>1029.69</td><td>593.80</td><td>33.49</td><td>0.67</td></tr>
<tr><td>6</td><td>1100</td><td>0.6</td><td>50</td><td>868.14</td><td>611.11</td><td>16.17</td><td>0.84</td></tr>
<tr><td>7</td><td>640</td><td>2</td><td>100</td><td>834.21</td><td>626.62</td><td>0.67</td><td>0.39</td></tr>
<tr><td>8</td><td>480</td><td>3</td><td>200</td><td>1089.45</td><td>626.30</td><td>0.98</td><td>1.51</td></tr>
<tr><td>9</td><td>2280</td><td>0.03</td><td>30 to 1</td><td>936.08</td><td>580.67</td><td>46.62</td><td>1.08</td></tr>
<tr><td>10</td><td>1800</td><td>0.1</td><td>50 to 1</td><td>982.68</td><td>587.10</td><td>40.18</td><td>1.07</td></tr>
<tr><td>11</td><td>2280</td><td>0.03</td><td>0.05</td><td rowspan="15">30 to 1</td><td>951.72</td><td>586.13</td><td>41.16</td><td>1.92</td></tr>
<tr><td>12</td><td>2280</td><td>0.03</td><td>0.15</td><td>930.86</td><td>591.49</td><td>35.79</td><td>0.97</td></tr>
<tr><td>13</td><td>2280</td><td>0.03</td><td>0.5-0.05</td><td>894.72</td><td>598.65</td><td>28.63</td><td>1.27</td></tr>
<tr><td>14</td><td>2280</td><td>0.03</td><td>0.15-0.05</td><td>990.82</td><td>595.22</td><td>32.06</td><td>1.92</td></tr>
<tr><td>15</td><td>2280</td><td>0.03</td><td>0.2-0.05</td><td>946.05</td><td>582.52</td><td>44.76</td><td>1.06</td></tr>
<tr><td>16</td><td>2280</td><td>0.03</td><td>0.01-0.2</td><td>900.92</td><td>590.96</td><td>36.32</td><td>0.73</td></tr>
<tr><td>17</td><td>2280</td><td>0.03</td><td>0.05-0.15</td><td>975.47</td><td>587.21</td><td>40.08</td><td>1.71</td></tr>
<tr><td>18</td><td>2280</td><td>0.03</td><td>0.05-0.2</td><td>970.81</td><td>580.86</td><td>46.43</td><td>1.75</td></tr>
<tr><td>19</td><td>2280</td><td>0.03</td><td>0.05-0.3</td><td>970.03</td><td>594.07</td><td>33.22</td><td>1.58</td></tr>
<tr><td>20</td><td>2280</td><td>20</td><td>0.06</td><td>0.1</td><td>938.89</td><td>594.99</td><td>32.29</td><td>2.56</td></tr>
<tr><td>21</td><td>2280</td><td>10</td><td>0.56</td><td>0.025</td><td>0.1</td><td>897.55</td><td>600.05</td><td>27.24</td><td>1.33</td></tr>
<tr><td>22</td><td>2280</td><td>10</td><td>0.55</td><td>0.05</td><td>40</td><td>0.1</td><td>947.99</td><td>602.66</td><td>24.62</td><td>0.98</td></tr>
<tr><td>23</td><td>2280</td><td>10</td><td>0.55</td><td>0.05</td><td>40</td><td>0.2</td><td>955.71</td><td>610.74</td><td>16.55</td><td>0.54</td></tr>
<tr><td>24</td><td>2280</td><td>15</td><td>0.04</td><td>0.1</td><td>40</td><td>0.1</td><td>908.95</td><td>585.55</td><td>41.73</td><td>0.26</td></tr>
<tr><td>25</td><td>2280</td><td>4</td><td>0.95</td><td>0.025</td><td>40</td><td>0.1</td><td>1049.91</td><td>599.92</td><td>27.37</td><td>1.80</td></tr>
</table>

Degree of destruction, dod

The dod indicates how many clients will be evaluated to be given a new scheme. So the dod can range from 1 to 879, which would be a completely new solution. A larger dod means a bigger search space whereas a smaller dod is more locally optimizing. It can be observed in Table 5 that larger dod values, such as 100 and 200, show only a slight improvement compared to smaller values of the dod. This could be explained by the complex interaction between the schemes and that by trying to give a large number of clients a new scheme it is challenging for the repair operators to find feasible solutions. By destroying only a smaller number of clients, the ALNS is able to find more feasible solutions. Additionally, a dod that decreases during the execution of the algorithm, is also tested. This is done by decreasing the dod linearly every time the temperature is updated. Experiment number 9 showed the largest improvement and therefore a decreasing dod from 30 to 1 is chosen.

Roulette wheel parameter, RW

The destroy and repair operators are chosen based on their weights. Their weights are updated based on their success in previous iterations. The roulette wheel parameter decides how much weight is given to the initial weights of operators and the success of operators in previous iterations. A value of 0 would mean all operators will keep the same weight as the initial weight and a value of 1 would mean that the previous set of iterations determines their weights. It is noticed in Table 5 that a lower value for the RW is preferred. This can be explained by the length of the Markov Chain length and the number of operators. Since the Markov Chain length is not very high, some operators are not selected and a high RW would then lead to these operators even having less chance of being selected in the next set of iterations. Based on the results of the experiments, the RW is set to 0.1.

The other parameters, T_{start}, α, L and T_{stop}

Firstly, a higher starting temperature allows for greater exploration at the beginning of the ALNS. This is due to the fact that the chances of accepting worse solutions is higher. This greater exploration could lead to higher quality of solutions. However, the greater exploration also means that it will take longer for the ALNS to converge and the run time will be higher. Secondly, the decrease factor determines how fast the temperature goes down. It is important for balancing between solution quality and the run time of the algorithm. A high decrease factor will mean faster convergence, but possibly lower quality solutions. In contrast, a low decrease factor would mean slower convergence which allows the algorithm to search the solution space more thoroughly and could give possibly higher quality solutions. Furthermore, the Markov Chain length has the same trade off as the starting temperature and the decrease factor. A longer Markov Chain results in a higher running time but can produce higher quality solutions, while a shorter Markov Chain reduces running time but may lead to lower quality solutions. Lastly, when these three parameters are set, then the stopping temperature, can be calculated, based on the required running time. As can be seen in Table 5 the largest improvement was achieved for $T_{start} = 10$, $T_{stop} = 0.03$, $\alpha = 0.05$ and $L = 20$. This gave a run time of 936 seconds. However, for the final experiments, we aim for a run time of 1 hour since this is a medium-term planning problem, and immediate solutions are not required. Since we want to conduct several final experiments, 1 hour is a suitable duration that balances thoroughness and efficiency. Thus, we have adjusted the parameters to $T_{start} = 20$, $T_{stop} = 0.068$, $\alpha = 0.025$, and $L = 40$. These settings will yield approximately 9000 iterations and a run time of about an hour.

5.2.3 Other input parameters

There are several parameters that need to be set for the ALNS. We have chosen to tune several of them which is discussed in Section 5.2.2. The values added to the success score of the destroy and repair operator still need to be determined. Firstly, σ_1 is added when the new found solution is the best one found so far. Secondly, σ_2 is added when the new found solution is an improvement, but not the best solution found so far. Thirdly, σ_3 is added when the new found solution is not better than the current solution, but it is still accepted. Lastly, when the solution is rejected, σ_4 is added. The values have to be merit-based and they are set as $\sigma_1 = 3$, $\sigma_2 = 2$, $\sigma_3 = 1$ and $\sigma_4 = 0$. These values σ_1, σ_2, σ_3 and σ_4 remain untuned, as tuning all parameters is resource-intensive and the other parameters in Section 5.2.2 are deemed more influential. Additionally, the weights of the destroy and repair operators are all set to 1, since all operators should have the same weight at the beginning. In addition, the probabilities of the destroy and repair operators are also configured to ensure they are equally likely to be chosen.

5.3 Different scenarios

In this section, we will examine various instances provided to the ALNS to evaluate the impact of different scenarios. In Table 6 an overview of the different scenarios are given. The first column gives the experiment number and the second column gives a description of the experiment. Then in the in third column the number of daily nurses working are given and in the fourth column the number of time slots on a day is given. Furthermore, in the fifth column the total number of time slots over the planning horizon is given and in the sixth column the number of clients is given. Moreover, the seventh column shows whether clients can be moved from their starting week or not. In the last column the options for the schemes are given. A dash indicates that all the schemes are possible considering whether moving is allowed or not.

Table 6: Overview of the different scenarios

Exp. num.	Description of experiment	Num. of nurses	Num. of slots/day	Num. of tot. slots	Num. of clients	Moving possible?	Schemes
1	Flexible clients	2	2	520	879	Yes	-
2	Complete flexibility	2	2	520	879	Yes	All options[1]
3	Number of daily routes	3	2	390	879	Yes	-
4	Clients choose one slot	3	2	390	879	Yes	Clients choose one slot
5	Clients choose two slots	3	2	390	879	Yes	Clients choose two slots
6	No movement of clients	3	2	390	879	No	-
7	More clients	3	2	390	1400	Yes	-

[1] Infeasible in practise.

5.3.1 Flexible clients

In this scenario, it is as in the Injection Team currently with the freedom of planning clients in time slots as we would prefer. So we have 2 nurses working on each day and the clients can be moved from their current starting week as discussed in Section 5.1.1.

5.3.2 Complete flexibility

In this scenario, we do not take into account the starting weeks of clients. This means that we have complete freedom to plan appointments in any week for all clients. This would show the result of what could be achieved if an appointment of a client could be freely moved. However, note that this solution is infeasible in practice. Since not all clients can have their appointment in any week.

5.3.3 Number of daily routes

At the moment, the Injection Team has two nurses working on a day which results into two daily routes. However, as we have noticed, in weeks when a lot of clients need an injection, the demand is too high. So in this scenario we will test the influence of having three daily routes. All the other settings remain the same as the previous scenario.

5.3.4 Clients choose one slot

In this scenario, it will be tested what the influence of the flexibility of clients in their appointment time is. We will test the scenario in which clients can freely choose one time slot in which they want their appointment. In this case, three daily routes will be used, as it is not possible to provide this service with two daily routes and still have feasible solutions. This is because the demand for two nurses is already high.

5.3.5 Clients choose two slots

This scenario is similar to the previous one, except now clients choose two time slots. Then ExpertCare decides which one of the two chosen time slots, will be assigned to the client. In this case, again three daily routes will be used, as it is not possible to provide this service with two daily routes and still have feasible solutions.

5.3.6 No movement of clients

In this scenario, we will test the influence of not being able to move clients to other weeks than their starting week. The number of nurses is in this case increased to 3 nurses, otherwise the model is not able to plan the clients. This is due to very busy weeks. This is tested to see what the influence is of using the possibility of moving clients to other weeks. ExpertCare wants to use an extra nurse in the future next to the 3 current daily routes. This extra nurse will handle all disturbances to the schedule and handle the first appointment of new clients. The extra nurse can also make sure that a client is moved to a more preferable week. In this experiment we can see how much that will help the planning.

5.3.7 More clients

In this scenario, we will add clients to the current client pool to see what the influence is on the solution. Given the steady growth in the number of clients of the Injection Team over the past years, as shown in Figure 3, it is important for ExpertCare to assess the implications of managing a higher client load. The decision has been made to increase the number of clients to 1400 in collaboration with the Injection Team. This is done by artificially creating new clients and assigning them a location based on the distribution of the locations of the current client pool. Moreover, a frequency is assigned to the clients based on the occurrence of frequencies observed by the current client pool. Lastly, based on their frequencies a random starting week is selected for the generated clients.

5.4 Conclusion

This section concludes this chapter where we focused on the research question: *How can we tune the solution approach to ensure high-quality outcomes?*. Firstly, we discussed the data specific input of the Injection Team needed for our model. Secondly, we discussed the input parameters needed for the ALNS. Lastly, we described the different experiments we will test with the ALNS. The results of this will be described in the next chapter.

6 Results

This chapter discusses the results of the experiments conducted. It answers the following research question: *How does the solution approach perform under experimental settings?*. In Section 6.1 the performance of the ALNS is analyzed. Additionally, Section 6.2 shows the obtained results from all the different experiments conducted which we described in Section 5.3. Lastly, Section 6.3 concludes this chapter.

6.1 Performance Adaptive Large Neighborhood Search

In this section we will explore the performance of the ALNS. Firstly, in Section 6.1.1 the ALNS is compared with an exact model. Secondly, in Section 6.1.2 the convergence of the ALNS is analyzed. Lastly, in Section 6.1.3 the performance of the operators of the ALNS is examined.

6.1.1 Comparison exact model

To see how well the ALNS performs it is compared with the exact model developed in Section 4.1. The exact model (integer linear programming, ILP) without the time windows is developed using Python and Gurobi (version 11.0.0). For both models it is possible to move clients to a different week than their starting week. It is compared with our ALNS, having time slots lasting a whole day. The following settings are used for the ALNS: $T_{start} = 10$, $T_{stop} = 0.01$, $\alpha = 0.05$, $L = 20$, RW$= 0.1$ and the dod$=1$. The dod is set lower than before, since we are testing on a smaller client pool. The results are displayed in the table below. The solutions are given in hours and the run time is given in seconds. The last column shows the difference between the ILP solution and the ALNS solution, namely ALNS solution minus ILP solution. We have tested the model for different instances. In the first one all clients that have a weekly or bi-weekly frequency are tested. That are 4 clients with frequency 1 and 9 clients with frequency 2. This is done over a planning horizon of 26 weeks. The second one also includes the 8 clients with a frequency of 3. The third one is the same as the previous one, only now the planning horizon is 4 weeks. In the last one also the clients with a frequency of 4 weeks are included. These amount to 148 additional clients. It is seen from the results that the ALNS is able to generate solutions close to the solutions of the ILP. In the more complex instances, it is clear that the run time of the ILP increases much and in the last instance is not able to produce a solution in an hour and also no upper bound is found since it is out of memory.

Table 7: Results comparison ALNS with ILP

Description	ILP solution	ILP run time	ILP integrality gap	ALNS solution	ALNS run time	Δ ALNS and ILP
Frequency 1 and 2 over 26 weeks	68.46	13.7	0%	68.86	55.88	0.40
Frequency 1, 2 and 3 over 26 weeks	75.27	3600[1]	13.90%	84.44	65.53	9.17
Frequency 1, 2 and 3 over 4 weeks	11.52	3600[1]	6.06%	13.17	22.33	1.65
Frequency 1, 2, 3 and 4 over 26 weeks	No solution	3600[1]	-	302.88	228	-

[1] The model was terminated after 3600 seconds.

6.1.2 Convergence of Adaptive Large Neighborhood Search

To see the performance of the ALNS the solutions found by the ALNS and the best solution found so far per iteration is shown in Figure 16. This is for a repetition from experiment 1. The other repetitions can be seen in Appendix F in Figures 28, 29, 30 and 31. It is seen that these behave similarly. On the y-axis the value represents the total travel time in hours over half a year and the x-axis shows the number of iteration. It is evident that the ALNS algorithm converges to an optimized solution. Initially, it accepts more worse solutions, while later in the process, it becomes more selective, accepting fewer worse solutions. This is due to the decrease in the temperature of the ALNS. In Figure 17 the performance of the ALNS for all repetitions of experiment 1 in one plot is displayed. This shows the spread of the ALNS. We can see that repetition 2 is a little higher than the others in the last iteration. The others are relatively close to each other.

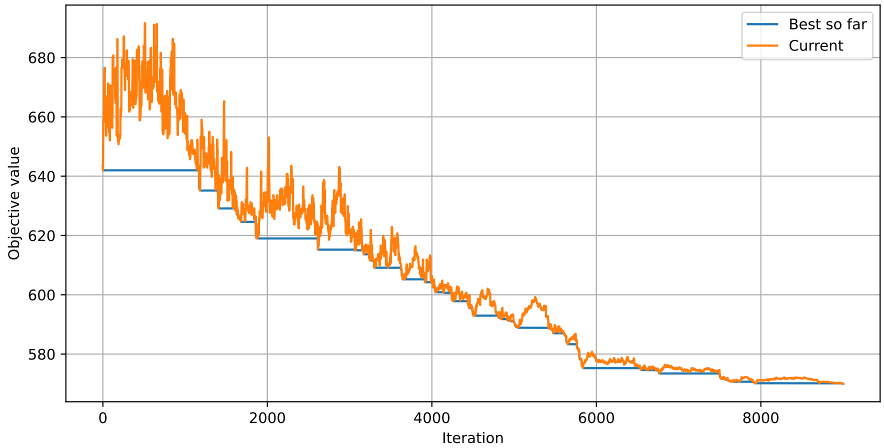

FIGURE 16: Performance of the ALNS for the first repetition of experiment 1

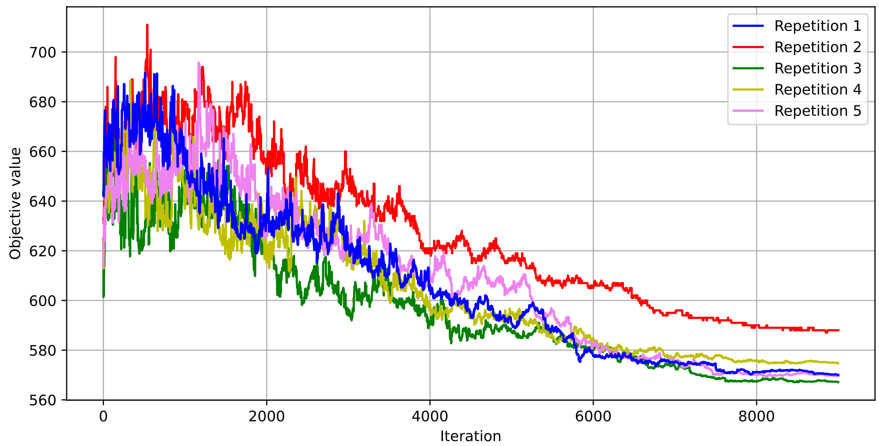

FIGURE 17: Performance of the ALNS for all the repetitions of experiment 1

6.1.3 Operator performance of Adaptive Large Neighborhood Search

In Figure 18 the destroy operator performance of the ALNS for a repetition of experiment 1 is shown. The other repetitions can be seen in Appendix F in Figures 32, 33, 34 and 35. It shows the success rate π_j of each operator j over the iterations. It can be seen that all operators are increasing and not converging yet. Meaning that all operators are chosen throughout the execution of the ALNS. Also, it can be noticed that the worst destroy is increasing more slowly as the number of iterations increases. This is even more clear in other repetitions of experiment 1. This may be because the clients located farthest away are already scheduled efficiently. Additionally, in Figure 19 the repair operator performance of the ALNS for a repetition of experiment 1 is displayed. The other repetitions can be seen in Appendix F in Figures 36, 37, 38 and 39. It is evident that the random repair operator is outperformed by the other two operators. The other two operators are similar in their performance. This can be explained by the similarity of the two operators.

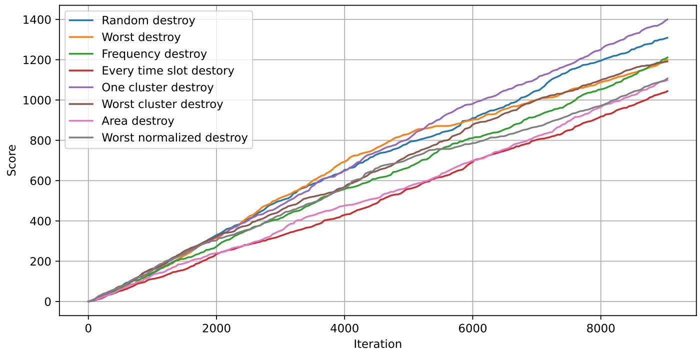

FIGURE 18: Performance of the destroy operators for the first repetition of experiment 1

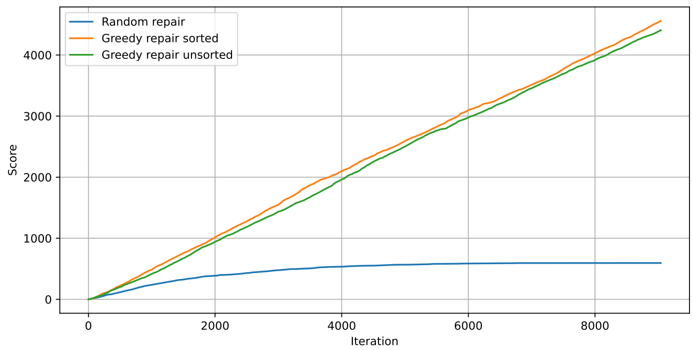

FIGURE 19: Performance of the repair operator for the first repetition of experiment 1

6.2 Results experiments

The previously discussed experiments have been executed in Python (version 3.11). Each experiment has been done for 5 repetitions. We have chosen for the final output 5 repetitions instead of 3 repetitions in the fine tuning in Section 5.2.2, because we wanted to speed up the fine tuning process and we want more robust and reliable output in the final output. Additionally, from the low standard deviation observed in the fifth column of Table 8 it is clear that our outputs are stable. The average results over all repetitions is shown in Table 8. The first column shows the number of the experiment and the second column gives a description of the experiment. The results of the travel time are shown in Table 8, in the third column. This is given in hours. The fourth column is the standard deviation of the solution of the ALNS of all the repetitions. The standard deviation is also given in percentage in the fifth column. Additionally, the run time is given in seconds in the sixth column. Moreover, the average traveling hours per nurse per day is given in the seventh column. Furthermore, the average number of appointments per nurse per day is given in the eighth column and the average number of nursing hours per nurse per day is given in the ninth column. Lastly, the average number of days in which routings are done is given in the last column. For the first two experiments this can be at most 260, because 2 nurses work 5 times a week for 26 weeks. Additionally, for the other experiments it can be at most 390, because 3 nurses work 5 times a week for 26 weeks. The second row shows the values of the Injection Team over 2023 which we obtained in Section 2.5. It is not an experiment, but it is included in the table for comparative purposes. The results are assessed on the three performance indicators discussed in Section 2.5, namely the travel time, the continuity of care and the balanced working days for the nurses. We will first discuss the performance indicators briefly and then we will discuss the results of each experiment in more detail.

Travel time

Firstly, an important performance indicator is the travel time. In Figure 20 the average number of traveling hours per nurse per day is shown for each experiment during the first repetition. The results of all repetitions of experiment 1 and the other experiments can be found in Appendix G and H, respectively. Recall that in Section 2.5.1 it was determined that on average the travel time per day per nurse is around 2.63 hours for the Injection Team. Additionally, considering two nurses working each day and having 130 working days in half a year, this amounts to 684.08 hours of traveling time in half a year.

Continuity of care

Secondly, another performance indicator we discussed is the continuity of care. However, we do not show the results for the continuity of care of our experiments, because the coc would always be 1 and the ratio of nurse switches 0. This is due to the x-weekly schedules of clients, combined with nurses having fixed working days. This means that each client will always be helped by the same nurse and therefore the coc would be equal to 1 and the ratio of nurse switches 0. Only when adjustments are made to the medium-term planning the coc would be affected. So for the Injection Team it is easier to ensure continuity of care compared to the literature. Since in the literature clients often need care multiple times per week, which makes it is more difficult to ensure it is the same nurse every time.

Balanced working days for the nurses

Thirdly, the balanced working days for the nurses is another performance indicator that we discussed. In Figure 21 the average number of appointments per nurse per day is shown for each experiment during the first replication. Additionally, Figure 22 shows the average number of nursing hours per nurse per day for each experiment during the first replication. The results of all repetitions of experiment 1 and the other experiments can be found in Appendix G and H, respectively. Recall that in Section 2.5.3 it was determined that the Injection Team had on average 8.02 appointments per nurse per day and 5.92 hours of nursing time per nurse per day. This would mean an average working day of a nurse is around $5.92 + 2.63 = 8.55$ hours. This is higher than the desired average 8 hours of working for nurses. In experiment 1 and 2 which also have 2 daily nurses working have a much higher average number of nursing hours per nurse per day, namely for both 6.94 hours. This increase compared to the Injection Team in 2023 can be explained by the increase in clients over 2023 which resulted into more appointments and higher nursing hours. The working hours found in experiment 1 and 2 are 9.15 and 8.98, respectively. This indicates that an extra daily nurse is needed to decrease the workload for the nurses. The working hours are found by taking the sum of the average traveling hours and the average nursing hours per nurse per day.

Table 8: Results of all the experiments

Exp. num.	Description of experiment	Solution ALNS	St. dev.	St. dev. %	Run time	Avg. travel.	Avg. app.	Avg. nurs.	Avg. days
-	Injection Team	684.08	-	-	-	2.63	8.02	5.92	260.00
1	Flexible clients	573.67	8.18	1.43	3824.38	2.21	9.26	6.94	260.00
2[1]	Complete flexibility	530.14	10.96	2.07	10931.24	2.04	9.25	6.94	260.00
3	Number of daily routes	610.68	8.37	1.37	4650.58	1.88	7.39	5.54	324.60
4	Clients choose one slot	785.87	8.78	1.12	2039.16	2.02	6.18	4.64	390.00
5	Clients choose two slots	727.27	7.37	1.01	2817.76	1.89	6.27	4.70	384.00
6	No movement of clients	607.88	6.43	1.06	3777.97	1.71	6.86	5.14	356.00
7	More clients	928.36	17.73	1.91	5045.50	2.38	9.74	7.30	389.40

[1] Infeasible in practise.

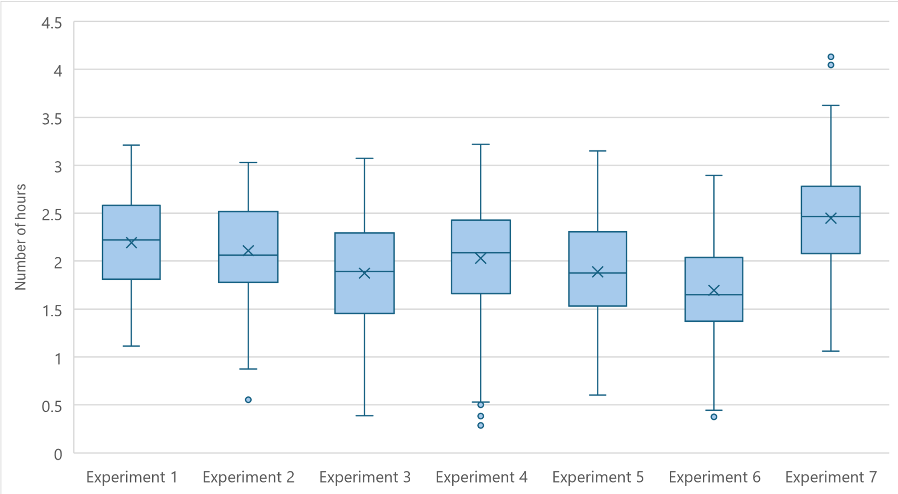

FIGURE 20: Average number of traveling hours per nurse per day of all experiments

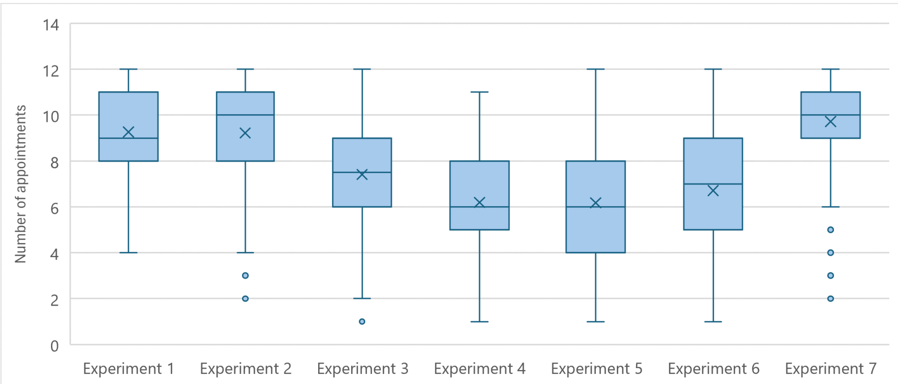

FIGURE 21: Average number of appointments per nurse per day of all experiments

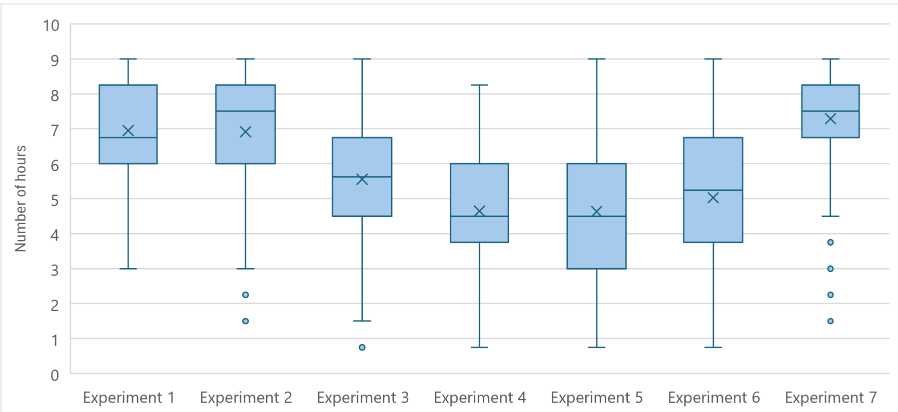

FIGURE 22: Average number of nursing hours per nurse per day of all experiments

6.2.1 Flexible clients

In this scenario, the planners can freely choose the ideal appointment slot for each client, considering time windows, working hours and the starting weeks of clients. The results can be seen in Table 8 in experiment number 1. The ALNS provides a total travel time of 573.67 hours over half a year. This gives a reduction of 16.14% compared to the 684.08 hours of the Injection Team in 2023. Considering this reduction in travel time, the productivity of ExpertCare would increase to 60.40%, which is an improvement of 1.7%. The total number of days nurses worked are 260 and therefore the average travel time per day per nurse is around 2.21 hours. This gives a reduction of 0.42 hours per nurse per day compared to the values found of the Injection Team over 2023. Additionally, the average number of appointments per nurse per day is 9.26 and the average number of nursing hours per nurse per day is 6.94. These values are both a bit higher than the values found of the Injection Team over 2023. This can be caused by the steady increase of clients, which we have seen in Figure 3. This means that the average working hours of nurses per day is around 9.15 hours which is higher than the nurses' wishes of around 8 hours on a day. This means that an additional nurse is needed to fulfill this wish. The according clustering is shown in Figure 23. To see the individual clusters more clearly and have additional data of the clients in the clusters, the reader is referred to Appendix G. It can be noticed that each cluster contains clients in the same region. In cluster 4 a large amount of clients with frequency 12 is seen, whereas in cluster 3 and 7 a large amount of frequency 26 is seen. Clients with frequency 4 is present in each cluster and is more evenly distributed among the clusters. Moreover, as discussed in Section 4.3.2 a penalty was introduced to ensure that the model is not biased towards schemes with fewer appointments. However, there are still approximately 33 appointments less than the initial input. This can be explained by the fact that the weeks with the most clients cannot be handled and therefore have to be moved to another week. Thus, the solution remains on average around 2408.20 appointments for it to be feasible and the penalty makes sure it does not decrease more.

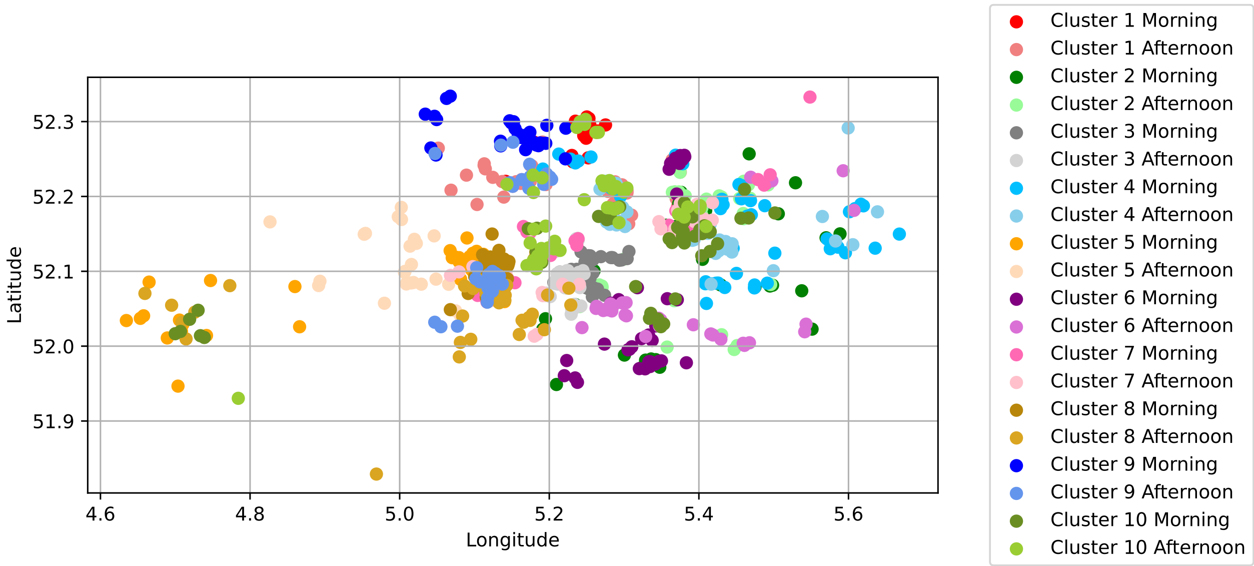

FIGURE 23: Clustering for a repetition of experiment 1

6.2.2 Complete flexibility

In this scenario the starting weeks of clients is not taken into account. This means that an appointment of a client can be planned in any week. However, as discussed before this is not feasible in practise. The results are shown in Table 8 in experiment number 2. It is clear this experiment has the lowest outcome and the highest running time. This makes sense, because this one has more flexibility and therefore more options for each client to consider. This increases the run time and decreases the total travel time. This reduction in the total travel time highlights the importance of having the ability to move clients away from their starting week. The average number of appointments per nurse per day is 9.25 and the average number of nursing hours per nurse per day is 6.94. The clustering is shown in Figure 24. It is again seen that the clients in each cluster are close to each other in the region.

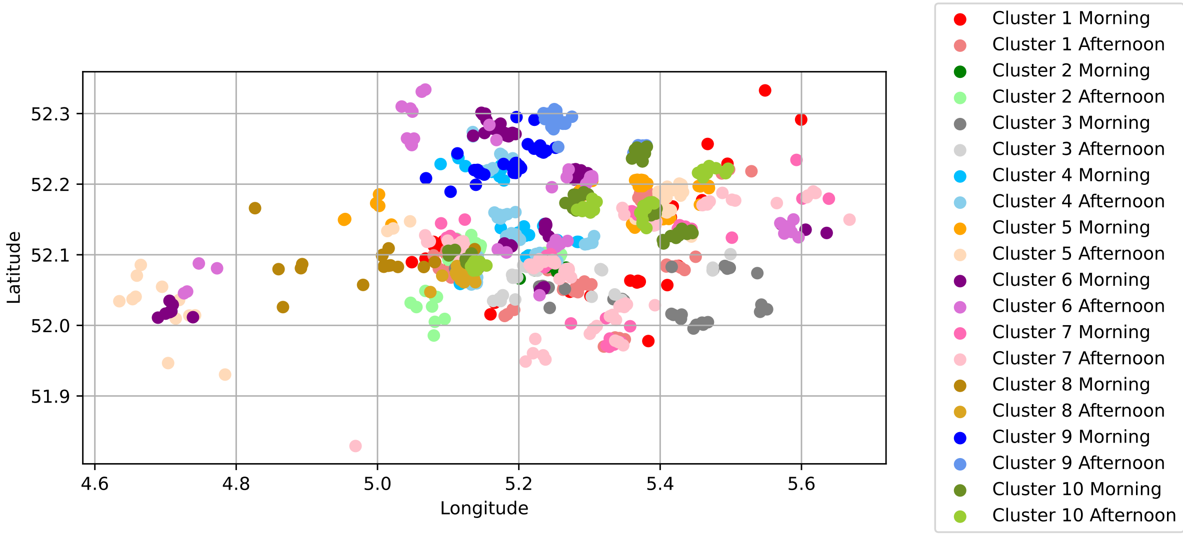

FIGURE 24: Clustering for a repetition of experiment 2

6.2.3 Number of daily routes

In this scenario the number of daily routes is increased to three. The results can be seen in Table 8 in experiment number 3. Comparing it to experiment number 1 we see that the total travel time has increased. This could be caused by the many available slots in the schedule. The repair operators can assign clients more easily to empty slots, even though assigning it to other slots is more efficient. This is especially true for the random repair operator. When the clients are more spread out over all the time slots, the algorithm struggles to move them to fewer time slots, since there is no destroy operator that destroys an entire time slot. This could explain the increase in the total travel time. To support our claims we compared the average number of appointments per nurse per day to experiment 1. The average number of appointments per nurse per day is 7.39, whereas in experiment 1 it is equal to 9.26. This is almost two clients less on each day per nurse compared to experiment 1. Additionally, the average nursing time per nurse per day is 5.54 hours, which is also lower since there are less appointments on each day. Additionally, it is seen that the performance of the random repair operator is very low in Figure 48 in Appendix H. This can be explained by the same argument. The random repair operator has a larger chance of assigning clients to empty slots which is not always preferred.

6.2.4 Clients choose one slot

In this scenario the clients can freely choose one time slot in which they would like to have their appointment. This experiment also has three daily routes to provide feasible solutions. The results are shown in Table 8 in experiment number 4. It is seen that the total travel time is much higher than in all the other experiments, except the last one. This makes sense since there is no flexibility at all possible apart from moving clients from their starting weeks. Additionally, the run time is much lower. This makes sense since the number of options of schemes for each client is greatly reduced. The average number of appointments per nurse per day is 6.18 and the average number of nursing hours per nurse per day is 4.64 hours. The number of appointments and nursing hours per nurse per day have decreased. This is because the travel time has increased since the options for scheduling clients is less flexible. The clustering is shown in Figure 25. It is clear by comparing it to Figure 23, the clusters are a lot more spread out. This is indeed expected as the travel time is much higher. The performance of the destroy operators is also displayed in Appendix H in Figure 49. It can be noticed that the worst destroy and worst normalized destroy almost completely converge after half the iterations. However, this does only happen in half of the repetitions. Apparently, after enough iterations, the "farthest" clients cannot be assigned a better scheme.

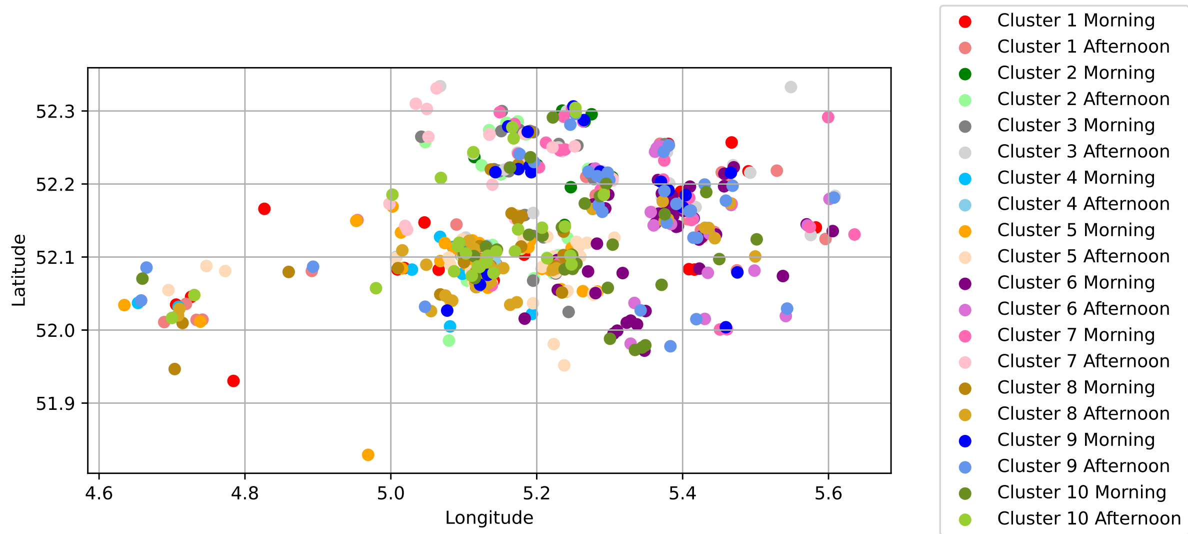

FIGURE 25: Clustering for a repetition of experiment 4

6.2.5 Clients choose two slots

This scenario is similar to the previous one, except now a client can choose two slots. The results are shown in Table 8 in experiment number 5. It is noticed that the travel time has decreased compared to the previous scenario, which is due to the higher flexibility in planning clients. It gives a reduction of 58.6 hours over half a year when clients choose two time slots instead of one time slot. The average number of appointments per nurse per day is 6.27 and the average number of nursing hours per day is 4.70, which is comparable to the previous experiment.

6.2.6 No movement of clients

In this scenario the case of not moving the clients to another week than their starting week is tested. The results of this can be seen in Table 8 in experiment number 6. Comparing it to experiment number 3 the travel time is only slightly lower and the run time is lower. This can be explained by the fact that in this experiment, there are fewer schemes for clients since they cannot be moved away from their starting week. This speeds up the algorithm. We would have expected experiment 6 to have a higher travel than experiment 3, since experiment 3 has more flexibility. However, the difference is 2.8 hours over half a year which is not much. Additionally, the average number of appointments per nurse per day is 6.86 and the average number of nursing hours per nurse per day is 5.14 hours. These are lower than in experiment 3, but that can be explained by experiment 3 having more days in which there are no appointments. Again, like in experiment 3 the random repair operator performs poorly, as can be seen in Figure 50 in Appendix H.

6.2.7 More clients

In this scenario the number of clients have been increased to 1400. The results of this can be seen in Table 8 in experiment number 7. This solution has the highest travel time among all the experiments, namely 928.36 hours. This can be explained by the fact that the number of clients has increased by more than 50%. This results in more appointments and therefore a higher travel time. The average number of appointments per nurse per day is 9.74 and the average number of nursing hours per nurse per day is 7.30 hours. This is much higher than seen in the other scenarios. It gives on average 9.68 working hours for the nurses. This is much higher than 8 hours and therefore an extra nurse would be needed to decrease the workload.

6.3 Conclusion

This section summarizes the results, concentrating on the research question: *How can we improve the medium-term planning in order to enhance the productivity of the Injection Team of ExpertCare?*. Firstly, we investigated the performance of the ALNS. Secondly, we conducted several experiments using our proposed ALNS. We tested the option of freely planning the clients which gave a travel time of 573.67 hours over half a year which amounts to a reduction in the travel time of 16.14% compared to the Injection Team in 2023. However, this schedule was observed to be very full. Additionally, we explored a scenario with three nurses working each day and analyzed the influence of clients choosing one or two time slot(s) they prefer. Furthermore, we examined the impact of not exploring the option of moving clients from their starting week and the influence of having 1400 clients.

7 Conclusion and recommendations

This chapter concludes the paper and answers the following research question: *What are the conclusions we can draw and the recommendations we can give to ExpertCare based on the results of this research?*. Firstly, in Section 7.1 the final conclusion of the paper is given. Secondly, in Section 7.2 recommendations for ExpertCare are given. Thirdly, Section 7.3 gives possibilities for future research. Lastly, Section 7.4 discusses the practical and academic contributions of this thesis.

7.1 Conclusions

This research revolves around the core problem "the manual creation of the blueprint relying solely on the planners' knowledge and intuition". The goal is to increase the productivity of the Injection Team of ExpertCare. The productivity in 2023 was reported to be 58.7%. The analysis of the Injection Team's current situation shows that the average travel time per nurse per day is around 2.63 hours. The continuity of care was found to be around 0.65-0.67 and the ratio of nurse switches around 0.57-0.63. Additionally, the average number of appointments per nurse per day is approximately 8.02 and the average nursing hours per nurse per day is around 5.92 hours. The literature review conducted gave insights in how to design a medium-term planning of clients in the context of home health care. It was seen that the visitation frequency of clients of the Injection Team is much lower than seen in the literature. Moreover, in the literature it was seen that an adaptive large neighborhood search is a successful method. Therefore, we have developed a two-phase approach. In the first phase a greedy heuristic constructs an initial solution. Then, in the second phase the ALNS improves the initial solution. The results has been compared with an exact model to assess how the ALNS performs in Section 6.1.1. It was found that the ALNS provides solutions close to the exact model and that for more complex instances the ALNS can still produce solutions within a reasonable time compared to the exact model.

Allowing clients' appointments to be freely planned resulted in a total travel time of 573.67 hours over half a year, averaging 2.21 hours per day. This led to a reduction of 0.42 hours of travel per nurse per day, equating to a 16.14% decrease in travel time compared to the Injection Team in 2023. Consequently, ExpertCare's productivity increased by 1.7%, reaching 60.40%. While the productivity target of 63% was not met, reducing nearly half an hour of travel per nurse each day is a large improvement. Furthermore, it has been concluded that the x-weekly planning of clients and nurses having a fixed working day each week, lead the coc to be 1 and the ratio of switches to be 0. Only when changes are made to the medium-term this is affected. In addition, the average number of appointments and the average number of nursing hours have increased by 1.24 and 1.02, respectively. This can be explained by the steady growth in the number of clients over time. At the beginning of 2023, the number of clients was around 700 and grew to around 900 clients. The average working hours per day per nurse is 9.15 hours which is more than the desired 8 hours. From this we conclude that an additional nurse is needed to fulfill this. Also, the results of the clustering of clients have been analyzed to determine patterns in the clustering of clients. However, we were not able to derive any clear patterns to cluster clients. Consequently, we cannot provide guidelines for assigning clients to specific clusters based on their visitation frequency, location, or appointment week. This is due to the complex interactions between various client parameters, such as visitation frequency, location, and appointment week. Thus, it is concluded that to improve the medium-term planning to enhance the productivity a more sophisticated planning system is needed. We will elaborate on this more in Section 7.2.

Additionally, several other experiments have been conducted to assess the influence of these instances. The case of complete flexibility in planning client appointments is analyzed. In this case the the clients can be planned in any week which is in practise not possible. This resulted in a total travel time of 530.14 which gives a reduction of 153.94 hours compared to the Injection Team in 2023. This reduction in total travel time highlights the importance of using the possibility of moving clients away from their starting weeks. When adding an extra daily route to the planning the total travel time is increased, since the algorithm spreads the clients too much over all the time slots. Another option that is explored is clients having the possibility to choose their time slot as they prefer. This gave a total travel time over half a year of 785.87 hours. This is a large increase compared to freely planning clients. This can be explained by the fact that the options for scheduling clients is less flexible. Furthermore, the option of clients choosing two time slots and then ExpertCare deciding which one, resulted in a total travel time over half a year of 727.27 hours. So the travel time over half a year decreased by 58.6 hours when clients choose two time slots instead of one. Moreover, the case of not being able to move the client's from their starting week results also in an increase of the total travel time compared to freely planning the clients. This is again due to the algorithm inefficiently spreading clients out over the time slots. Lastly, we have increased the number of clients to 1400 to see the impact of this. The total travel time over half a year was found to be 928.36 hours. This gives an increase of 244.28 hours compared to the Injection Team in 2023. However, the average working hours for a nurse is 9.68 hours and therefore an additional nurse is needed to decrease the workload in this instance.

7.2 Recommendations

As concluded from our results a reduction of 16.14% in the travel time can be achieved, but there are no clear patterns found on how to cluster clients. Therefore, it is recommended for ExpertCare to use a more advanced planning system that takes into account all clients' parameters to improve their planning. An option would be to implement a system where each client is scheduled on a specific day for all their future appointments. This assignment of clients to a moment could be determined with our model if the moments when clients can be helped are known. This is then the medium-term planning. When a new client arrives, a model is needed that estimates the best time slots for them based on the current client pool. This model should account for the number of existing appointments on each day and prioritize scheduling on days when clients visit nearby locations. This would also avoid the great deal of time window and working time violations, since the model would be able to determine that a time slot is already full. Moreover, it is advised to consider avoiding a weekly planning method since only 4 clients have a weekly frequency, making them the only constant appointments every week. This means that having 5 time slots for Utrecht, could be too much in one week and not enough in another week. This makes it also complex to derive recommendations for clustering on a weekly basis. Instead, a monthly planning method could be considered. However, given the large number of half-yearly clients, a half-yearly planning system may be even more effective. Moreover, the average working hours for nurses are approximately 9.15 hours per day in our results. To reduce their workload, it is recommended that the Injection Team increase the number of daily nurses from two to three. This recommendation is also based on the consistent growth in the number of clients over time.

Moreover, it is advised for ExpertCare to establish their values and what they want to offer their clients, whether they would like to give the choice to the clients or they give the clients only a few options. From our results an indication is given on how this would change the total travel time. When clients choose one or two time slots the total travel time over half a year is increased by 101.79 and 43.19, respectively compared to the Injection Team in 2023. Thus, it could consider letting the clients decide their appointment. For example, a client could book a slot online which is available. This would mean that the planners and nurses do not have to plan a client's appointment which saves time. However, it does considerably increase the total travel time. Yet with the growing number of clients and the addition of more daily routes could make this option more realistic in the future. Another option is to let clients choose multiple options and then ExpertCare decides from these option. This is more a trade-off between client values and efficient routing. In this case it needs to be determined for each client what options they would like for an appointment. Then based on this data an efficient medium-term planning can be made.

Additionally, there are several interesting topics for future research. We will discuss these in the next section, but we will highlight two of them for the Injection Team of ExpertCare. Firstly, an interesting topic for future research for the Injection Team of ExpertCare is to introduce more time slots on a day. For example, three or four time slots on a day could be considered to see what exact implications this would have for the Injection Team. Secondly, exploring the feasibility of employing a spare nurse to manage disruptions, such as new clients during the week, could be a valuable area of study. It would be interesting to analyze the cost implications and benefits of having such a nurse available.

7.3 Possibilities for future research

This research is subject to certain limitations due to the defined scope and the complexity of ExpertCare's context. These limitations provide a foundation for future research topics. One area to explore is a stochastic planning approach for the problem. In this case the robustness of the planning could be tested by evaluating how flexible the schedule is when a new client arrives. For example, routes that are more spread out might be more adaptable to new clients, as there is a higher chance that they are close to the new client. Additionally, some clients prefer to be helped during a working day, necessitating nurse visits to their workplace. This is not currently considered due to a lack of data on their workplaces and occurs only occasionally. However, this option could be explored in future research. Another limitation is that certain medication patients were left out from our client pool. These are patients who only require care for two weeks or need care every day for a week. Including these patients could be interesting for future research. Furthermore, at the moment the travel time is an approximation. In the future, a better approach might consider traffic load and busy hours. Also, the routing in our ALNS uses the Clarke and Wright heuristic to construct routes. It could be interesting to explore other methods for routing on a day-to-day basis. Another potential topic is introducing flexibility in the model so clients do not necessarily have the same appointment day every week. This should first be considered to see if it is really needed by clients. Investigating the feasibility of having a spare nurse to handle disturbances, such as new clients during the week, could be another area of study. It would be interesting to see the cost implications and benefits of such a nurse. Finally, our model can only handle one or two time slots for appointments on a day. It could be worthwhile to analyze having three time slots for appointments in a day.

7.4 Practical and academic contributions

This section discusses the practical contributions of this research as well as the academic contributions of this research.

7.4.1 Practical contributions

This thesis has several practical contributions. Firstly, this thesis performed a thorough context analysis of the Injection Teams's current situation, providing valuable insights regarding their current client pool. It is also clear for ExpertCare what data can easily be found and what data is difficult or even unavailable to obtain. For example, there was no easy method to acquire the frequency data of the current client pool even though this information is crucial to have a view of how often visitations are required over time. Additionally, data of preferences of clients is often unavailable to acquire. It is always discussed at the first appointment when they would like to have their next appointment, but no where it is kept track of. Additionally, a prototype algorithm is made to efficiently construct a medium-term planning for the current client pool of ExpertCare. From this it is concluded that due to the complex interaction of clients' parameters no clear patterns for clustering clients is found. Therefore, it is recommended to use a more advanced planning system to enhance their productivity. This research shows that the proposed method of constructing a medium-term planning reduces the total travel time and does not violate client' time windows.

7.4.2 Academic contributions

The academic contributions of this work are twofold. Firstly, a planning horizon of 26 weeks is used in our research with much lower visitation frequencies of clients than seen in the literature. In the literature, the visitation frequency is often weekly, daily or multiple daily whereas in our problem the visitation frequency ranges from weekly to semi-annually. Therefore, the planning horizon used in our problem is much longer than seen in the literature. Secondly, the application of an ALNS on a pattern-based problem using schemes on a medium-term planning problem in the context of home care is novel and has not been previously documented in the literature. The model can be generalized to other similar problems, by adapting all the possible schemes a client can have. However, this is only for problems that have one time slot or two time slots on a day.

References

[1] Jonathan F. Bard, Yufen Shao, and Ahmad I. Jarrah. A sequential GRASP for the therapist routing and scheduling problem. *Journal of Scheduling*, 17:109–133, 2014.

[2] Jonathan F. Bard, Yufen Shao, and Huan Wang. Weekly scheduling models for traveling therapists. *Socio-Economic Planning Sciences*, 47(3):191–204, 2013.

[3] E. J. Beltrami and L. D. Bodin. Networks and vehicle routing for municipal waste collection. *Networks*, 4(1):65–94, 1974.

[4] John Bowers, Helen Cheyne, Gillian Mould, and Miranda Page. Continuity of care in community midwifery. *Health Care Manag Sci.*, 18(2):195–204, 2015.

[5] Kris Braekers, Katrien Ramaekers, and Inneke Van Nieuwenhuyse. The vehicle routing problem: State of the art classification and review. *Computers & Industrial Engineering*, 99:300–313, 2016.

[6] Ann Melissa Campbell and Jill Hardin Wilson. Forty years of periodic vehicle routing. *Networks*, 63(1):2–15, 2014.

[7] Paola Cappanera and Maria Grazia Scutellà. Joint assignment, scheduling, and routing models to home care optimization: A pattern-based approach. *Transportation Science*, 49(4):830–852, 2014.

[8] Safa Chabouh, Adnen El-Amraoui, Sondes Hammami, and Hanen Bouchriha. A systematic review of the home health care planning literature: Emerging trends and future research directions. *Decision Analytics Journal*, 7, 2023.

[9] Mohamed Cissé, Semih Yalçındağ, Yannick Kergosien, Evren Şahin, Christophe Lenté, and Andrea Matta. OR problems related to home health care: A review of relevant routing and scheduling problems. *Operations Research for Health Care*, 13-14:1–22, 2017.

[10] G. Clarke and J. W. Wright. Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research*, 12(4):568–581, 1964.

[11] Jalel Euchi, Malek Masmoudi, and Patrick Siarry. Home health care routing and scheduling problems: a literature review. *4OR*, 20:351–389, 2022.

[12] Amir Mohammad Fathollahi-Fard, Abbas Ahmadi, Fariba Goodarzian, and Naoufel Cheikhrouhou. A bi-objective home healthcare routing and scheduling problem considering patients' satisfaction in a fuzzy environment. *Applied Soft Computing*, 93:106385, 2020.

[13] Thomas A. Feo, Mauricio G. C. Resende, and Stuart H. Smith. A greedy randomized adaptive search procedure for maximum independent set. *Operations Research*, 42(5):860–878, 1994.

[14] Fariba Goodarzian, Hossein Shokri Garjan, and Peiman Ghasemi. A state-of-the-art review of operation research models and applications in home healthcare. *Healthcare Analytics*, 4, 2023.

[15] Chris Groër, Bruce Golden, and Edward Wasil. The consistent vehicle routing problem. manufacturing & service operations. *Manufacturing & Service Operations*, 11:630–643, 2009.

[16] Daniela Guericke. *Routing and scheduling for home care services: Solution approaches for static and dynamic settings.* PhD thesis, Universität Paderborn, 2016.

[17] Daniela Guericke and Leena Suhl. The home health care problem with working regulations. *OR Spectrum*, 39:977–1010, 2017.

[18] Mike Hewitt, Maciek Nowak, and Nisha Nataraj. Planning strategies for home health care delivery. *Asia-Pacific Journal of Operational Research*, 33(5):91–102, 2016.

[19] Sarah Hurkmans, M. Yousef Maknoon, Rudy R. Negenborn, and Bilge Atasoy. An integrated territory planning and vehicle routing approach for a multi-objective residential waste collection problem. *Transportation Research Record*, 2675(7):616–628, 2021.

[20] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.

[21] Grigorios D. Konstantakopoulos, Sotiris P. Gayialis, and Evripidis P. Kechagias. Vehicle routing problem and related algorithms for logistics distribution: a literature review and classification. *Operational Research*, 22:2033–2062, 2022.

[22] Attila A. Kovacs, Bruce L. Golden, Richard F. Hartl, and Sophie N. Parragh. The generalized consistent vehicle routing problem. *Transportation Science*, 49(4):796–816, 2014.

[23] Tarık Küçükdeniz and Beril Şevval Develioğlu. Schedule and routing in home healthcare system using clustering analysis and multi-objective optimization. *Artificial Intelligence-based Healthcare Systems*, pages 1–21, 2023.

[24] M. Y. Lin, K. S. Chin, and K. L. Tsui. Patient assignment and grouping problem on home health care: A case in hong kong. In Cher Ming Tan and Thong Ngee Goh, editors, *Theory and Practice of Quality and Reliability Engineering in Asia Industry*, pages 151–164. Springer Singapore, 2017.

[25] P.A. Maya Duque, M. Castro, K. Sörensen, and P. Goos. Home care service planning. the case of landelijke thuiszorg. *European Journal of Operational Research*, 243(1):292–301, 2015.

[26] Daniel Merchán and Matthias Winkenbach. An empirical validation and data-driven extension of continuum approximation approaches for urban route distances. *Networks*, 73(4):418–433, 2019.

[27] N. Mladenović and P. Hansen. Variable neighborhood search. *Computers & Operations Research*, 24(11):1097–1100, 1997.

[28] Isabel Méndez-Fernández, Silvia Lorenzo-Freire, Ignacio García-Jurado, Julián Costa, and Luisa Carpente. A heuristic approach to the task planning problem in a home care business. *Health Care Management Science*, 23:556–570, 2020.

[29] Stefan Nickel, Michael Schröder, and Jörg Steeg. Mid-term and short-term planning support for home health care services. *European Journal of Operational Research*, 219(3):574–587, 2012.

[30] Erfaneh Nikzad, Mahdi Bashiri, and Babak Abbasi. A matheuristic algorithm for stochastic home health care planning. *European Journal of Operational Research*, 288(3):753–774, 2021.

[31] Delaram Pahlevani, Babak Abbasi, John W. Hearne, and Andrew Eberhard. A cluster-based algorithm for home health care planning: A case study in australia. *Transportation Research Part E: Logistics and Transportation Review*, 166:102878, 2022.

[32] Inmaculada Rodríguez-Martín, Juan-José Salazar-González, and Hande Yaman. The periodic vehicle routing problem with driver consistency. *European Journal of Operational Research*, 273(2):575–584, 2019.

[33] Stefan Ropke and David Pisinger. An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation Science*, 40(4):455–472, 2006.

[34] Paolo Serafini. Simulated annealing for multi objective optimization problems. In *Multiple Criteria Decision Making*, pages 283–292, New York, NY, 1994. Springer New York.

[35] Yufen Shao, Jonathan F. Bard, and Ahmad I. Jarrah. The therapist routing and scheduling problem. *IIE Transactions*, 44(10):868–893, 2012.

[36] Paul Shaw. Using constraint programming and local search methods to solve vehicle routing problems. In Michael Maher and Jean-Francois Puget, editors, *Principles and Practice of Constraint Programming — CP98*, pages 417–431, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.

[37] Kenneth Sorensen, Florian Arnold, and Daniel Palhazi Cuervo. A critical analysis of the "improved clarke and wright savings algorithm". *International Transactions In Operational Research*, 26(1):54–63, 2017.

[38] Jacqueline Wirnitzer, Iris Heckmann, Anne Meyer, and Stefan Nickel. Patient-based nurse rostering in home care. *Operations Research for Health Care*, 8:91–102, 2016.

[39] Semih Yalçındağ, Paola Cappanera, Maria Grazia Scutellà, Evren Şahin, and Andrea Matta. Pattern-based decompositions for human resource planning in home health care services. *Computers & Operations Research*, 73:12–26, 2016.

[40] Lu Zhen, Jiajing Gao, Zheyi Tan, Gilbert Laporte, and Roberto Baldacci. Territorial design for customers with demand frequency. *European Journal of Operational Research*, 309(1):82–101, 2023.

A Research design

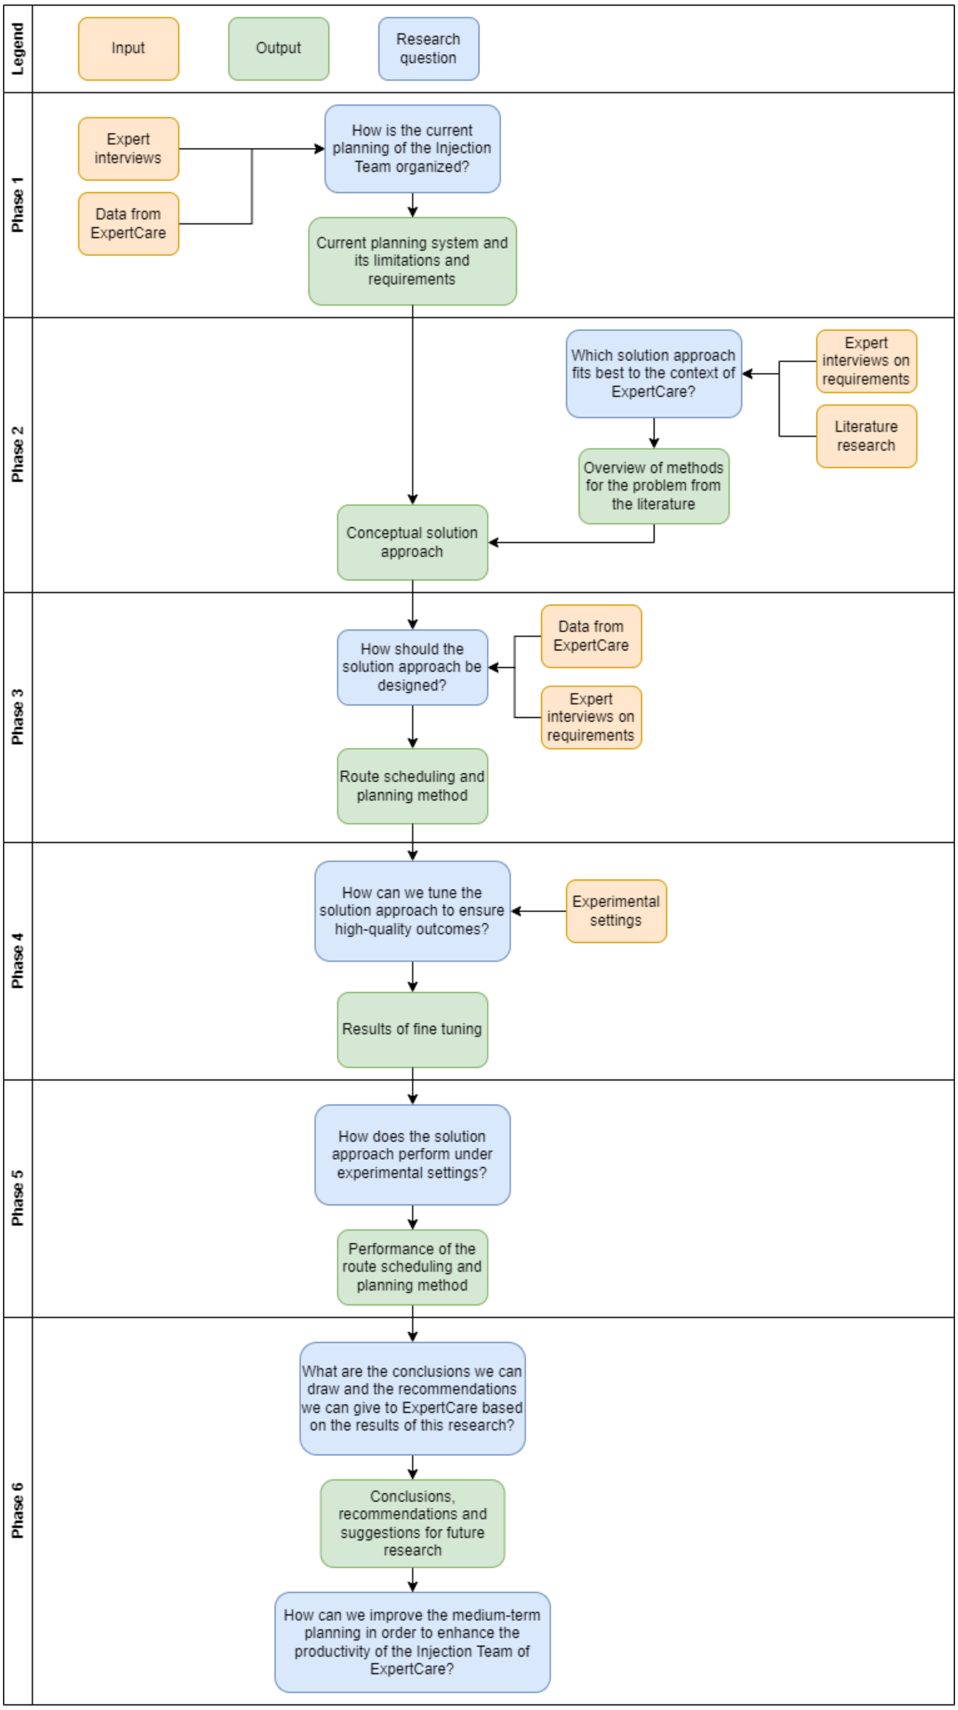

Figure 26: Research design

B Measurements travel time

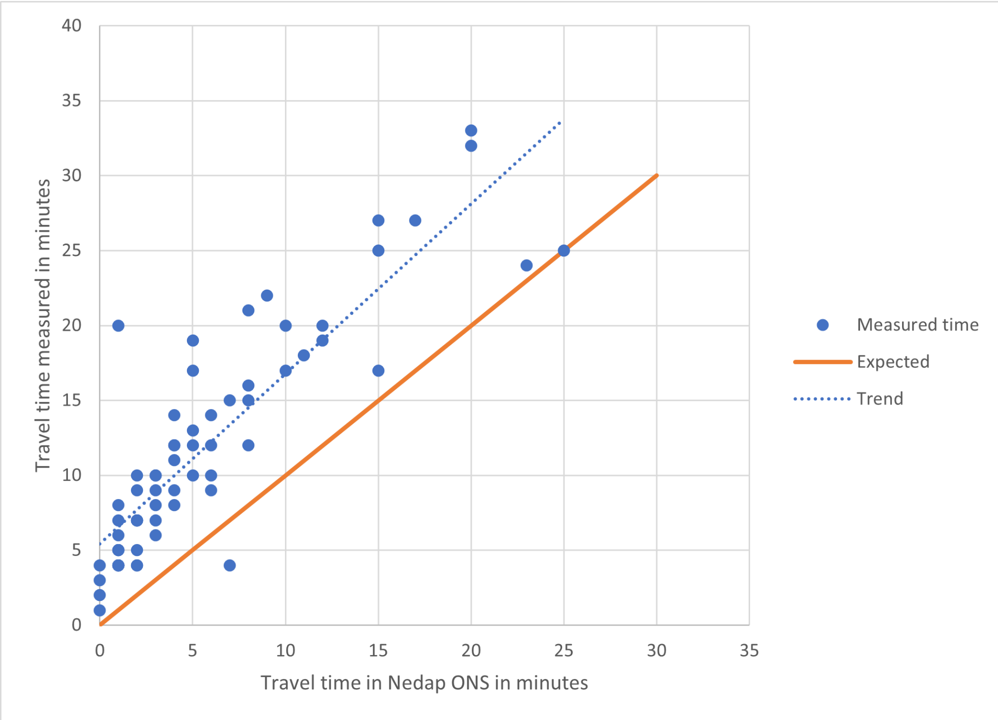

Figure 27: Travel time measurements graph

Table 9: Travel time measurements

Time in Nedap	Measured time	Extra time
10	20	10
3	9	6
2	9	7
12	20	8
1	5	4
5	12	7
6	12	6
4	11	7
0	2	2
8	12	4
20	33	13
17	27	10
2	4	2
2	7	5
8	16	8
11	18	7
1	4	3
2	7	5
15	25	10
1	6	5
1	4	3

Continued on next page

Table 9 – continued from previous page

Time in Nedap	Measured time	Extra time
1	4	3
2	10	8
2	4	2
15	27	12
4	12	8
0	4	4
10	17	7
1	5	4
8	21	13
2	7	5
1	8	7
3	10	7
6	10	4
12	19	7
1	20	19
4	14	10
5	17	12
8	15	7
3	8	5
9	22	13
4	12	8
0	1	1
4	9	5
7	15	8
4	8	4
1	4	3
5	10	5
5	19	14
6	14	8
3	10	7
15	17	2
7	4	-3
25	25	0
3	6	3
5	13	8
2	5	3
6	9	3
0	3	3
20	32	12
1	5	4
23	24	1
3	7	4
2	7	5
1	7	6

C Cluster analysis

Table 10: Cluster analysis current schedule of the Injection Team

Cluster	**Number of clients**	**Average number of visits in half year**
Monday Morning 1	54	160.54
Monday Afternoon 1	51	141.77
Monday Morning 2	19	74.69
Monday Afternoon 2	72	154.92
Tuesday Morning 1	30	117.46
Tuesday Afternoon 2	28	84.15
Tuesday Morning 2	55	126.38
Tuesday Afternoon 2	18	81.15
Wednesday Morning 1	47	93.15
Wednesday Afternoon 1	73	176.08
Wednesday Morning 2	19	76.85
Wednesday Afternoon 2	62	166.00
Thursday Morning 1	65	155.46
Thursday Afternoon 1	38	127.31
Thursday Morning 2	50	105.85
Thursday Afternoon 2	19	56.23
Friday Morning 1	65	203.31
Friday Afternoon 1	30	81.54
Friday Morning 2	56	163.69
Friday Afternoon 2	28	120.62

D Home Health Care Routing and Scheduling literature overview

Table 11: Literature overview of HHCRSP

Author(s)	Horizon	Obj	Constraints	PB	Stochasticity	Frequency	Coc	Solved by:
Bard et al. [2]	Week	Min TC	-TW -P -CTW -Worktime -Overtime -Skill -Break -Holiday	No	No	Once per week	No	MILP
Bard et al. [1]	Week	Min TC	-HTW -Overtime -Skill -Break	Yes	No	Multiple times per week	No	MILP - GRASP
Bowers et al. [4]	30 days simulated	Min TT	-P -Coc -WB	Yes	Arrival of new patients included (by Poison distribution)	4 times per 10-15 days (for 10-15 days)	Yes	Clarke & Wright - Sim
Cappanera et al. [7]	Week	Min TT BW	-HTW -Skill -Coc	Yes	No	Multiple times per week	Yes	Sim - MILP
Guericke [16]	At least a week	Min TT Min WTN	-HTW -Worktime -Overtime -Skill -Break -Holiday	No	No	-	Yes	MILP - ALNS
Fathollahi-Fard et al. [12]	1 day to 3 weeks	Min TC Max PS	-HTW	No	Uncertainty in time windows, service time, travel time and patient's preferences	-		MILP
Hewitt et al. [18]	2-3 months	Min S Min TD	-Coc -Worktime	No	Arrival of new patients included (by simple point estimate)	2-3 times per week (for 60-90 days)	Yes	ConVRP Record to Record
Lin et al. [24]	4 weeks with RH	BW	-PP -Overtime -Skill -Coc -STW	No	No	1 time per week 1 time per 2 weeks 1 timer per 4 weeks	Yes	MILP

Continued on the next page.

Author(s)	Horizon	Obj	Constraints	PB	Stochasticity	Frequency	Coc	Solved by:
Maya Duque et al. [25]	Several weeks	Max PS Min TD	-STW -PP -Coc -CTW -Worktime -Skill	Yes	No	1-5 times per week	Yes	2-stage solution approach
Nickel et al. [29]	Week	Min TD OC UT	-HTW -CTW -Holidays -Skill -BW	Yes	Arrival of new patients included	Multiple times per week	Yes	ALNS - CP
Nikzad et al. [30]	Week	Min S Min TC	-Skills -HTW -Worktime -Coc	No	Uncertain travel and service times	Multiple times per week	Yes	2-stage solution approach
Shao et al. [35]	Week	Min TC BW	-HTW -Skill	Yes	No	Multiple times per week	No	GRASP
Wirnitzer et al. [38]	4 weeks	Max coc	-P -Coc -Worktime -Overtime -Break -Holidays -Skill	No	No	1-21 times per week	Yes	MIP
Yalçındağ et al. [39]	Week	Min TC BW	-Skill -Coc -Worktime	Yes	No	Multiple times per week	Yes	2-phase solution approach

Obj = Objective(s), TC = travel cost, TT = travel time, BW = balancing workload, WTN = waiting time of nurses, PS = patients' satisfaction, S = staff, TD = travel distance, OC = overtime cost, UT = unscheduled tasks, Coc = continuity of care, TW = time window, P = preferences, CTW = care workers' TW, HTW = hard TW, STW = soft TW, PP = patients' preferences, PB = pattern-based, MILP = mixed integer linear program, GRASP = greedy randomized adaptive search procedure, Sim = simulation, ALNS = adaptive large neighborhood search, CP = constraint programming, MIP = mixed interger program

E Pseudocode greedy heuristic

Algorithm 3a Greedy heuristic: minimize max workload

```
Sort clients by frequency
for each client i do
    for each feasible scheme s in S_i do
        Score = 0
        for each time slot t in s do
            Score = Score + number of appointments already in t
        end for
        Normalize Score with number of appointments in s
        Save Score to scheme s
    end for
    Randomly select a scheme among the ones with lowest Score
end for
```

F Results Adaptive Large Neighborhood Search performance experiment 1

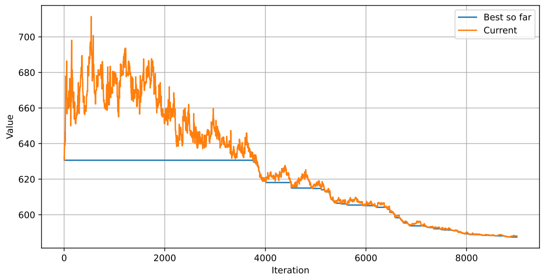

FIGURE 28: Performance of the ALNS for the second repetition of experiment 1

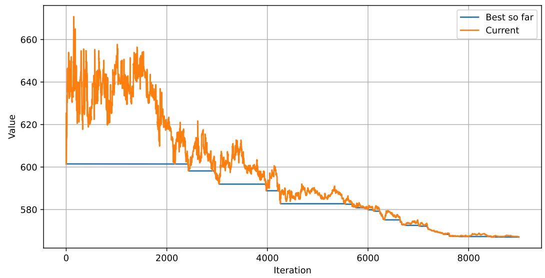

FIGURE 29: Performance of the ALNS for the third repetition of experiment 1

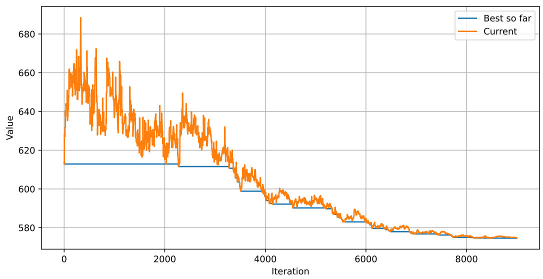

FIGURE 30: Performance of the ALNS for the fourth repetition of experiment 1

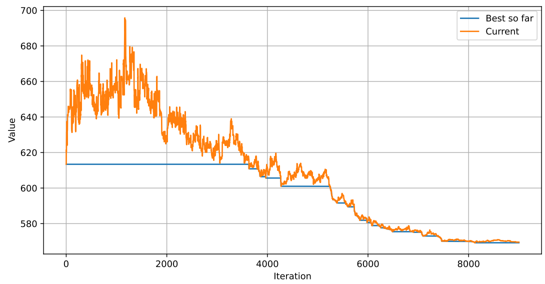

FIGURE 31: Performance of the ALNS for the fifth repetition of experiment 1

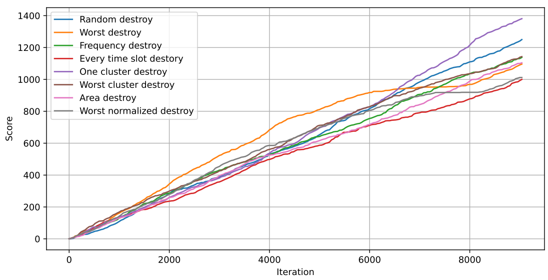

FIGURE 32: Performance of the destroy operators for the second repetition of experiment 1

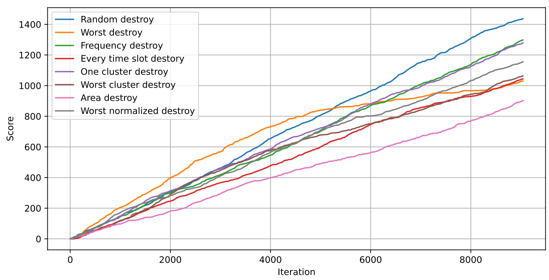

FIGURE 33: Performance of the destroy operators for the third repetition of experiment 1

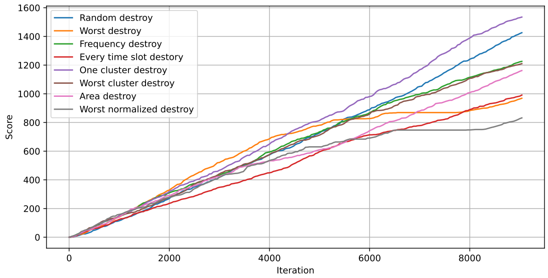

Figure 34: Performance of the destroy operators for the fourth repetition of experiment 1

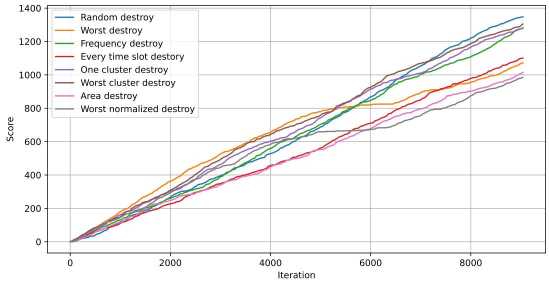

Figure 35: Performance of the destroy operators for the fifth repetition of experiment 1

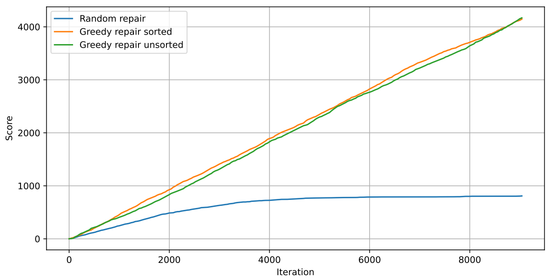

Figure 36: Performance of the repair operators for the second repetition of experiment 1

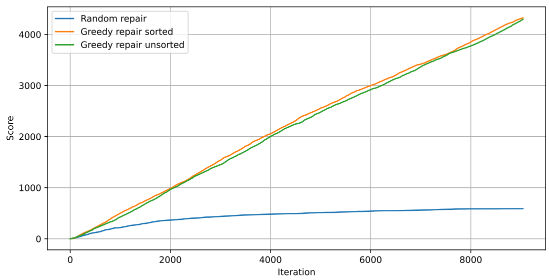

Figure 37: Performance of the repair operators for the third repetition of experiment 1

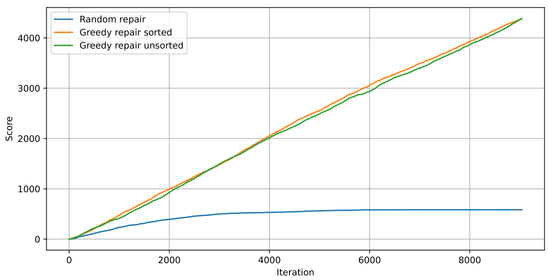

Figure 38: Performance of the repair operators for the fourth repetition of experiment 1

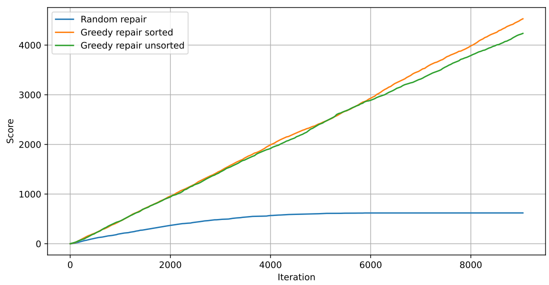

Figure 39: Performance of the repair operators for the fifth repetition of experiment 1

G Results experiment 1

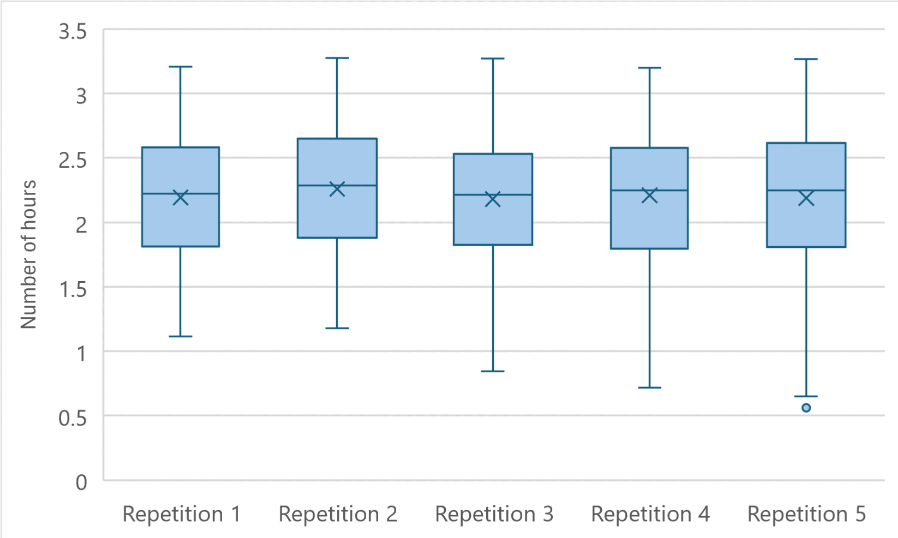

FIGURE 40: The average traveling hours per nurse per day for all repetitions of experiment 1

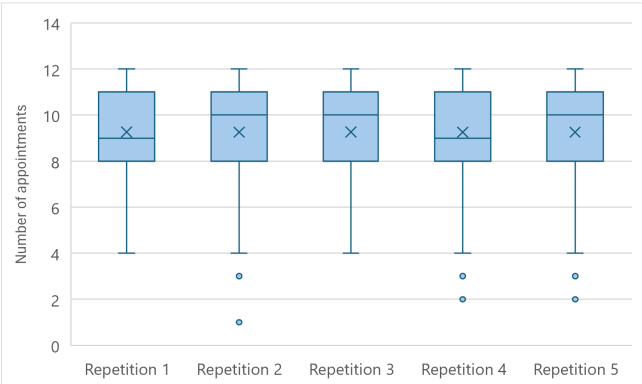

FIGURE 41: The average number of appointments per nurse per day for all repetitions of experiment 1

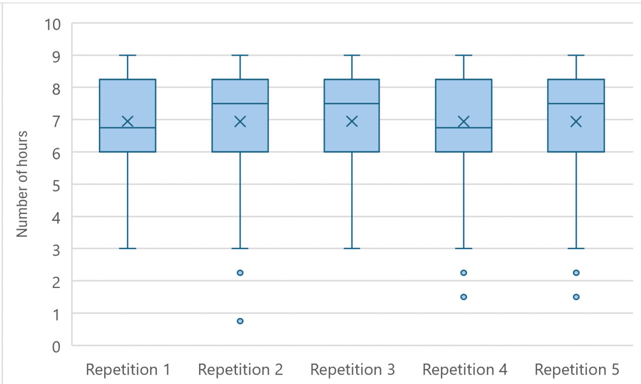

FIGURE 42: The average number of nursing hours per nurse per day for all repetitions of experiment 1

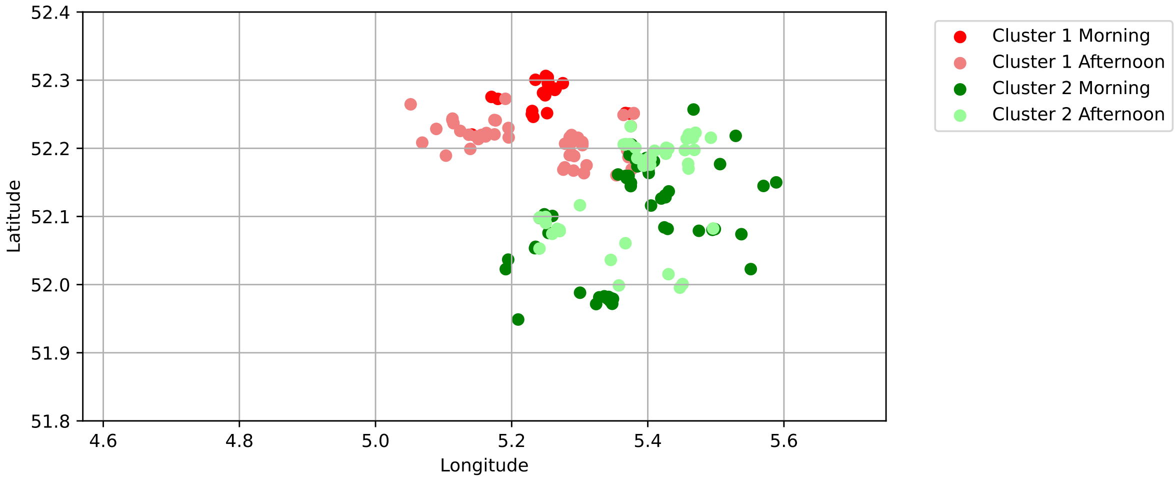

FIGURE 43: 2 clusters for a repetition of experiment 1

TABLE 12: Number of clients per frequency per cluster

	Cluster 1 Morning	**Cluster 1 Afternoon**	**Cluster 2 Morning**	**Cluster 2 Afternoon**
Frequency 1	1	0	0	0
Frequency 2	0	0	0	0
Frequency 3	1	0	0	0
Frequency 4	10	9	6	11
Frequency 6	1	1	0	0
Frequency 8	0	0	0	2
Frequency 12	9	36	33	14
Frequency 26	0	4	15	17

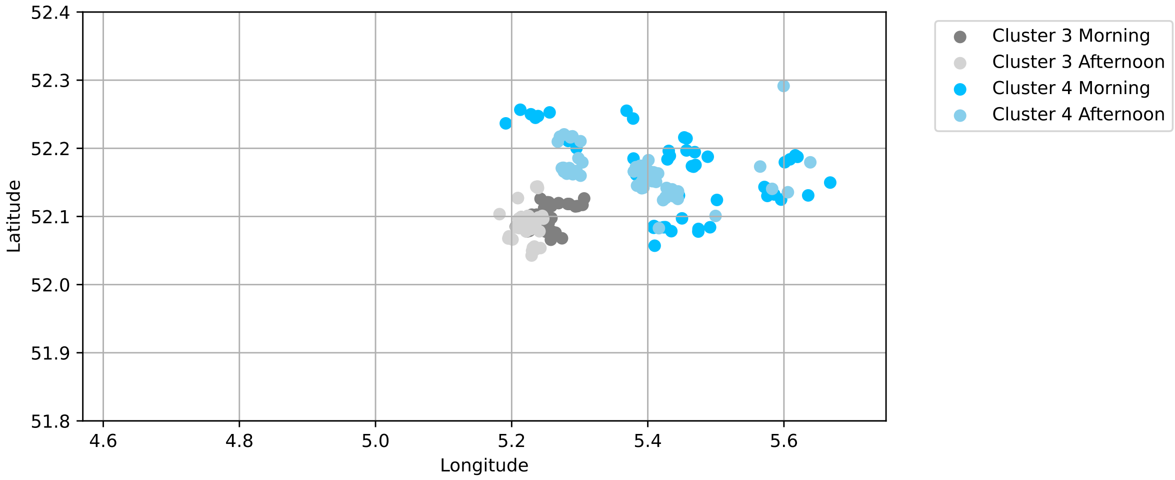

FIGURE 44: 2 clusters for a repetition of experiment 1

TABLE 13: Number of clients per starting week per cluster 3 and 4

	Cluster 3 Morning	**Cluster 3 Afternoon**	**Cluster 4 Morning**	**Cluster 4 Afternoon**
Frequency 1	1	0	0	0
Frequency 2	0	0	0	0
Frequency 3	4	0	0	0
Frequency 4	6	10	5	4
Frequency 6	0	0	0	0
Frequency 8	0	0	0	0
Frequency 12	7	4	35	44
Frequency 26	33	34	8	9

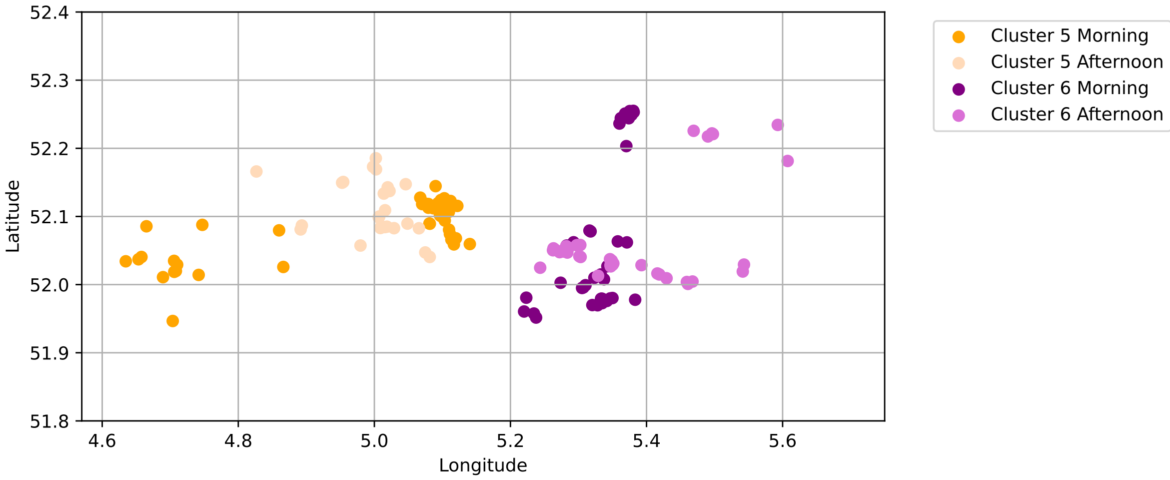

FIGURE 45: 2 clusters for a repetition of experiment 1

TABLE 14: Number of clients per starting week per cluster 5 and 6

	Cluster 5 Morning	**Cluster 5 Afternoon**	**Cluster 6 Morning**	**Cluster 6 Afternoon**
Frequency 1	0	1	0	0
Frequency 2	1	0	0	1
Frequency 3	1	0	0	0
Frequency 4	4	9	10	7
Frequency 6	0	0	0	0
Frequency 8	2	0	0	0
Frequency 12	20	7	9	9
Frequency 26	21	8	24	22

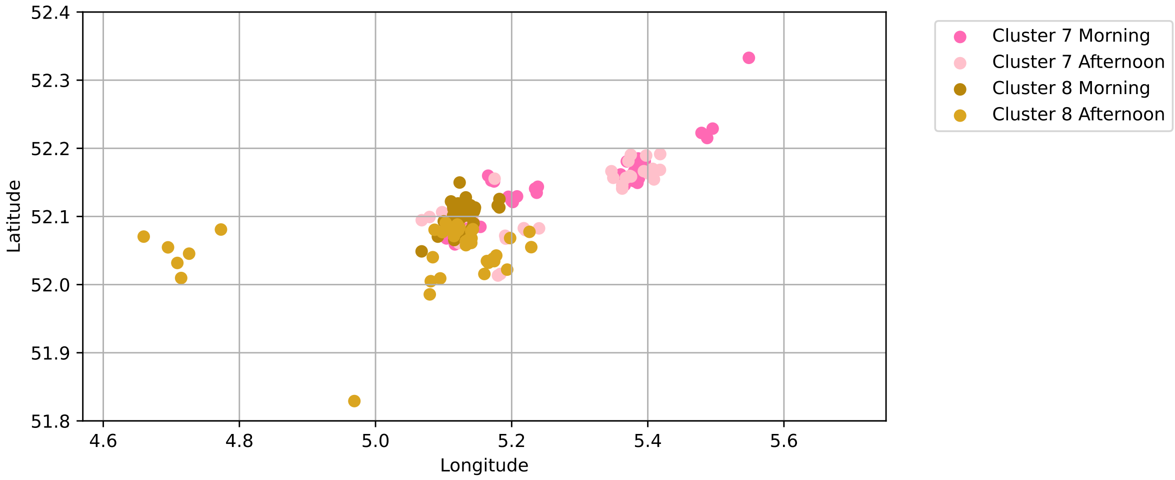

FIGURE 46: 2 clusters for a repetition of experiment 1

TABLE 15: Number of clients per starting week per cluster 7 and 8

	Cluster 7 Morning	**Cluster 7 Afternoon**	**Cluster 8 Morning**	**Cluster 8 Afternoon**
Frequency 1	0	0	1	0
Frequency 2	0	2	2	0
Frequency 3	0	0	0	1
Frequency 4	7	7	9	6
Frequency 6	0	0	0	0
Frequency 8	1	1	1	0
Frequency 12	26	10	7	19
Frequency 26	28	26	15	15

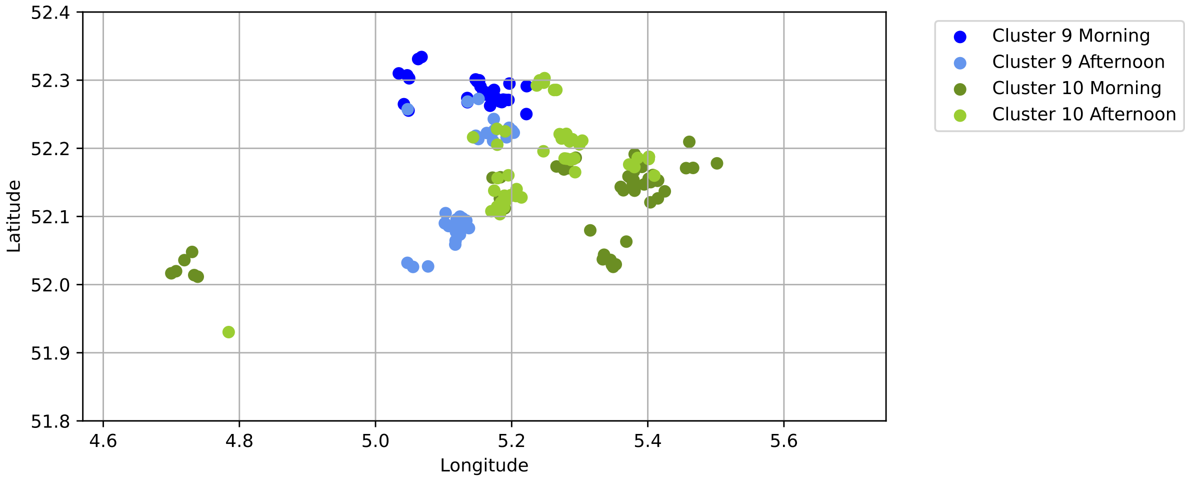

FIGURE 47: 2 clusters for a repetition of experiment 1

TABLE 16: Number of clients per starting week per cluster 9 and 10

	Cluster 9 Morning	**Cluster 9 Afternoon**	**Cluster 10 Morning**	**Cluster 10 Afternoon**
Frequency 1	0	0	0	0
Frequency 2	2	0	0	1
Frequency 3	0	0	1	0
Frequency 4	6	13	2	7
Frequency 6	0	0	1	0
Frequency 8	0	0	0	0
Frequency 12	21	15	31	28
Frequency 26	1	9	16	11

H Results other experiments

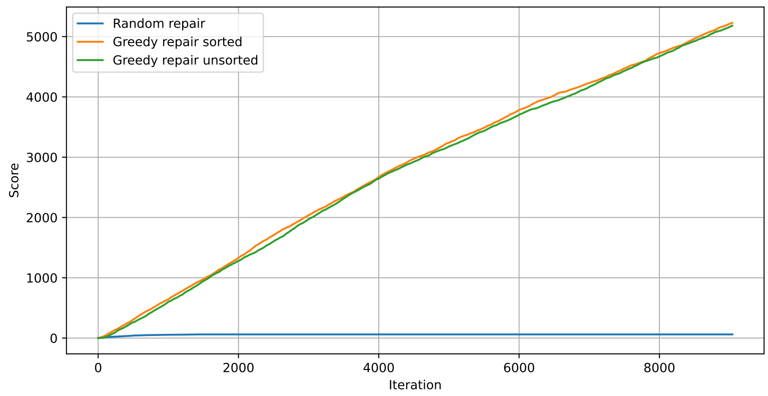

FIGURE 48: Performance of the repair operator for the first repetition of experiment 3

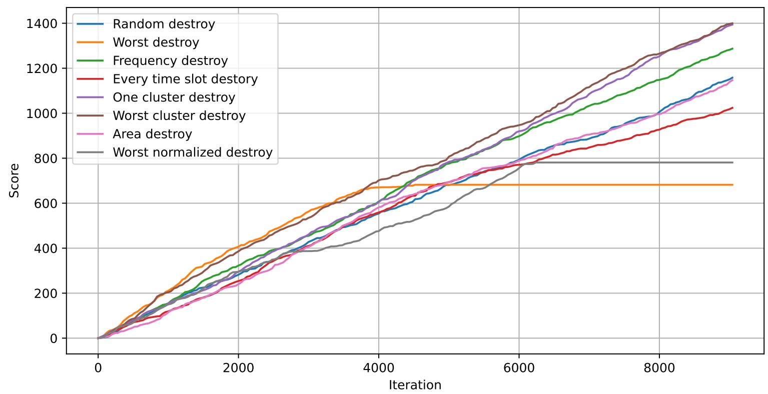

FIGURE 49: Performance of the destroy operator for the first repetition of experiment 4

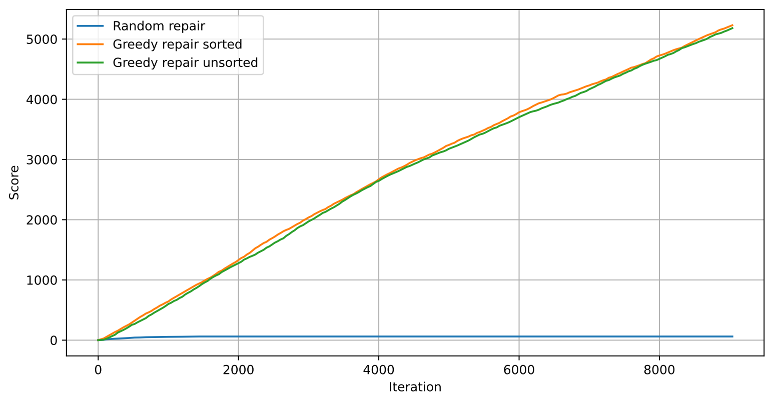

FIGURE 50: Performance of the repair operator for the first repetition of experiment 6

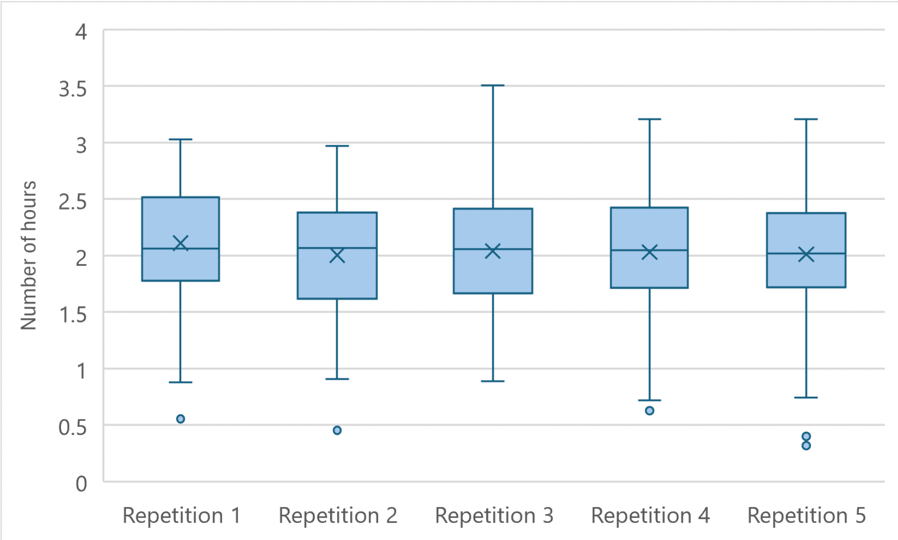

FIGURE 51: The average traveling hours per nurse per day for all repetitions of experiment 2

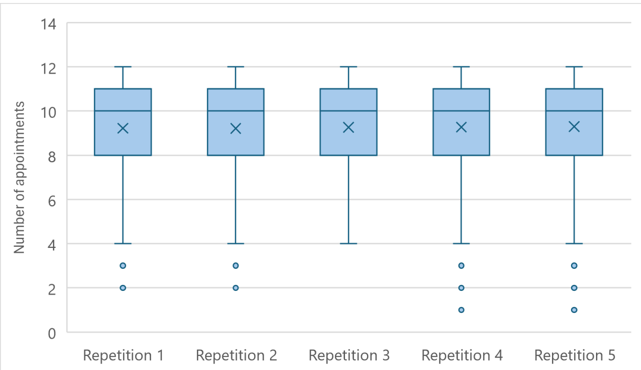

FIGURE 52: The average number of appointments per nurse per day for all repetitions of experiment 2

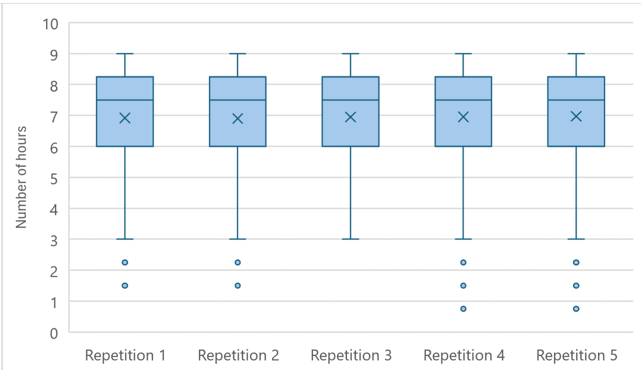

FIGURE 53: The average number of nursing hours per nurse per day for all repetitions of experiment 2

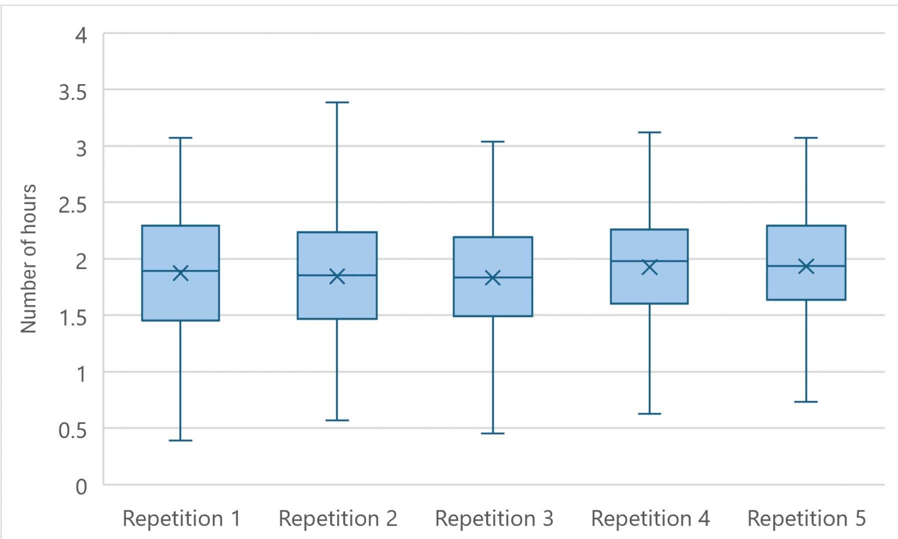

FIGURE 54: The average traveling hours per nurse per day for all repetitions of experiment 3

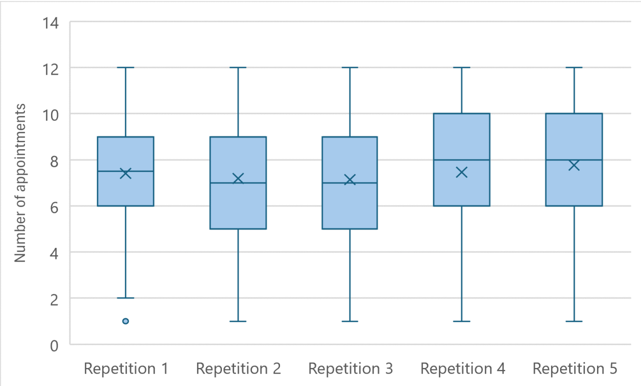

FIGURE 55: The average number of appointments per nurse per day for all repetitions of experiment 3

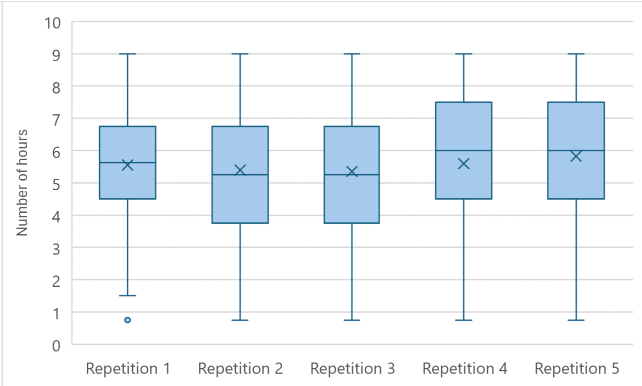

FIGURE 56: The average number of nursing hours per nurse per day for all repetitions of experiment 3

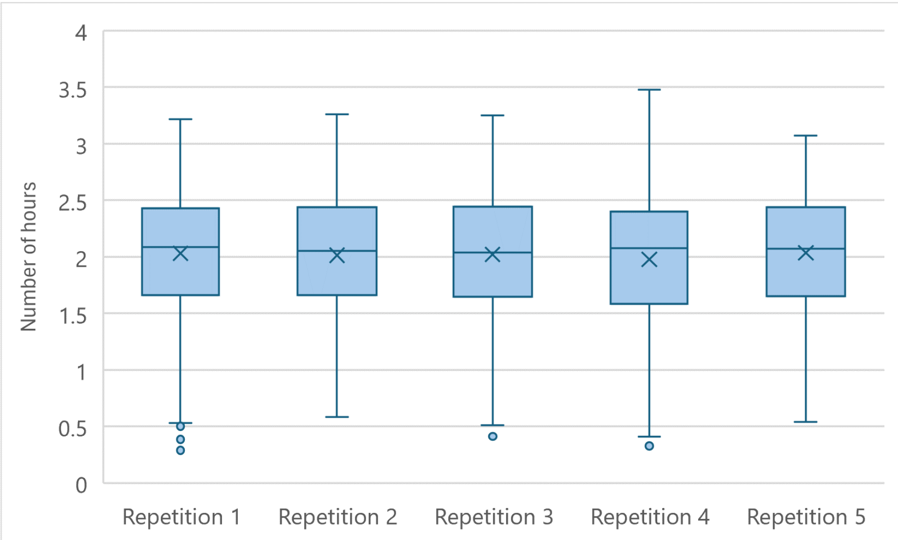

FIGURE 57: The average traveling hours per nurse per day for all repetitions of experiment 4

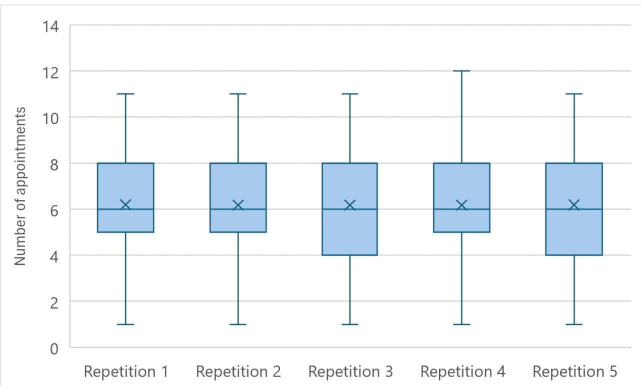

FIGURE 58: The average number of appointments per nurse per day for all repetitions of experiment 4

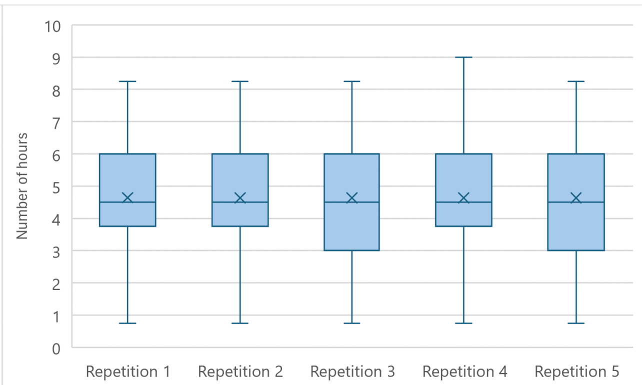

FIGURE 59: The average number of nursing hours per nurse per day for all repetitions of experiment 4

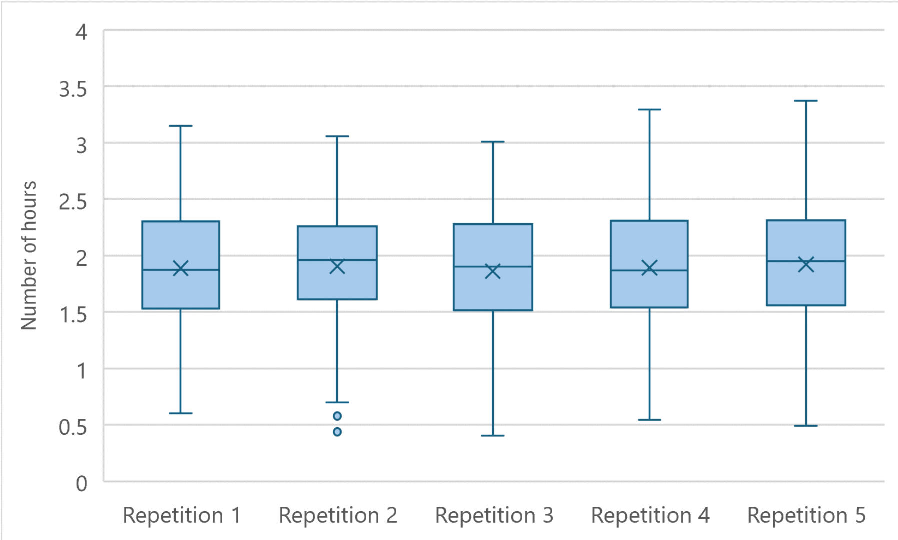

FIGURE 60: The average traveling hours per nurse per day for all repetitions of experiment 5

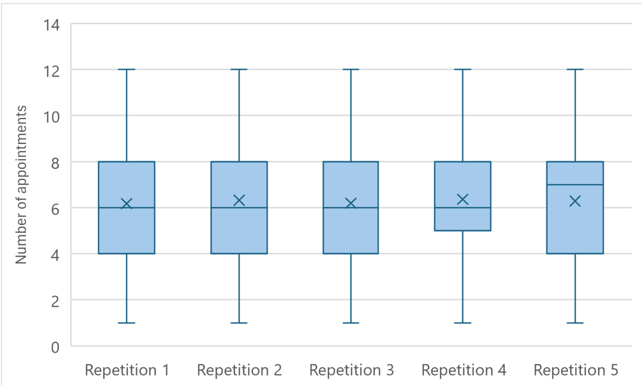

FIGURE 61: The average number of appointments per nurse per day for all repetitions of experiment 5

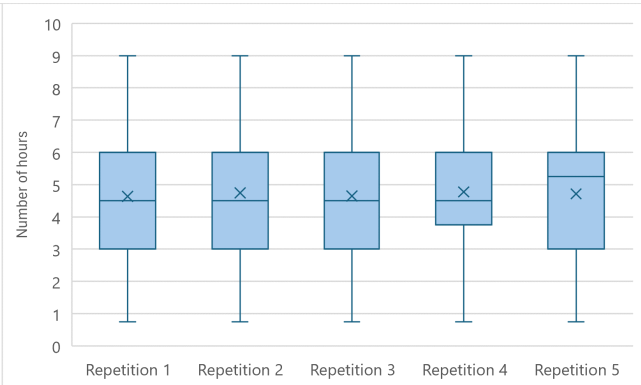

FIGURE 62: The average number of nursing hours per nurse per day for all repetitions of experiment 5

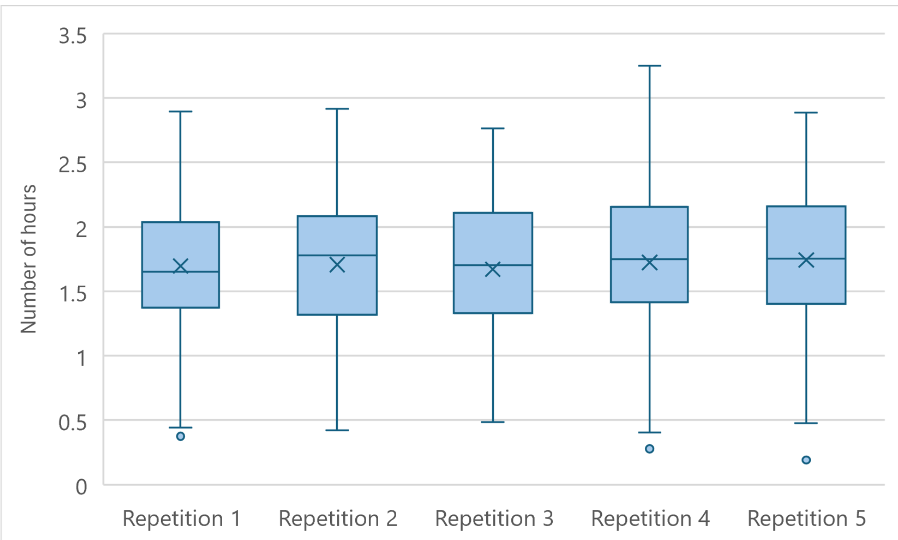

FIGURE 63: The average traveling hours per nurse per day for all repetitions of experiment 6

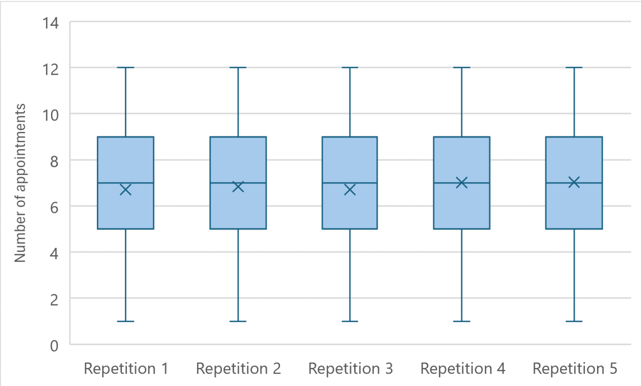

FIGURE 64: The average number of appointments per nurse per day for all repetitions of experiment 6

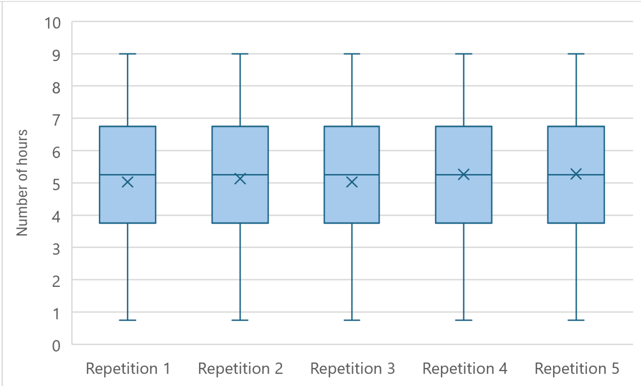

FIGURE 65: The average number of nursing hours per nurse per day for all repetitions of experiment 6

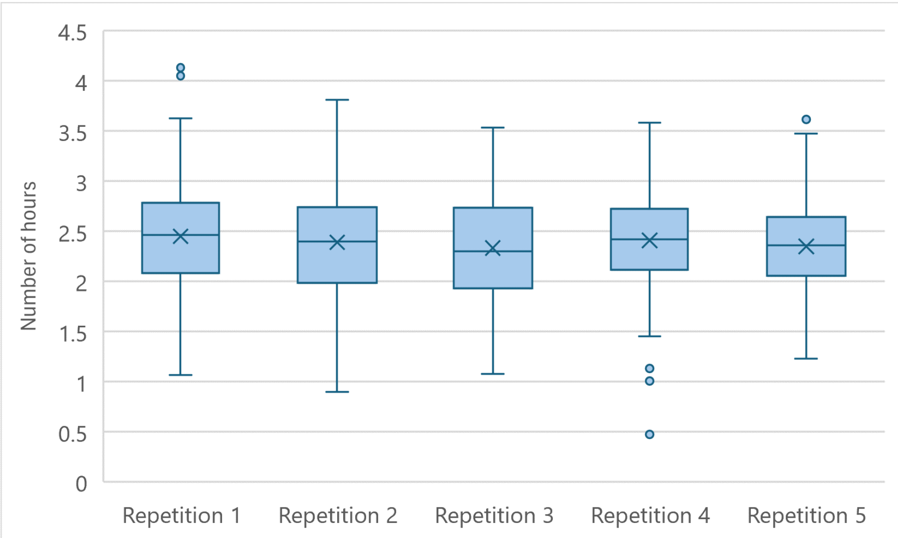

FIGURE 66: The average traveling hours per nurse per day for all repetitions of experiment 7

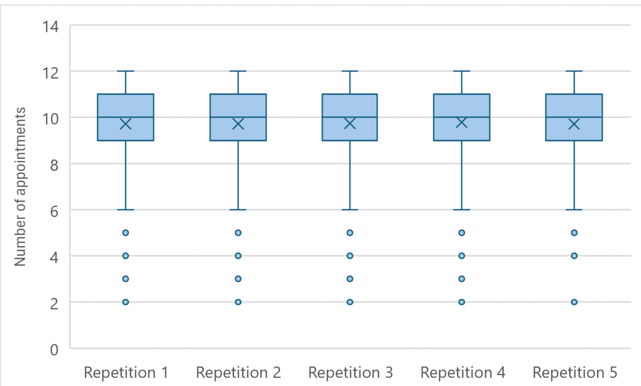

FIGURE 67: The average number of appointments per nurse per day for all repetitions of experiment 7

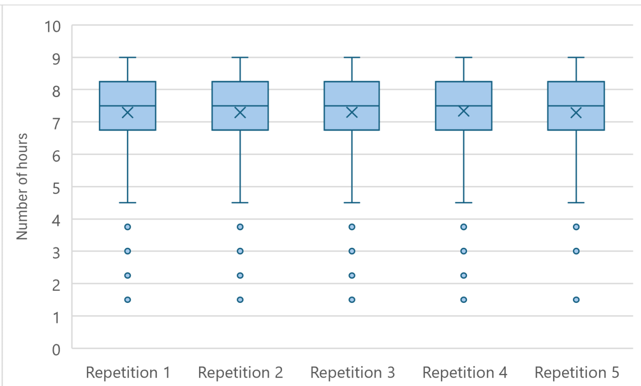

FIGURE 68: The average number of nursing hours per nurse per day for all repetitions of experiment 7