MSc. Thesis
Industrial Engineering & Management

Tactical allocation of resources in an outpatient clinic to treat patients close to home and improve room planning

Luuk Klein Nagelvoort

Supervisors UT:
S. Rachuba &
A.G. Leeftink

Supervisors Isala:
A.E. Roelofs &
J. Doornebal

Department of Behavioural,
Management, and Social Science

December 19, 2024

UNIVERSITY OF TWENTE.

Preface

With this thesis, I will conclude my Master Industrial Engineering and Management at the University of Twente. The months leading up to this have been a huge learning experience for me, and have definitely kept me busy. I have written something that I am genuinely proud of, and for that I have many people to thank.

Thank you Ton Roelofs for your enthusiasm, and patience, throughout my thesis. I always felt that I was supported, and that any question I asked would receive consideration and immediate attention. You provided me with enough freedom to express my curiosity, while steering me back on track whenever I got too lost in the details. For this I am very grateful. Thank you as well to Joan Doornebal. The enthusiasm you have for your work was genuinely inspiring, and on many different occasions rubbed off on me. This was particularly the case around the midpoint of my thesis, when I got slightly stuck on the data-analysis.

Thank you as well to Sebastian Rachuba for your valuable insights and guidance. Our talks repeatedly provided me with renewed confidence and motivation whenever I got stuck, which I really did need at certain moments. Thank you as well to Gréanne Leeftink for your help and extensive feedback. I believe your advice surrounding the literature and structure of my thesis have had substantial effects on the eventual quality, and for this I am grateful. I also want to thank Erwin Hans, although he may not read this, for convincing me to enroll in this Master.

Lastly, but definitely not least, I want to thank all my family and friends, and in particular my housemates, for their support. I have to thank Kim specifically, for the daily coffee breaks.

To anyone reading, I hope you may enjoy this thesis as much as I have enjoyed working on it,

Luuk Klein Nagelvoort,

Enschede, December 2024

Management Summary

Delivering care as close to patients' homes as possible is pivotal in improving accessibility of care, which is a national priority for the Netherlands. This accessibility can be facilitated through effective employee capacity allocation policies, which subsequently requires an effective capacity planning methodology to ensure feasibility. With our research we show how hospital appointment data can be used to inform decision-making on employee allocation. We also show how mathematical optimization through Simulated Annealing (SA) can be used to ensure hospital planning goals.

We demonstrate the effectiveness within the internal medicine department of the Isala hospital cluster. We use department appointment data from 2023 to retroactively determine a capacity allocation that would have ensured that all patients would have been treated as close to home as possible. The analysis uses the google maps distance matrix API to determine the closest location for each previous appointment. Given this information, the department can determine the potential workload that could have been scheduled in each location on a proximity basis.

These calculations can be used to inform decisions on future allocations, given the stable nature of patient demand in the department. They can also be used to determine current performance. We show, for instance, that the department treated 54.7% of their patients as close to home as possible in 2023. Continued analysis of the data using our methodology will allow the department to evaluate how many patients are being treated close to home on an ongoing basis. The department can then incorporate this information as an additional performance indicator, in addition to scheduling- and resource constraints.

The previous information allows the department to determine where capacity should be available given the formulated objective, but does not directly tell them how to schedule the capacity throughout the week. We show how planning objectives and constraints can be ensured through mathematical optimization, using Simulated Annealing, a computational optimization. We demonstrate the effectiveness of this technique specifically by improving planning performance in Zwolle. We show that the approach can reduce the unnecessary use of undesirable, windowless rooms by 100% in two evaluated sub-departments. The planning can be created while complying with employee preferences and unavailability. Given its performance in Zwolle, we suggest that the approach can be expanded straightforwardly to include specific location-time constraints.

These results demonstrate the effectiveness of this approach in a hospital environment. The optimization ensures current planning objectives, while providing a holistic framework for incorporating new objectives into the planning. Given the developed method, the department can ascertain all their objectives if they have exact knowledge of firstly all employee preferences and obligations, and secondly an overview of the consultations that need to be scheduled in each location. For the second part, the department can refer back to the data analysis that showcases the potential volume

of patient care in each location. Using this analysis in coordination with employees will allow them to determine which employees would be available for satellite consultations, and generate a list of the consultations to be scheduled. The department can then coordinate availability of rooms with each location.

To underline the utility of these approaches within the internal medicine, we develop a baseline by benchmarking current performance, and indicating other potential improvements. We show that under current conditions, the department treats 54.7% of patients in the closest location. Changes to the scheduling approach could increase this to, at most, 80.9% under perfect conditions. If instead, the department changes their capacity allocation to reflect the calculations from our data analysis, they can expect to treat 95.3% of patients in the satellite locations if each timeslot contains two hours and thirty minutes of physical patient appointments. The above numbers show that a data-driven approach to the employee capacity allocation can ensure that this planning goal is achieved, and that changes to the allocation follow directly.

To show the utility of the planning approach, we demonstrated that the current room schedule for general internal medicine schedules an average of 36.2 timeslots of consultations in rooms without windows. For hematology and oncology this is 5.2 timeslots. We show that decreasing the total number of timeslots by increasing the number of appointments per timeslot has slight effects on the number timeslots scheduled in windowless rooms for GIM, but does not reduce them to zero. For hematology and oncology, the number of windowless slots decreases more substantially with longer consultations. We show that using simulated annealing can reduce the average number of windowless consultations to 0.17 and 1.71 respectively through improved planning, with the remaining dayparts caused directly by limited room capacity. It can also ensure that switching between rooms during the week is minimized, and that sharing of rooms by multiple employees is minimized. Additionally, the approach can incorporate employee availability constraints.

To conclude. We recommend that the internal medicine department includes additional dimensions in their database. The department should be able to distinguish between appointments that do require a consultation room, and appointments that do not. To replicate our analysis further, the data should be able to distinguish between the desired distribution of DBC's across specialties. Lastly, the closest (- or desired) location for each appointment should be included. Given these dimensions, the department can analyze the planned work in each location versus the potential work in each location. Given this information, we recommend that the department investigates creating a centralized document which contains employee availability and preferences, such that the planning can always be evaluated on this basis. The department should also concretely map out information on room availability in the satellite locations. Using this information the department can construct a new planning approach based either on our MILP formulation or heuristics code. Either of these approaches would be able to create a new planning for the department that can simultaneously ensure selection of specific rooms, employee- and room availability, and improve accessibility for patients. The only requirement is the correct information.

Contents

Acronyms & Abbreviations

IZA: Integral Care Agreement (Integraal Zorg Akkoord)
NP: New patient
CP: Control patient
TC: Telephonic consultation
NPTC: New patient telephonic consultation
NS: Nurse specialist (Verpleegkundig specialist)
PA: Physician Assistant
DBC: Diagnosis Treatment Combination (Diagnose-Behandelcombinatie)
HiX: Healthcare Information eXchange (system)
the department: the Internal Medicine Department
GIM: the General Internal Medicine sub-department
H&O: the Hematology and Oncology sub-department

Dutch terms

Algemeen interne: General internal medicine specialty
Endo/Endocrinologie: Endocrinology
Infec/Infectiologie: Infectious disease
Nephro/Nefro/Nefrologie: Nephrology
Vascu/Vascular/Vaculaire Geneeskunde: Vascular Medicine
Elderly/Ouderengeneeskunde: Elderly Care
Hemat/Hematologie: Hematology
Onco/Oncologie: Oncology
Aandachtsgebied: Specialty

1 Introduction

Isala Hospital aims to develop a new capacity allocation for its Internal Medicine department to improve the expression of national strategic healthcare goals. The objectives are specifically to treat patients as close to home as possible, and to ensure adequate workspaces for Isala's employees. We first introduce Isala's Internal Medicine department in Section 1.1. As the problem owner, they requested and managed the execution of this research. We continue with the motivation for our research in Section 1.2. In Section 1.3, we elaborate on the specific challenges which the Internal Medicine department faces, and summarize these in a problem cluster. We formulate the research questions in Section 1.4, and describe the action problems and associated evaluation criteria in Section 1.5.

1.1 Company description

In the region of Zwolle, the Netherlands, the Isala hospital cluster is responsible for 675.000 people across twenty different municipalities. Although the hospital provides both clinical and outpatient care across a range of departments, the Internal Medicine department (from now on: the department) focuses mainly on outpatient care. The department consists of two sub-departments: general internal medicine (GIM), and oncology/hematology (H&O). H&O covers the specialties of *oncology* and *hematology*, and GIM covers the specialties *acute care, endocrinology, infectious disease, nephrology, elderly care*, and *vascular medicine*. In total, the department is made up of eight specialties.

The department employs over a hundred healthcare professionals, supported by a range of other staff members. It operates across multiple locations: the main hospital in Zwolle, a smaller hospital in Meppel, and satellite locations in Kampen, Heerde, and Steenwijk. Although the department performs ancillary treatments, such as dialysis and chemotherapy, most of the activities involve supportive care procedures. These procedures consist mainly of diagnostic appointments for new patients, and monitoring for patients with chronic conditions. Since most of these activities are basic care services without the need for specialized rooms, most of the care is administered in shared consultation rooms.

1.2 Research motivation

In the Netherlands, the aging population is driving a rising demand for healthcare services that outpaces the modest growth in healthcare personnel (Centraal Bureau voor de Statistiek, 2022). With one in six individuals already employed in healthcare, further expansion is neither feasible nor desirable (VWS, 2022). The Ministry of VWS (Public Health, Welfare and Sports) has constructed an agreement that outlines a set of national goals for healthcare in the Netherlands (VWS, 2022). This Integral Care Agreement (Integraal Zorg Akkoord - IZA) was constructed in cooperation with most of the major hospital organizations, the organization for mental healthcare (GGZ), and multiple elderly care providers. Among the objectives are increasing accessibility of care for patients, as well as ensuring employees' physical and mental health (VWS, 2022).

Literature shows that longer travel distances reduce accessibility of care. Mseke et al., (2024), for instance, illustrate that even in countries with well-developed infrastructure, longer distances correlate with increased no-shows and cancellations. Kelly et al., (2016) reviewed 108 studies relating travel distances to health outcomes for patients, and found that these deteriorated with increased travel distance in 77% of the evaluated studies. Reducing travel times may therefore not only improve access to care, but also directly positively affect patient health. Isala own vision also directly aligns with the idea of minimizing travel: "Close if possible, far away if necessary." Achieving this goal, however, requires a structured approach to match employee capacity with patient demand across the locations. This had led to the following research question:

"*How can the Internal Medicine department ensure that patients are treated as close to home as possible?*"

Equally important for Isala is to ensure adequate workspaces for its employees, as these have been linked to higher engagement and reduced stress (Teoh et al., (2021)). Daylight exposure, in particular, has been shown to correlate with reduced stress, decreased medical errors, improved professional performance (Ulrich, 2008), as well as more well-regulated circadian rhythms (Ulrich, 2008; Huisman et al., 2012). Studies like Alimoglu et al., (2005) highlight the role of daylight exposure in improving job satisfaction and lowering stress in nurses in particular, while Boubekri et al., (2014) found higher job satisfaction, lower stress-levels and reduced negative health outcomes on SF-36 metrics (Ware et al., 1996) in office workers.

These ideas have already informed the design of the Isala building in Zwolle, where consultation rooms were designed to maximize daylight exposure (Burger, 2013). Nevertheless, healthcare professionals are still unable to avoid the use of windowless rooms under current capacity conditions, which undermines these benefits. This had lead to the following research question:

"*What strategies can be used by the Internal Medicine department to reduce the use of windowless rooms in Zwolle?*"

1.3 Problem context

In this section we describe the current situation in the department, in order to to narrow down what the actual core problems are.

The main reason is a lack of information. The department does not know how many patients could be treated in each location, and can therefore not base its employee capacity directly on patient demand. Instead, the capacity at each location has been historically determined. Decisions about when and where a doctor would work have often been based on prior immediate needs and personal doctor preferences. Commonly, this decision was made at the time of employment and never reevaluated. Since then, vacated capacity in a certain location has frequently been straightforwardly replaced by equivalent capacity, without considering potential modifications. As a result, the schedule is primarily based on historical precedents.

Given optimal capacity use, an allocation that would align this capacity with the demand in each location would have many benefits. The first is alignment with national strategic goals according to the IZA, as well as alignment with Isala's own vision. The resulting reductions in travel time would free up more time for patients, and could, therefore, have a substantial impact especially for those patients that have to be treated on a regular basis. Reaching the hospital in Zwolle also requires travelling through dense traffic. Zwolle, therefore, is less accessible for patients who are unable to handle this type of traffic, such as elderly patients who have trouble driving. Facilitating treatment at different locations might, therefore, be particularly beneficial for those specific patients. Lower travel distances would also reduce national fuel usage, contributing to sustainable policy.

There is another possible benefit, which ties into the second problem. In treating more patients at the satellite locations, we could also treat fewer patients in the main locations of Zwolle and Meppel. This could alleviate capacity issues in those locations, decreasing pressure on the planning process. Capacity is limited, and employees' external obligations and personal preferences make scheduling a very difficult task in Zwolle in particular. This is further complicated by the fact that the planning is made manually using Excel, meaning that staff members may spend days puzzling together a feasible schedule every time a change is required.

This problem is aggravated further by the lack of a clear policy to determine whether or not an employee is granted a room for a consultation. The department operates with set consultation *slot* durations, but has no policy that governs how many patients should be treated during that time. This means that whenever a room is occupied, that room is occupied for the entire time slot, even if the employee only treats a small number of patients. This causes issues when doing so was not strictly necessary. In this scenario, other employees who could have made better use of the room capacity will not be able to do so. Consequently, windowed rooms can not always be found for employees, resulting in either the use of windowless rooms, or the postponement of appointments.

Given these issues and our proposed research questions, we arrive at the following research goal.

"*To develop a standardized methodology for allocating capacity across Isala's Internal Medicine locations that ensures that patients are treated close to home, while improving room planning to reduce consultations in windowless rooms.*"

The problem landscape, as described above, has been visualized in Figure 1A. This figure shows the problems and their mutual relationships in a cluster. The green and black problems represent the core problems (Heerkens et al., (2021)), meaning that these problems have no direct cause. We will be solving the green core problems directly, while the black problem is out of scope. The orange problems represent the action problems (Heerkens et al., (2021)), which are those problems that can be observed on the surface. The blue problems are used to describe the connection between the core and action problems.

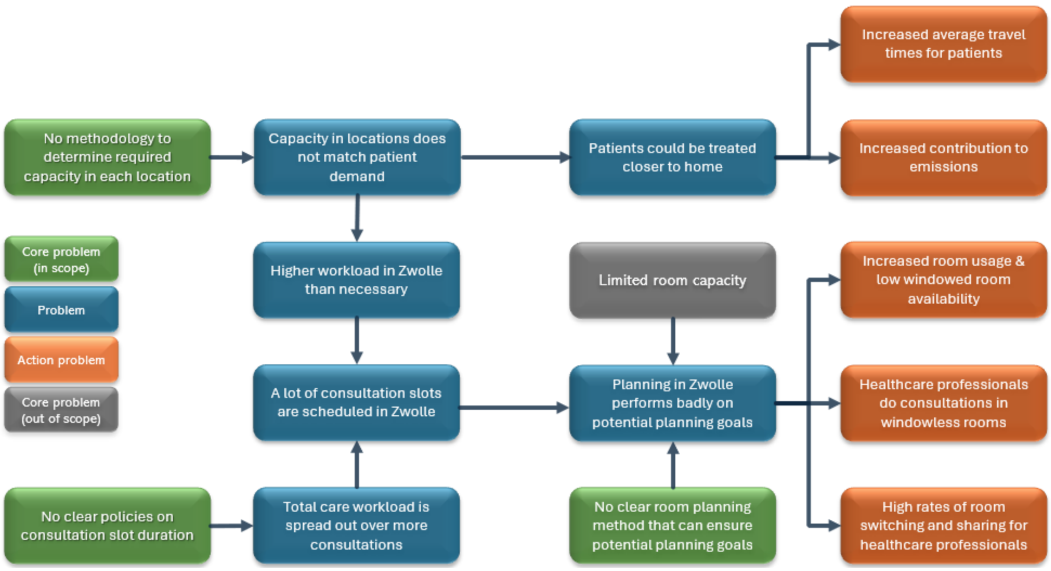

Figure 1A: Problem cluster

1.4 Research questions

We present two separate research questions. The first question covers the core problem "no clear methodology to determine required capacity in each location". The second research question covers the other two core problems within Zwolle specifically: "planning is done manually using Excel", and "no clear policies on consultation slot duration."

(I) "*How can the Internal Medicine department ensure that patients are treated as close to home as possible?*"

(II) "*What strategies can be used by the Internal Medicine department to reduce the use of windowless rooms in Zwolle?*"

1) Context analysis

(a) *What are the available resources?*

(b) *What are the characteristics of the current room schedules?*

2) Literature review

(a) *How can we determine the distance from the patient to each hospital location?*

(b) *What planning problems exist in literature with a similar structure?*

(c) *What methods would be effective in creating a new room planning?*

3) Data analysis

(a) *What data is available for the analysis?*

(b) *Which parts of the data are relevant?*

(c) *Which employees, specialties, patients, and appointments should be considered in the analysis?*

(d) *What are the assumptions that we make in the data analysis?*

4) Solution design

(a) *How can we determine patient demand for each location?*

(b) *How can we determine the number of consultations required per specialty and profession?*

(c) *How can we decrease the total number of slots used in the room schedules?*

(d) *What planning policies and methods should we choose?*

5) Results analysis

(I) Performance analysis of the location allocation

(a) *What is the performance of the current location capacity allocation?*

(b) *How many consultations can be scheduled per specialty and profession in each location?*

(c) *What is the performance of the new location capacity allocation?*

(d) *How does the reallocation across the locations affect the room schedule in Zwolle?*

(II) Performance analysis of the room schedule

(a) *What is the performance of the current planning?*

(b) *What is the performance of decreasing the total number of consultations?*

(c) *What is the performance of new planning approaches?*

(d) *What is the performance of policies on the duration of consultations combined with a new planning approach?*

6) Conclusions and recommendations

(I) Conclusions for the location allocation

(a) *How can the department improve the performance of the current allocation?*

(b) *How could the department change the current capacity allocation?*

(c) *What would be the effects on the number of patients treated close to home?*

(d) *What changes do we recommend to the department?*

(II) Conclusions room planning

(a) *How do we recommend that the department decreases the complexity of its planning?*

(b) *What strategies can the department implement to improve its planning approach?*

7) Discussion and future research

(a) *What are the limitations of the analysis?*

(b) *What assumptions did we make that should be considered in the decision-making?*

(c) *How can the results be more broadly applied?*

(d) *What expansions do we recommend for future research?*

1.5 Evaluation criteria

To measure the performance on our action problems, we use six evaluation criteria. The action problems and associated criteria are shown in Table 1a below.

Action problem	**Evaluation criteria**
Increased travel times for patients	*Percentage of patients treated close to home*
Increased contribution to emissions	CO_2 *emission reductions*
Increased room usage & low room availability	*Total number of consultations planned in Zwolle.*
Healthcare professionals do consultations in windowless rooms	*Total number of consultation slots in windowless rooms.*
High rates of room switching for healthcare professionals	*Number of times healthcare professionals have to switch rooms.*
High rates of room sharing for healthcare professionals	*Number of rooms used by more than one healthcare professional.*

Table 1a: Evaluation criteria for Action Problems

2 Context analysis

Before we can design solutions to the problems identified in the previous section, we require a better understanding of the department's resources. In this chapter we analyze both those resources (Section 2.1) and the current planning methods (Sectio n2.2) used by the internal medicine department.

2.1 Available resources

In this section, we explain the different resource dimensions within the department. In order, we describe the locations, rooms, specialties, and employees. We then describe the different types of appointments and how the department deals with holidays.

2.1.1 Locations

The department operates across five locations, with its main hospital in Zwolle and a smaller hospital in Meppel. The three surrounding locations—Kampen, Heerde, and Steenwijk—serve as satellite facilities, each designed to increase capacity, expand patient access within the region, and reduce travel time for appointments. A map showing the locations can be seen in Figure 2A. This map also shows hospitals outside of the Isala cluster (shown in red). As patients reside farther away from Zwolle, their choice of hospital naturally expand to also include those locations.

Although the majority of the appointments are handled in Zwolle due to its higher capacity, most care services can be carried out across all of the Isala locations, with only a few exceptions. We describe all the exceptions in detail in Appendix A.3.

Other differences between the locations revolve around the practicality for patients. Firstly, Zwolle is a city, meaning that there is increased traffic. Secondly, parking is paid, and thirdly the larger location means that it takes longer to get to the consultation room from the car.

The main difference between the locations, however, is the subjective experience of patients and employees, which is related to the atmosphere at each location. Although preferences vary, employees report myriad individual examples of patients preferring the calmer atmosphere in the satellite locations. The same holds for the employees themselves, although there are also some reasons why they instead might prefer Zwolle. Examples are that satellite locations are often farther away for employees, and that the main location motivates collegiality and ease of access to information and meetings.

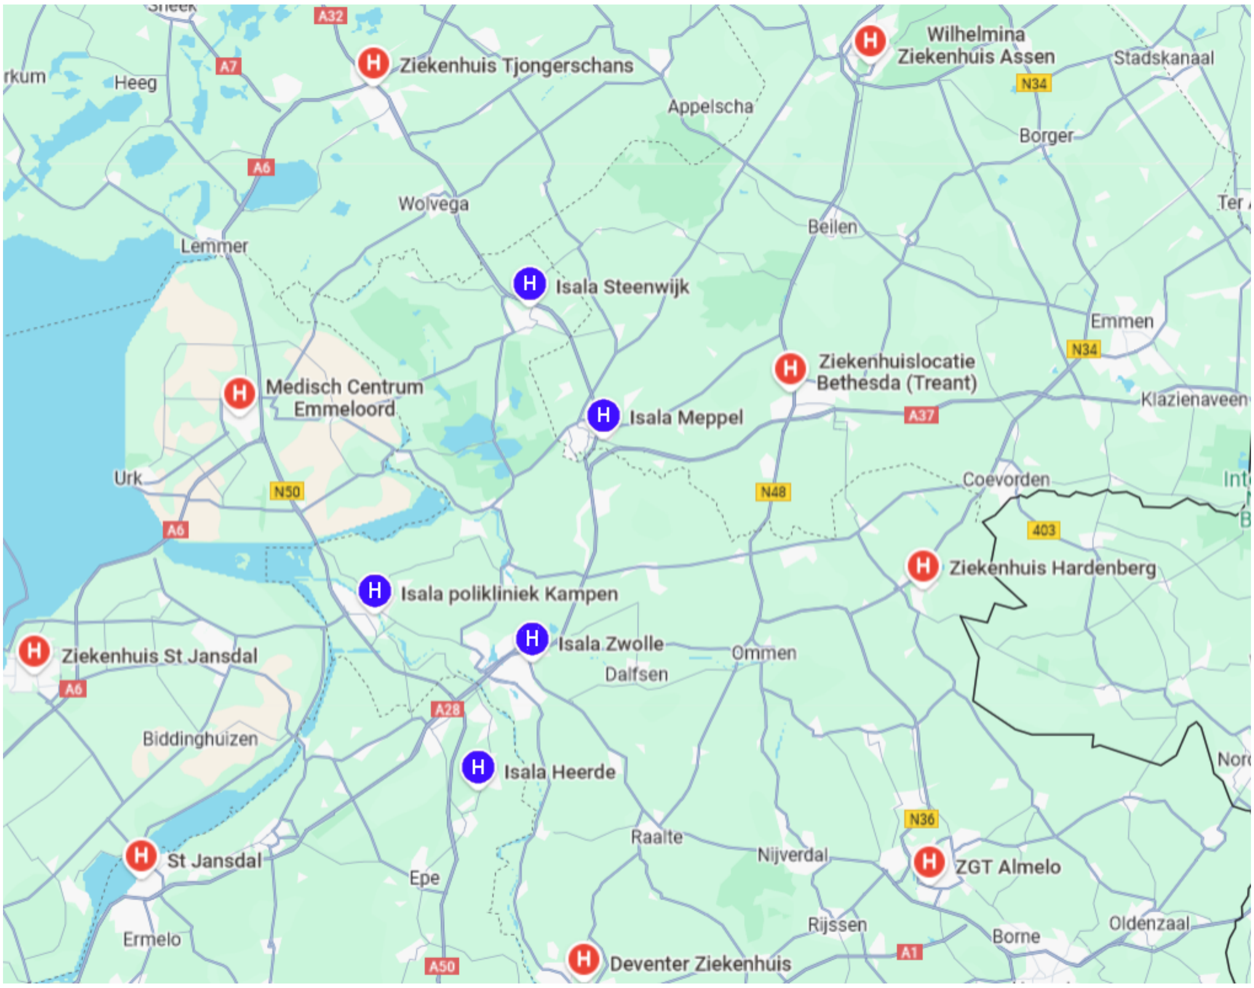

Figure 2A: Map showing Isala locations and surrounding hospitals

2.1.2 Rooms

During the design of the hospital in Zwolle, a pivotal goal was to facilitate as much daylight as possible within consultation rooms. To this end, the architects incorporated butterfly structures into their design (Figure 2B - Burger, 2013). These structures allow exposure to daylight on both sides of the building, with the inner side of the building looking out on a garden. The internal medicine department uses the top half of 'vlinder 2' (butterfly 2). Within this space, 34 rooms with daylight exposure are available to the department. Of these, eleven are used by H&O, and 23 are used by GIM. Three additional windowed rooms are used by the planning bureau of GIM, the secretariat of GIM, and Innofeet, an external party. Inside the building, in between the windowed rooms, we also find additional windowless rooms used as working rooms, storage rooms, and several diagnostics/treatment facilities. Due to a perceived lack of capacity, the department has started using some of the windowless working rooms for their consultations as well. Eleven of these windowless rooms are typically considered in the schedule. An overview of the distribution of rooms can be found in Figure 2C. This room diagram only considers the rooms that are explicitly included within the room schedules of the sub-departments, and thus excludes a range of windowless rooms used for other purposes.

Figure 2B: Design of the department work space

The reader may have noticed that the rooms are grouped under the specialties. This is due to the fact that every specialty has its own front desk, such that specialties mostly use rooms close to their own front desk. To still use capacity as efficiently as possible, the specialties share their room capacity dynamically when necessary. This happens mostly within the sub-departments, but can also happen between sub-departments. The barrier for doing this between sub-departments is larger, however, since the room planning is made by two separate parties.

The reader may also have noticed that the diagram shows the number of rooms multiplied by a factor of ten. This due to the fact that the department works with a set consultation slot duration of a half-day, one in the morning and one in the afternoon. Although evening slots are possible, they are not a common occurrence and therefore not a reliable resource. This means that a total of ten of these *slots* are then typically available per room in each week. In Zwolle, the rooms used for internal medicine are available exclusively to the department, and each room therefore contributes exactly ten *owned* slots. In the other locations, this is not the case, and it therefore makes more sense to express their capacity in terms of the number of slots than a fractional number of rooms. In order to be consistent, we express the capacity in Zwolle in slots as well.

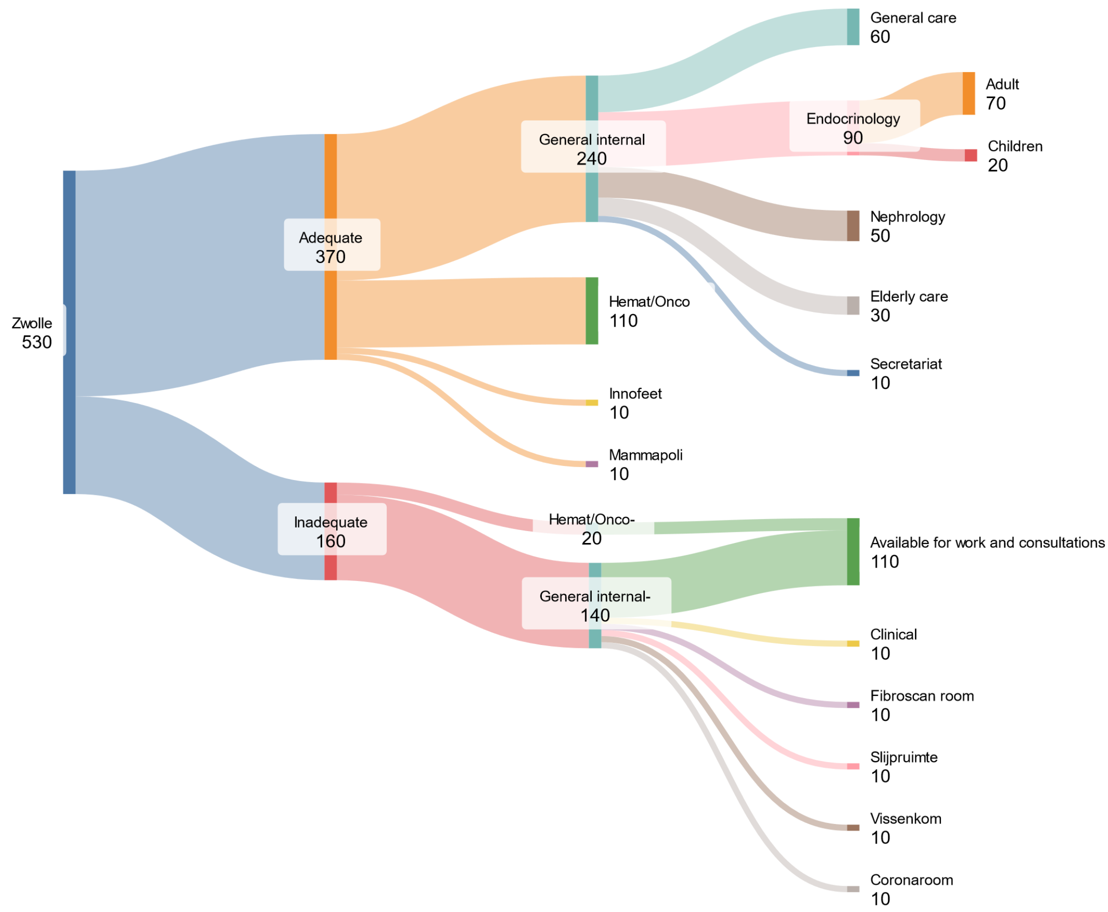

Figure 2C: Division of windowed (adequate) and windowless (inadequate) room slots within Zwolle

For the satellite locations, the capacity is as shown in Table 2a. We differentiate between slots already in use by internal medicine and slots that are still available. A number of slots in the satellite locations have not been dedicated to any department in the standard weekly schedule, meaning that the department could request use of those additional slots. The use and availability of slots varies between even and odd weeks, as can be seen in the table. We exclude slots specifically used for oncology nurses in Kampen and Heerde (4&2 respectively). The telephonic slots in Meppel occur in windowless rooms.

Location	**Meppel**		**Kampen**		**Heerde**		**Steenwijk**	
Week type	**Even**	**Odd**	**Even**	**Odd**	**Even**	**Odd**	**Even**	**Odd**
Owned slots	80	80	22	22	10	13	9	9
Free slots	-	-	5	7	12	11	29	34
(Telephonic slots)	(15)	(15)	-	-	-	-	-	-
Total windowed	**80**	**80**	**27**	**29**	**22**	**24**	**38**	**43**

Table 2a: Available consultation slots by location and week type

2.1.3 Specialties

Internal medicine is a broad department, consisting of eight specialties across two sub-departments. Most employees only see patients from their own specialty, and it is therefore important to determine the capacity allocation on a per-specialty basis. An exception to this rule exists for doctors from the specialties within the GIM sub-department, as these doctors can handle appointments for most of the simpler cases of the other specialties in their sub-department. Some of the appointments can be handled by all six specialties, which we then consider *general* appointments grouped under a new specialty 'Algemeen Interne'. Other appointments can be handled by a subset of specialties, and others are specific to a certain specialty.

A previous study by van der Wel, (2023) has determined this for the most relevant diagnoses within the department. They analyzed the most frequent diagnoses that cumulatively account for 90% of all appointments within the Isala GIM sub-department, and determined how these were distributed across the specialties. Based on their findings, the department has established a target distribution of the diagnoses for each GIM specialty. Although this target does not match the current distribution, the hospital aims to align them in the future. Because of this, we present all the workload per specialty on the basis of the diagnoses that each specialty is expected to handle. Notably, about 30% of the work managed by GIM doctors falls within a general pool of appointments that can be handled by any GIM doctor. Employees that are not doctors —as well as doctors from the H&O sub-department— only draw from their own individual specialties, aside from a handful of cases.

2.1.4 Employees

A number of distinct professions are required to handle the various demands of the care process in a hospital. Of those professions, the doctors, nurse specialists, and physician's assistants are mainly involved in the *diagnosis* of patients. Nurses are in charge of the day-to-day care and support of patients. Doctors' assistants mainly provide supportive care and assist with administrative work. They may draw blood, or do an initial check-up. Co-assistants are medical students in training, requiring direct doctor supervision. In most specialties, the majority of the work is done by doctors. There are a few exceptions, however. Endocrinology, for instance, relies heavily on their nurses and nurse specialists, with less than 20% of their total work done by doctors.

2.1.5 Appointments

The specialties all use an expansive variety of appointment codes to qualify the type of appointment. Although this information may be relevant for the department, we are only interested in the classification of the appointment into two different categories. The first category is whether the appointment handles a new- or returning (control) patient, the second category is whether the appointment was physical or telephonic. For this we use four different abbreviations: CP (control patient), NP (new patient), TC (telephonic appointment), NPTC (new patient telephonic appointment). While new patients can effectively be allocated to any of the employees of the appropriate profession, returning patients generally stick with the employee that handled their first appointment, unless specifically requested.

2.1.6 Reduction weeks

One last thing that we have to consider in our analysis is capacity reduction due to holidays. This happens most during the national holidays, which the hospital has aptly named 'reduction-weeks'. Employees work 44 weeks a year, with eight weeks of vacation, and tend to align their vacations with the reduction-weeks. The department does not want to lose all their capacity at once, however, and therefore makes sure that vacations are relatively spread out over different weeks within specialties. Additionally, employees may choose to go on holiday outside of the designated periods. As a result, the capacity floats slightly above the average capacity (105% of the average) outside of reduction-weeks, and scales down to on average approximately 84% of the average during the actual reduction-weeks. The reduction-weeks in 2023 fell on weeks 1, 9, 17-18, 30-35, 43, and 52. We show this in Table 2b as well.

Holiday / Reduction-Week	Week Number
Kerstvakantie 2022-2023	1
Voorjaarsvakantie	9
Meivakantie	17-18
Zomervakantie	30-35
Herfstvakantie	43
Kerstvakantie 2023-2024	52

Table 2b: Dutch holidays

2.2 Current planning methods

The department tracks its scheduling activities using a program called HiX, an interface used to visualize the appointments of the employees. The individual calendars of employees are available to them within this system, but a *complete overview* of all consultation slots only exists within an Excel sheet used by the planning department. This same Excel sheet is used to allocate the rooms to the employees, with a screenshot sent to them by mail on a weekly basis.

The underlying planning procedure within the department happens on two levels, and starts with the creation of a tactical blueprint in Excel. This blueprint is shown as a grid, wherein employees are allocated a slot for half-days during the week. The planning department is aware of most of the employee preferences and puzzles the blueprint together manually, but the working preferences of employees are incomplete and spread over a large number of documents on the department drive. The current blueprint has existed in its approximate form for a long time, and minor adjustments are only made when employee availability changes. This includes change of preferences, employees leaving, and new hires.

These grids as mentioned earlier are filled operationally with appointments. Since the planning does not change much, the planning horizon of patient appointments may extend even until a year into the future. Depending on the appointment type, the scheduling is managed by different staff members. New patient appointments (NP) for all locations and specialties are managed by the planning bureau, while follow-ups (CP) are scheduled by the secretariats at each respective location and specialty. When employees are unavailable for part of a session, they can block off that part of a consultation slot within the HiX program. Most of the time, this is not incorporated directly into the main Excel schedule.

The actual room schedules are made separately from the HiX program, and closer to the consultation dates. This is a pragmatic decision due to the experienced capacity issues. In doing so, employee absence will be mostly known at the time of room allocation and the planning department can puzzle the schedule together more easily.

As for the capacity allocation across the locations, there is no concurrent knowledge that describes how consultation slots at the satellite locations were distributed amongst employees in the past. Regardless, the department currently only aims to maintain the current allocation; when an employee leaves, the same type of employee will be asked to replace the newly unfilled slots.

2.3 Chapter conclusion

In this chapter we described the available rooms, specialties, employees, and appointments. We also explained reduction weeks and the current planning methods. Now that we have elaborated on the practical information, we continue with our literature review. This will inform the theoretical part of our research, and enable the design of solutions to our established problems.

3 Literature review

In this chapter we evaluate relevant literature for the design of our solution approaches. For the location allocation, we first investigate whether any programs exist that can automatically calculate the distance between two ZIP codes for a large database (Section 3.1). On the contrary, no literature is required to determine how to reduce the number of consultations, since this is directly achieved by increasing the average duration of consultations. For the planning approach, we require a more extensive study. We aim to identify whether there are any relevant studies that have solved similar planning problems. We describe this in Section 3.2.

3.1 Location capacity allocation

We will explore the relevant literature for both problems. From the available data, we can already determine the workload that has been scheduled in each location. Before we can determine the demand in each location, however, we first need to develop a method for identifying the closest location for each appointment. Given this demand, we can immediately determine the capacity that would have been required. The literature review of this section is therefore aimed at finding a program that can be used to determine the distance between the patient- and hospital locations' ZIP codes.

Determining distance

Pgeocode (Symerio, 2023) is a Python library designed for geocoding and straight-linedistance calculation based on postal codes. It provides functionality for querying postal code data to extract location information, such as latitude and longitude, and for calculating the geodesic distance between two postal codes. The library is optimized for performance and includes support for datasets covering numerous countries, including the Netherlands. The syntax is simple, and the algorithms are quick and efficient (Symerio, 2023).

The Google Maps Distance Matrix API (Google, 2024) is a web service that calculates travel distances and times between multiple origins and destinations based on real-world routing data. It supports various transportation modes, such as driving, walking, and public transit. The API is commonly used in logistics, delivery services, and travel planning to optimize routes and schedules. The API is robust and scalable, but does require web querying. This likely makes it slower than pgeocode. Although pgeocode is easier to implement, the Google Maps Distance Matrix API (GM-API) provides real-world routing data, providing a more accurate estimation of the patient travel distance. Use of the API might be complicated due to the need to query the web for distance calculations, which can be computationally costly. After performing a feasibility study on a small number of ZIP codes, we decided to use the Google Maps Distance Matrix API .

3.2 Capacity planning for rooms

The problem we are attempting to study involves allocating employees to consultation rooms with a fixed slot duration. Slots have this fixed duration regardless of the room that they are placed in, and employees have a fixed number of slots that they need to complete. The number of rooms is a constrained resource in this problem, and we are aiming for schedule feasibility while minimizing the use of inadequate resources (i.e. windowless rooms).

Cardoen et al., (2010) reviewed the complexities of particularly the need to match staff schedules with available room time in an operating room context. Similarly, Hulshof et al., (2012) addressed the allocation of healthcare resources under limited availability, offering insights relevant to minimizing the use of inadequate rooms while ensuring sufficient staffing. Both studies highlight the difficulty in optimizing staff use under resource constraints.

3.2.1 Similarity to Parallel Machine Scheduling Problem (PMSP)

The problem structure itself is very closely related to fundamental machine scheduling problems. Graham et al., (1979) introduced a classification system for scheduling problems. According to their criteria in manufacturing, our structure resembles an open shop parallel machine scheduling problem. This means that each employee has a certain number of tasks (slots) that need to be scheduled into rooms (machines) in sequence without overlap.

Our model represents one of the most simplified versions of the problem. With no fixed rooms that employees have to be assigned to, and no fixed order, the problem becomes open shop instead of job shop. All slots also have the same duration regardless of the chosen room, and there are no release dates for the slots. The latter means that each slot can be placed at any time during the week as long as it does not overlap with another slot for the employee. Although this description maps closely to our problem, several adjustments are required.

Including room quality as an objective

The first adjustment is our actual objective. While machine scheduling problems often seek to minimize makespan or costs, our model seeks to ensure feasibility, minimize use of windowless rooms, and improve schedule quality criteria.

For the minimization of windowless rooms, the problem closely mirrors scheduling problems that involve different resource qualities. Kacem et al., (2002) propose hybrid models combining fuzzy logic and optimization to allocate tasks based on resource quality. Jeunet et al., (2020) developed a model for project planning optimization that minimized the makespan while incorporating penalties into the objective function for use of worker overtime. This could map to our problem if we include the use of windowless rooms as an additional minimization in our objective function.

Staff preferences

Another way in which our problem differs is by the inclusion of specific preferences for employees. Employees may for example have a preference for certain time-slots, and not be available at other times. Additionally, as we mentioned in Section 2.1.2, each specialty has a front desk near certain rooms, and employees from that specialty prefer rooms near that front desk. Finally, employees prefer to work in the same room as much as possible throughout the week. Another preference that we heard expressed in the department is that employees prefer to share rooms with a minimal number of other employees throughout the week.

To ascertain all these preferences, we can draw on ideas from existing workforce rostering techniques. Ernst et al., (2004) reviewed strategies for developing schedules that balance employee workloads while adhering to fixed shift durations. The study incorporated both time preferences and availability into the objective function.

Job-shop characteristics

Aside from preferences, the model also has to incorporate certain hard scheduling constraints. These include unavailability for employees, as well as requiring employees to work specific time-slots in certain locations. The latter is equivalent to job-shop characteristics where a certain task has to be completed on a specific machine at a particular time. These requirements will be included as hard constraints. The additional scheduling constraints and objectives may increase the complexity of our model. For this we can look at a study by Burke et al., (2006), who used meta-heuristic approaches to ensure compliance with scheduling constraints.

3.2.2 Complexity and associated solution options

Błażewicz et al., (2007) discusses how job shop scheduling problems, particularly those involving resource prioritization or eligibility constraints, are inherently difficult to solve. The addition of room adequacy constraints in this study further complicates the problem. Furthermore, Martello et al., (1981) showed that even simplified assignments of tasks to resources can lead to complex, combinatorial challenges. Kravchenko et al., (1997) states that the parallel machine scheduling problem is NP-hard even for the case of two identical parallel machines, and developed a MILP solution to solve small instances. G.-de-Alba et al., (2022) developed a mixed integer formulation for the unrelated parallel machine scheduling problem.

Our problem is complicated further through the scheduling constraints. Van den Bergh et al., (2013) comprehensively reviewed healthcare planning and scheduling, showing that Mixed-Integer Linear Programming (MILP) can be effective in staffing with operational constraints, but they do acknowledge the scalability issues. This is corroborated by Ernst et al., (2004). From the previous paragraph demonstrating the complexity of the PSMP, and the additional complexity due to scheduling constraints and room quality, we explore more scalable methods as well.

Burke et al., (2006) explored metaheuristic approaches to ensure compliance with scheduling constraints. An example of such a meta-heuristic is Simulated Annealing (SA), explored for example by Cohn et al., (2009) to schedule medical residents. SA is effective at navigating complex solution spaces and is well-suited for minimizing multiple objectives at once. SA has also been effectively applied to minimize makespan in open shop parallel machine scheduling problems, as demonstrated by Naderi et al., (2011) and Li et al., (2024). The simplicity and adaptability of SA make it a promising solution, especially in balancing employee slot assignments with room quality constraints. Simulated Annealing has also been used in healthcare by Hulshof et al., (2013), who used it to assign surgeries to operating rooms under stochastic durations. For this they did not consider staff scheduling constraints.

3.2.3 Choice of method

The scalability concerns make the feasibility of exact methods such as MILP questionable, although the possibility of the *exact* results that these methods can generate always validates an attempt at their use. In our solution design, we will develop both an exact approach and an approach using SA.

For both approaches we rely on the ideas from Jeunet et al., (2020) and Ernst et al., (2004) for incorporating both windowless room use and staff preferences in our objective function. For the MILP formulation, we combine these formulations with ideas from Kravchenko et al., (1997) and G.-de-Alba et al., (2022), simplified for our solution. For Simulated Annealing we use ideas from Li et al., (2024), Naderi et al., (2011), and Hans et al., (2008).

Although classic Simulated Annealing is able to improve existing solutions, it is not inherently able to *create* an initial feasible solutions. We choose to use a construction heuristic to build this initial solution, specifically random construction. This type of construction explicitly incorporates randomness into the initial solution, which may be useful in avoiding certain local optima over multiple runs (Hans et al., 2008). Random allocation, however, has very little inherent intelligence, and may therefore not not be able to find feasible solutions for each instance. For this reason, we use intermittent local search during construction to ensure feasibility.

3.3 Conclusion and contribution

Operations research in healthcare has mostly focused on decreasing lead times and accomplishing feasibility, but tends to neglect the quality of the work environment and treating patients as close to home as possible. This study takes the first step for both of these goals, by developing a capacity allocation that would enable treatment close to home, as well as developing a framework that can incorporate workspace quality directly into planning.

Given the capacity allocation, including the desired allocation directly into the planning is the next critical step. Without it, the capacity allocation can not actually be fully enabled. Even if hospitals are able to accomplish the right resource allocation across various locations, sufficient facility resources should also be available in those locations. To this end, we develop a modular framework that can be expanded to include location constraints as well.

Although the current model focuses on refining scheduling at one site, we will develop a modular framework that may be expanded to include location-specific constraints. In doing so, the model would address an additional gap in healthcare resource allocation by including staff feasibility, workspace quality, and patient proximity all directly into a planning approach. This framework thereby lays the groundwork for future studies to plan allocations over multiple location.

Introducing this idea first within a single department of a single organization is first step in establishing this method as a solution to the national problem. The next step, for example, would be to expand the analysis to include the entire hospital and include additional objectives. The long-term goal is to facilitate accessibility across the entirety of the Netherlands by *enabling* an adequate distribution of required capacity. This combination of both allocation and planning will be vital in enabling this type of optimization in the Dutch healthcare system.

4 Solution design

In this chapter we describe the specifics of the methods that we will be using throughout the rest of this research. We explain all the steps we took to enable the development of our results. Firstly, we concisely introduce the available data in Section 4.1. This data allows us to determine the workload in hours that was planned in each location (relevant for research question I), as well as the workload in hours that was planned per employee per week (relevant for research question II).

To answer all research questions, we have to expand the use of the available data. The first goal is to determine a new capacity allocation over the locations, for which the approach is outlined in Section 4.2. The second goal is to recalculate the number of consultations that each employee could have executed based on new policies. We explain how we accomplish this in Section 4.3. Lastly, we require a planning method to create new schedules to compare to the current planning approach in the department. The design choices for this planning approach are laid out in Section 4.4, based on the literature described in the previous chapter.

4.1 Introduction available data

The dataset used is derived directly from Isala's database, and contains information on all appointments executed within the department from 2021-2023. Although most relevant information is included, the data does require some additional processing to exclude appointments that are irrelevant on the basis of our research questions. To accomplish this, we informed with the department to determine which appointments were relevant for either the *new capacity allocation* and the *room planning*. All of the steps we undertook to achieve these exclusions, and all of the data that were excluded, are described in Appendix A. We also explain the structure that is required of the data in that same appendix.

4.2 Location capacity allocation

In this section we outline the steps towards identifying number of consultations that could have been scheduled in each location per specialty and profession. Before we do this, we should first distinguish between physical and telephonic appointments (Section 4.2.1). Next, we need to develop a method to determine the closest location for each appointment (Section 4.2.2). We can then, subsequently, determine the number of consultations that could have been scheduled in each location, which we do in Section 4.2.3.

4.2.1 Telephonic appointments

Although a large percentage of the appointment workload involves telephonic appointments (37.18%), we only base the required capacity in the satellite locations on physical appointments (62.82%). This is because the main reason for scheduling any appointments at the satellite locations is to improve the situation for the patients. Doing telephonic appointments at the satellite locations does not reduce patient travel times, but instead increases travel times for employees. Capacity at the satellite locations is also limited, and would be better used for physical appointments from other departments. Employees also know this, and will be less encouraged to move slots to satellite locations if many of the appointments are telephonic. Even so, we can not completely exclude telephonic appointments. This is because the number of physical patients will never exactly match the capacity we schedule at each location, and the remaining capacity will have to be filled with telephonic appointments to avoid wasting capacity. After consulting with the department we decide that telephonic appointments in the satellite locations should cover at most 20% of the total consultation duration in that location for each specialty.

4.2.2 Determining the closest location

To determine the closest location for the physical appointments, it is necessary to obtain information on the place of residence for the patients involved in each appointment. To accomplish this, we add ZIP codes to each appointment through the patient number included in the appointment data. If multiple ZIP codes were associated with an appointment due to multiple referrals, we use the ZIP code that appears most frequently across that patients' referral data. In cases where the frequencies are tied, we select one ZIP code at random.

Analyzing the data, we see that there is one reoccurring entry in the ZIP codes that is not in a typical ZIP code format. This entry instead shows that the patient does not live in the Netherlands (0.16%). We assume this means that the patient will likely not have another appointment, or, if they do, they will be registered from then on. We therefore decide that we will not consider them for reallocation, and set their closest location as the location that they were treated in.

Now that all ZIP codes are included, we can determine the closest location for each appointment. We use the Google Maps Distance Matrix API [15] for this, which determines the travel distance between two given ZIP codes.

Since the ZIP codes of the patients are all in numbers-only format, the patient residences are taken approximately at the center of the districts they live in. We executed the calculation once for every unique ZIP code in the dataset, and then mapped the results back to the appointments. This saved a lot of time as compared to doing the calculation for every single appointment.

The approach was universal for each appointment, except for the appointments from the village of Wijhe and a small number of houses on the eastern side of the river Ijssel, which we allocated to

Zwolle. This is because, to travel to their actual closest location of Heerde, these patients would have to cross the river Ijssel by ferry. Although the route to Heerde is slightly shorter, it may instead take slightly longer depending on ferry departure times. We assume, therefore, that the inconvenience and cost of the ferry are enough to motivate patients to travel to Zwolle instead.. We are not aware of any other significant natural barriers to any of the locations.

Having done all this, we are able to determine what appointments would have been scheduled to each location if every appointment had been scheduled at the closest location. This, however, still does not allow us to determine a new capacity allocation. To accomplish this task, we need to know the number of *consultations* that could be scheduled in each location. We describe this in the next section.

4.2.3 Determining the number of consultations

We aim determine the number of consultations required on a *weekly* basis to match the planned work with the demand in each location. The amount of consultations should be based on the number of weekly hours of demand for each individual specialty and profession. We discussed before, in Section 2.1.5, that patients generally stick with the first employee that treated them. Moreover, if patients are treated in a certain location, they should be treated in that same location for their return appointments as well. This means that employees who treat patients in the satellite locations should be available for future appointments in those satellite locations. As a consequence, any consultation slot that is scheduled in a location should be recurring. Employee capacity is not available for every week of the year, however. In Section 2.1.6 we explained that employees typically work 44 weeks a year. This means that any hours allocated to employees are only available for a typical 44 out of 52 weeks. To compensate for this, we base the weekly capacity on the yearly workload divided by 44 weeks.

To calculate the number of consultation slots that we should allocate, we then divide the weekly required capacity by the typical slot duration in the satellite locations. Through consulting with the department we have determined that the analysis will be done using a standardized slot duration of two hours and thirty minutes. The approach using two hours and thirty minutes is based on current policies in Meppel, where employees are only allocated rooms if they aim for three hours of appointments and obtain at least 85% coverage. This equals two hours and 36 minutes, which we round down to two hours and thirty minutes to facilitate calculations. The calculation is shown in Equation 1.

$$RequiredConsultationsWeek = \frac{1}{2.5}\frac{NumberHoursYear}{44} \tag{1}$$

When determining the number of slots through the calculation outlined above, the outcomes mostly show fractional numbers. In this scenario, the demand of physical patients from that specialty will not be enough to fill up an exact round number of slots of two hours and thirty minutes. Although we can schedule in telephonic appointments, we do want to minimize this, as mentioned before in Section 4.2.1.

Allocating a slot to a location also means that the employee will have to travel to that location every week, *and* occupy that room capacity for the entire slot. We therefore strive to have a minimal percentage of physical appointments before we choose to schedule slots. Through consulting with the department, we decide that at least 80% of the appointments have to be physical across a profession within a specialty. If this is not the case, the fractional slot is instead added to the demand in *Zwolle*.

For *doctors* from the GIM sub-department, the approach is more complicated. In Section 2.1.3 (and Appendix A.2) we discussed that these doctors can also handle appointments that are grouped under the general specialty. This pool is quite large, and we can therefore first distribute it among the GIM doctors to fill up the fractional slots. We do this until all fractional GIM doctor consultations are filled to round numbers. The remaining general pool demand will then be allocated to the specialties with the lowest current percentage of general patients. After this procedure we can then check whether additional telephonic consultations need to be scheduled and use the same 80% physical requirement as before.

4.3 Policies on duration of consultations

In this section we outline the policy used to generate a new number of consultations for each employee on a weekly basis. The goal is to create new configurations that result in a lower number of total slots through a higher average consultation duration. We can aim for a certain average consultation duration, but determining the exact number of slots to acquire a certain average does not work directly. This is because employees do not have workloads that are all an exact multiple of a given duration. Instead, we aim for a certain *standard duration* per slot, using Equation 2. This equation retroactively determines the number of consultations for each employee given the number of hours of care that they executed. By summing the number of slots across employees we obtain the total number of slots.

$$NumberOfConsultationsWeek = \mathbf{RoundUp}(WorkloadWeek/StandardDuration) \tag{2}$$

Other policies could have also worked, but the specific configuration is not as relevant as the actual reduction in the number of slots. This is because we expect the number of consultations that should be scheduled per employee to be highly individualized, and to vary per specialty. The specific policy to implement will therefore be left to the department.

We test several values for the standard duration up to a total of two hours and forty-five minutes. This is due to the fact that the typical time that a room is available for either a morning and afternoon consultation is approximately three hours, and will not exceed it especially given the possibility for no-shows and the inevitable occurrence of unoccupied appointment slots. The lowest value for standard duration that we test is two hours, since this is the first value that slightly decreases the number of slots for H&O. For GIM, two hours provides a similar number of slots to the current planning.

4.4 New planning approach

In this section we describe the design choices for our new planning approach. We start by describing the solution requirements in Section 4.4.1, and then describe our mathematical problem formulation in Section 4.4.2. We created both an exact approach (described in Section 4.4.3), and a heuristics approach. For the heuristics approach we developed several construction heuristics (outlined in Section 4.4.4). To improve on the initial solutions generated by these heuristics, we use Simulated Annealing (SA), explained in Section 4.4.5.

4.4.1 Solution requirements

The purpose of this part of our solution is to develop a new planning approach. In order to do this, we require a tactical planning framework that mimics the current schedule's framework. This means that the model should allow us to schedule employees into specific rooms on designated days or slots, while adhering to various availability constraints, and including individual slot preferences.

The scheduling model should be able to handle instance sizes ranging from individual specialties to the whole department, with manageable run-times to enable multiple iterations. Visualization and export options are also needed for result evaluation. The most important requirement is that the quality of the schedule needs to be relatively independent of the size and complexity of the instance to be planned, which would indicate that comparisons between configurations are valid. That is; similar quality schedules need to be available under similar run-times regardless of instance size and complexity. In the next section, we develop the Mixed-Integer Linear Programming (MILP) model to represent the underlying framework of the planning.

4.4.2 Mathematical model formulation

Sets	
$\mathcal{E}$	Set of employees
$\mathcal{R}$	Set of rooms
$\mathcal{S}$	Set of slots
Parameters	
T_s	Duration of slot **s** (hours)
W_r	Boolean showing whether or not room **r** is windowed.
$A_{e,s}$	Boolean showing whether or not employee **e** is available during slot **s**
$P_{e,r,s}$	Boolean showing whether an employee **e** wants to be scheduled in room **r** during slot **s**
RC_e	Required capacity for employee **e** (hours)
M	Large number
Variables	
$x_{e,r,s}$	Binary variable that indicates whether or not employee **e** works in room **r** during slot **s**
$y_{e,r}$	Binary variable that indicates whether or not employee **e** uses room **r** during the week
v_e	Variable indicating overuse of rooms for an employee **e**
w_r	Variable indicating the number of employees in excess of *one* in room **r**
SC_e	Scheduled capacity for employee **e** (hours)

Table 4a: Model formulation

$$\min = c_1 \sum_{e\in\mathcal{E};r\in\mathcal{R};s\in\mathcal{S}} W_r\ x_{\mathbf{e,r,s}} + c_2 \sum_{e\in\mathcal{E}} v_{\mathbf{e}} + \sum_{r\in\mathcal{R}} w_{\mathbf{r}} - c_4 \sum_{e\in\mathcal{E};r\in\mathcal{R};s\in\mathcal{S}} P_{\mathbf{e,r,s}} x_{\mathbf{e,r,s}} \tag{a}$$

$$\sum_{r\in\nabla} x_{\mathbf{e,r,s}} \leq A_{\mathbf{e,s}} \qquad \forall e\in\mathcal{E}, \forall s\in\mathcal{S} \tag{b}$$

$$M y_{\mathbf{e,r}} \geq \sum_{s\in\mathcal{S}} x_{\mathbf{e,r,s}} \qquad \forall e\in\mathcal{E}, r\in\mathcal{R} \tag{c}$$

$$SC_{\mathbf{e}} = \sum_{r\in\mathcal{R};s\in\mathcal{S}} T_{\mathbf{s}}\ x_{\mathbf{e,r,s}} \qquad \forall e\in\mathcal{E} \tag{d}$$

$$\sum_{r\in\mathcal{R}} x_{\mathbf{e,r,s}} \leq 1 \qquad \forall e\in\mathcal{E}, s\in\mathcal{S} \tag{e}$$

$$\sum_{e\in\mathcal{E}} x_{\mathbf{e,r,s}} \leq 1 \qquad \forall r\in\mathcal{R}, s\in\mathcal{S} \tag{f}$$

$$SC_{\mathbf{e}} \geq RC_{\mathbf{e}} \qquad \forall e\in\mathcal{E} \tag{g}$$

$$v_{\mathbf{e}} \geq \sum_{r\in\mathcal{R}} y_{\mathbf{e,r}} - 1 \qquad \forall e\in\mathcal{E} \tag{h}$$

$$w_{\mathbf{r}} \geq \sum_{e\in\mathcal{E}} y_{\mathbf{e,r}} - 1 \qquad \forall r\in\mathcal{R} \tag{i}$$

$$x_{\mathbf{e,r,s}}\ ;\ y_{\mathbf{e,r}} \in \{0,1\}, \qquad \forall e\in\mathcal{E}, r\in\mathcal{R}, s\in\mathcal{S} \tag{j}$$

$$v_{\mathbf{e}}\ ;\ w_{\mathbf{r}}\ ;\ SC_{\mathbf{e}} \geq 0, \qquad \forall e\in\mathcal{E}, r\in\mathcal{R} \tag{k}$$

Above we work out the base model LP formulation, which allows us to schedule employees into slots within rooms. The objective function is based on our KPI's to ensure minimization of room sharing and switching, all while minimizing use of windowless rooms. In an exact approach, run-times increase exponentially when we try to identify switches between consecutive slots, and we therefore instead minimize the *use of different rooms* for employees. c_1, c_2, c_3, c_4 are the weights of the different objectives. Within the model formulation, (b) makes it so that an employee can only be scheduled in a slot if they are available during that slot (c) defines variable y as being equal to one if employee e uses room r at any point during the week, which is used to define v and w. (d) defines the scheduled workload for employees as being equal to the occupied slots multiplied by the slot durations. (e) makes sure that employees use a maximum of one room per slot. (f) makes sure that a maximum of one employee in each room per slot. (g) ensures that the required capacity is scheduled per employee. (h) defines v as the number of unique rooms in excess of one that employee **e** uses. (i) defines w as the number of unique employees in excess of one that are scheduled into room **r**. (j) ensures the binary nature of decision variables x and y. (k) ensures that remaining decision variables are positive integers.

4.4.3 Exact approach

We implement the model exactly according to the MILP formulation of Section 4.4.2. Although the exact approach may be able to ensure a solution that is at the global optimum, it is generally slower than a heuristic. This means that an exact approach might fall short in computation times and solution quality for larger and more complex instances. Given the combinatorial nature of our problem, it can reach quite a high dimensionality. We therefore choose to perform all our experiments only using our heuristics approach. The exact approach has been implemented, and future exploration of this approach may be useful for operational use in the hospital. For the purpose only of this thesis, however, we ran a very large number of experiments (every week of 2023 for each respective sub-department multiple times). For this we required a much faster planning approach, which is why we use the planning approach described in 4.4.4.

4.4.4 Heuristics approach

This section lays out the decisions we made to construct our planning approach. We explain the relevant algorithms used for the construction of an initial solution, as well as for the simulated annealing heuristic.

Construction heuristics for SA

We mentioned in Section 3.2.3 that we use random allocation to construct an initial solution. This algorithm schedules random available slots for each employee until they reach their quota in terms of the number of consultations. The pseudocode for random construction is shown below in Algorithm 2. The algorithm finds random slots and checks whether they are available, and whether the employee can be added to the found slot according to preferences and availability. To check the latter, we use Algorithm 3. This algorithm checks whether the considered employee already has a consultation in the considered slot, and checks whether the room is occupied in the considered slot. If not, the consultation is added to the employees schedule. This algorithm, $AddToWorkAndRoomSchedule$, also updates all delta-objective values that are stored in the employee and room objects.

If no slot can be found, the algorithm outputs an error message. After the planning of each employee, we run an iteration of DecreaseRoomsUsed, which rearranges the schedule so the current schedule makes use of fewer different rooms. After the planning of each week, we verify the schedule, calculate the objective function, and store the solution as the new best solution. UpdateForm is used to show relevant information to the user.

In Algorithm 8 of the Appendix, we also show an alternative construction heuristic, which uses sequential allocation. The algorithm places employees into one room as much as possible. The selection of employees is done in descending order according to the number of consultations that need to be scheduled. This algorithm plans employees according to their availability in one room from Monday morning up till Friday afternoon according to the number of slots that have to be scheduled for them. The next employee then follows the same procedure in the next room. This is shown in Algorithm 8. This algorithm can achieve good standalone performance, but does not have the added benefit of randomness. The department may consider using the ideas from this algorithm for their planning rules.

Algorithm 2 Random Construction

```
procedure ScheduleAllEmployeesRandomly
const MaxAttempts ← 1000 {Maximum number of attempts to find a slot}
variables: d, r, sℤ {Indices for day, room, slot}
           Week, Employee {Week and Employee objects}
           attempts ∈ ℤ {Tracks attempts for scheduling}
           WorkloadFactor, CutoffWorkload ∈ ℝ
Randomize() {Initialize random number generator}
for Week ∈ (Weeks) do
  for Employee ∈ (Week.Employees) do
    while Employee.RequiredWorkload > Employee.PlannedWorkload do
      attempts ← 0
      repeat
        d, r, s ← Random
        attempts ← attempts + 1
      until (RoomIsEmpty and VerifyAddition(Week.WeekNr, Employee.ID, r, d, s))or
      (attempts ≥ MaxAttempts)
            or (attempts ≥ MaxAttempts)
      DecreaseRoomsUsed
      if attempts ≥ MaxAttempts then
        LogMessage('No suitable schedule found for ' + Employee.Name)
        break
      else
        AddToWorkAndRoomSchedule(Week.WeekNr, Employee.ID, r, d, s)
      end if
    end while
  end for
  VerifySchedule
  CalcObjectiveFunction
  StoreBestSolution
  UpdateForm
end for
```

Algorithm 3 Verify Addition

```
function VerifyAddition(w, empToSchedule, r, d, s): Boolean
Result ← True {Assume true unless a conflict is found}
// Step 1: Check if the employee is already occupied in the target timeslot
if Week.Employees[empToSchedule].WorkSchedule[d][s] = Empty then
  Result ← False
  return Result {Cannot add; employee is busy}
end if
// Step 2: Check if the room is already occupied in the target timeslot
if Week.Rooms[r].PlannedEmployeeIDs[d][s] = None then
  Result ← False
  return Result {Cannot add; room is busy}
end if
return Result {If no conflicts, addition is possible}
```

4.4.5 Simulated Annealing (SA)

Simulated Annealing was inspired by the physical annealing process in metallurgy, and uses a temperature parameter to balance the rate of exploration and exploitation during its optimization. The heuristic allows for exploration at the start to find a broad range of potential solution spaces. When lower temperatures are reached, the heuristic reduces to a gradient descent (exploitation), where the heuristic moves towards the local optimum in the found solution space.

The heuristic obtains new solutions using neighborhood operators. These are functions that slightly perturb the current solution, for example by moving a consultation from one day to another, or from one room to another. The heuristic then evaluates the perturbed solution to see if it has a lower objective value (meaning it is objectively better). If so, it will keep the new solution. If not, it will keep the new solution with a certain acceptance probability. For this acceptance probability, we use the following Equation 3, based on the Boltzmann distribution (Kirkpatrick et al., 1983). Here NewOV is the objective value (OV) of the solution after a swap has taken place, and CurrentOV is the objective value of the solution before the swap.

$$\begin{aligned} NewOV - CurrentOV > 0: &\qquad P_{accept} = e^{-\frac{(NewOV - CurrentOV)}{T}} \\ NewOV - CurrentOV \leq 0: &\qquad P_{accept} = 1 \end{aligned} \tag{3}$$

We should choose our *neighborhood operators* and *cooling scheme* carefully, since they determine how efficiently the heuristic is able to find better solutions. An ineffective choice of neighborhood operators, in particular, can limit the possible solutions that the heuristic can find as a result of getting stuck in local optima.

Neighborhood operators

Our initial solution will have some slots that are filled, and (most of the time) some slots that are empty. Knowing this, we can think immediately of two possible neighborhood operators. The first is to swap two different slots, and the second is to move a certain slot to an empty spot. We know that situations may emerge where no single swap or move will immediately improve the situation anymore. To avoid this situation, we also enable the heuristic to perform either two or three operations at once before evaluating the new objective value. The number of operations is chosen with a certain probability; one swap with 50% probability, two swaps with 30% and three swaps with 20%.

Description of SA algorithms

We show our local search procedure in Algorithm 4. This algorithm manages all required steps in obtaining new solutions. It calls Algorithm 5 on each iteration, which finds a random neighborhood solution. The local search algorithm then restores the current solution with a certain probability if the neighborhood solution has a higher objective value.

To find a new solution, Algorithm 5 calls Algorithm 6, which verifies whether a swap is possible. Swapping an employee from an old slot to a new one is possible under two conditions. Firstly, the employee needs to be available in the new slot. Secondly, the room needs to be available in the new slot. The swap verification is used to verify the swap from each respective employees perspective (so we verify it two times, both ways).

Algorithm 4 Local Search Algorithm for Simulated Annealing

```
procedure LocalSearchSA(CurrentTemp, w, MarkovChainLength)
var i, j: Integer, NumChanges: Integer
for i ← 0 to MarkovChainLength − 1 do
  Application.ProcessMessages {Process pending application messages}
  Increment NumAttempted {Track the number of attempted neighbor solutions}
  NumChanges ← DetermineNumChanges {Determine number of perturbations}
  for j ← 0 to NumChanges − 1 do
    SwapSlotsRandomly {Perform random slot swaps}
  end for
  CalcObjectiveFunction {Calculate the new objective value given values stored in employ-
  ees and rooms}
  if Week.OV < Week.CurrentOV then
    StoreCurrentSolution {Update current solution to the neighbor solution}
    Increment NumAcceptedCurrentBetter {Track better solutions found}
    if Week.OV ≤ Week.BestOV then
      StoreBestSolution {Update best solution}
      Increment NumAcceptedBest {Track best solutions found}
    end if
  else
    if Week.CurrentOV = Week.OV then
      Increment NumAcceptedCurrentEqual {Track equivalent solutions}
    end if
    if Random < exp((Week.CurrentOV − Week.OV) / CurrentTemp) then
      if Week.CurrentOV ≠ Week.OV then
        Increment NumAcceptedCurrentWorse {Track worse solutions accepted}
      end if
      StoreCurrentSolution {Accept the worse neighbor solution}
    else
      RestoreCurrentSolution {Reset to the original solution}
    end if
  end if
  Update Form1.Edit9.Text ← FloatToStr(Week.CurrentOV) {Display current objective
  value}
  Update Form1.Edit10.Text ← FloatToStr(Week.BestOV) {Display best objective value}
end for
```

Algorithm 5 Swap Slots Randomly

```
procedure SwapSlotsRandomly(w)
const MaxAttempts ← 100 {Maximum number of attempts to find a valid swap}
var d1, d2, r1, r2, s1, s2, e1, e2, attempts: Integer
Randomize() {Initialize random number generator}
attempts ← 0
repeat
   // Step 1: Randomly select two different slots
   d1, d2, r1, r2, s1, s2, e1, e2 ← Random
   while (d1 = d2) ∧ (r1 = r2) ∧ (s1 = s2) do
      d2, r2, s2 ← Random
   end while
   // Step 2: Retrieve employee IDs assigned to these slots
   e1 ← Week.Rooms[r1].PlannedEmployeeIDs[d1][s1]
   e2 ← Week.Rooms[r2].PlannedEmployeeIDs[d2][s2]
   // Step 3: Verify if the swap is valid
   if VerifySwap(w, e1, e2, r2, d1, s1, d2, s2) and VerifySwap(w, e2, e1, r1, d2, s2, d1, s1)
   then
      RemoveFromWorkAndRoomSchedule(w, e1, r1, d1, s1)
      RemoveFromWorkAndRoomSchedule(w, e2, r2, d2, s2)
      AddToWorkAndRoomSchedule(w, e1, r2, d2, s2)
      AddToWorkAndRoomSchedule(w, e2, r1, d1, s1)
      break {Exit loop if swap is successful}
   end if
   attempts ← attempts + 1
until attempts ≥ MaxAttempts
```

Algorithm 6 Verify swap from old to new (one-way)

function VerifySwap(w, $empToSwap$, $empToSwapWith$, $newRoom$, $oldDay$, $oldSlot$, $newDay$, $newSlot$): **Boolean**
  Result ← **True** {Assume true unless a conflict is found}
  // Step 1: Check if the slot to swap to has an employee
  if $empToSwap \neq None$ **then**
    // Step 1.1: Check if the timeslot to swap has an employee in it, since empty slots are always eligible for swapping
    if ($oldDay = newDay$) **and** ($oldSlot = newSlot$) **then**
      return Result
    end if
    // Step 1.2: Verify if workload conservation is maintained
    if $oldSlot \neq newSlot$ **then**
      if ($oldSlot = Morning$) **and** ($Week.Employees[empToSwap].PlannedWorkloadWeek - DurationMorning + DurationAfternoon < Week.Employees[empToSwap].RequiredWorkload$) **then**
        Result ← **False**
        return Result
      else if ($oldSlot = Afternoon$) **and** ($Week.Employees[empToSwap].PlannedWorkloadWeek + DurationMorning - DurationAfternoon < Week.Employees[empToSwap].RequiredWorkload$) **then**
        Result ← **False**
        return Result
      end if
    end if
    // Step 1.3: If the employee is occupied at the timeslot we want to swap to, then the swap isn't available, unless we are swapping the same employee
    if $Week.Employees[empToSwap].WorkSchedule[newDay][newSlot] \neq -1$ **then**
      if $empToSwap \neq empToSwapWith$ **then**
        Result ← **False**
        return Result
      end if
    end if
    // Step 1.4: Check if the new room and timeslot are occupied
    if $Week.Rooms[newRoom].PlannedEmployeeIDs[newDay][newSlot] \neq -1$ **then**
      Result ← **False**
      return Result
    end if
  end if
  return Result {Return true if empToSwap can be swapped to the new slot}

5 Experimental design

In this chapter we describe the experiments used to answer our research questions. We discuss the required data, the experimental parameters, and the evaluation criteria. The experiments cover the current capacity allocation over the locations in Section 5.1, as well as the new allocation in Section 5.2. Section 5.3 covers the experiments on the *current* room schedule, while the remaining experiments describe approaches towards improving it. The experiment in Section 5.4 covers the policy to decrease the total number of consultations in the schedule, as described before. Section 5.5 describes how we evaluate the performance of the new planning approach. Section 5.6 covers the experiments combining the previous two methods. We summarize the evaluation criteria for each experiment at the end of the chapter.

5.1 Experiment 1: Current location capacity allocation

The goal of the experiments introduced in this section is to determine the performance of the current location capacity allocation. We will show the total work planned in each location per specialty, and how this compares to the total demand per specialty. Based on the results, we will identify options for improving the planning of appointments given the current allocation. The results are obtained by analyzing the appointment data from the department given the processing described in Appendix A.

We evaluate this allocation for each sub-department, mainly based on the percentage of patients treated in the closest location. We also identify estimates of potential reductions in CO_2 emissions, as well as potential reductions in the number of consultations planned in Zwolle.

5.2 Experiment 2: New location capacity allocation

The purpose of this experiment is to determine a new capacity allocation in terms of professions, specialties, and locations. We apply the approaches described in Section 4.2.3 to obtain this allocation. We deliver the results through two tables. In the first table we show the number of hours before any of the analyses, with the intent of enabling repetition of the calculations given different parameters. In the second we show the number of slots given a standardized slot duration of two hours and thirty minutes, as discussed in Section 4.2.3.

This experiment uses the same data as the previous experiment, and we once again analyze results in terms of the expected percentage of patients that can be treated in the closest location, estimated CO_2 reductions, and reductions in the number of consultations in Zwolle. The analysis is additionally extended in one way with respect to the previous experiment. Specifically, we also analyze how the newly determined allocation would have performed in 2023 given *weekly* demand variability, which has not been explicitly accounted for in our approach.

5.3 Experiment 3: Current room planning

We evaluate the performance of the current planning in this experiment. These results are not based on the appointment data, but instead use the Excel room schedules described in Section 2.2. The available schedules were weeks 4-52 of 2023 for GIM, excluding week 8, and weeks 32-2023 to week 13-2024 for H&O. The remaining schedules were either incomplete or had, at this point, been deleted by the planning department.

For each schedule, we determine the total number of consultations planned during each week, as well as the total number of windowless consultation slots. Due to time limitations, we did not write a program that automatically determined the number of room switches and counts of room-sharing. Instead, these metrics are evaluated for one individual week for each respective sub-department, specifically week 7 for GIM and week 43 for H&O. Most other weeks would have been viable too, except for reduction-weeks and weeks in 2024. Reduction weeks have an atypical number of consultation slots for employees. Moreover, we had no access to appointment data of 2024, which means that choosing a schedule from 2024 prohibited comparison to schedules created based on appointment data in further experiments.

Our analysis involves one additional step: determining how the number of windowless consultations changed with the total number of consultations in the current schedule. We relate these metrics to each other in a graphical format for each week and determine the linear regression showing their interdependency. We exclude weeks with zero windowless consultation slots from the regression. This is because any week with fewer than eighty slots for H&O always results in zero windowless slots in the available data. This, therefore, presents a different regime of the data. We therefore exclude weeks 32, 33, 44, 52 of 2023 and week 1 of 2024 from the H&O regression. The last criteria that we evaluate the current planning on are the number of times that a room was occupied for anything other than a consultation slot, and the number of mistakes in the manual planning.

5.4 Experiment 4: Policies on consultation durations

This section's experiment serves to evaluate the performance of a decrease in the total number of consultation slots in the schedules. A policy to generate this decrease in total slots has been described in Section 4.3. The values used for the standard duration vary up to two hours and 45 minutes for both sub-departments, as we also described in that section. We use the appointment dataset described before (Appendix A to provide input data for WorkloadWeek. This workload is based on the total workload *planned* in Zwolle, both physical and telephonic. The reason for using the planned workload instead of the demand as calculated, is that this demand will likely not be realized in the near future. This means that using the current planning is more relevant for generating insights into Isala's immediate options for planning improvements.

We calculate the range of expected total consultations across all weeks in 2023, given the various standard durations. Given these values, we *estimate* the number of windowless consultation slots using the linear regression from the previous experiment. We show the average and maximum number of consultations (total and windowless). The number of room switches and room shares can not be predicted for this policy, since we only know these values for weeks 7 of GIM and 43 of H&O in the current planning.

5.5 Experiment 5: New planning of current room schedules

The goal of the experiments in this section is to determine the performance of a new planning approach. The specific technique used is Simulated Annealing, as outlined in Section 4.4. We use the values derived from the current room schedules as input for the room schedule, such that we create a new planning for every available schedule. The available schedules were, again, weeks 4-52 of 2023 for GIM, excluding week 8, and weeks 32-2023 to week 13-2024 for H&O. The number of employees for these schedules varies between 39 and 67 for GIM and between 13 and 24 for H&O. The number of rooms used for each schedule of GIM is 23 windowed rooms and fourteen windowless rooms. The number of rooms used for each schedule of H&O is eleven windowed rooms and two windowless rooms, representing the current structure of the room schedules. (Figure 2C on page 19)

We again evaluate room switching and room sharing counts in weeks 7 (GIM) and 43 (H&O) to compare the new planning approach to the current planning approach. We also evaluate these metrics for weeks 12 and 16 for GIM and weeks 40 and 50 for H&O to show whether performance in weeks 7 and 43 is generally replicable for the specific planning approach we used.

5.6 Experiment 6: Combination of policies and SA scheduling

This experiment combines the approaches in the previous two experiments. The goal is to investigate the joined effects. We therefore use the input from decreasing the total number of slots with standard durations, covering all of 2023 for both sub-departments. This workload is again based on the total workload *planned* in Zwolle, both physical and telephonic. The planning will be created using SA. Schedules will again be evaluated on the number of windowless slots in each week. Out of time restrictions, we only use a standard duration of two hours to evaluate *all* weeks. The number of windowless consultation slots for longer durations (i.e. fewer total slots) will subsequently be *estimated* on this basis. Again, the number of rooms used for each schedule of GIM is 23 windowed rooms and fourteen windowless rooms, and eleven windowed rooms and two windowless rooms for H&O. Week 7 for GIM and 43 for H&O will both be evaluated on room switching and counts of room sharing for *each* of the mentioned standard durations.

Experiment	*Percentage planned closest 2023*	*Reductions CO_2*	*Number of consultations planned each week*	*Number of windowless consultations planned each week*	*Room switches and counts of room sharing: week 7 GIM week 43 H&O*	*Room switches and counts of room sharing: week 12/16 GIM week 40/50 H&O*
Current location allocation	x	x (est.)				
New location allocation	x	x (est.)	x (est.)			
Current planning			x	x	x	
New slots			x	x (est.)		
Rescheduling				x	x	x
New slots 2 hrs + SA				x		
New slots (2 hrs – 2.75 hrs) + SA				x (est.)	x	

Table 5a: Evaluation criteria for each experiment, est. refers to estimated

5.7 Choice of SA cooling scheme for experiments 5 and 6

Having created the SA program, we have to decide on our cooling scheme parameters. Our approach for determining these has been described in Appendix C. We use a starting temperature of 400, a stopping temperature of 10^{-5}, a cooling rate of 0.86, and a Markov chain length of 10^4. The calculated parameter configuration should allow for approximately 10^6 total swaps.

6 Results & Analysis

In this chapter we demonstrate the results of the experiments described in the previous chapter. The results cover the current capacity allocation over the locations in Section 6.1, as well as the new allocation in Section 6.2. Section 6.3 covers the performance on the *current* room schedule, while the remaining sections describe the performance of our improvement approaches. Of these, Section 6.4 covers the policy to decrease the total number of consultations in the schedule, and Section 6.5 evaluates the results of implementing the new planning approach. Section 6.6 evaluates the combination of the previous two methods. We analyze the results in Section 6.7.

6.1 Experiment 1: Current location capacity allocation

We compare the planned workload to the potential workload in each location, per specialty and sub-department. We then describe options for improving the allocation given the current capacity. From here on, we express the current situation in terms of *Planned Work* and the distribution resulting from our analysis in terms of *Demand*. This demand is only an expression of our location analysis. A dashboard showing further supporting insights into the planning of the department is shown in Appendix D.

Performance evaluation

Percentage of patients planned to closest location: Table 6a summarizes the demand and planned work (in hours) for each location throughout 2023. It also includes the percentage of demand correctly scheduled to each location. For Zwolle, the majority (94.5%) of the *physical* demand was correctly allocated to Zwolle. In Meppel, while the planned physical work closely matches the demand (4156 vs. 4043 hours), only 58.4% of the *actual* demand was scheduled. For Kampen, Heerde, and Steenwijk, the planned work accounted for 25.2%, 12.5%, and 18.7% of their respective physical demands.

Reduction in CO_2 emissions: The average duration of appointments in the department is twenty minutes and twelve seconds. Therefore, we estimate that each hour of care involves approximately three patients. We simultaneously estimate that each patient scheduled to a closer location saves five kilometres on a round trip. This is a conservative estimate, based on analyzing several locations on the boundary between multiple locations. An example is Mullegenweg 1, shown in Appendix A.9 in Table A-VII. Given these values, we estimate that each hour of care moved to a closer location saves 2.1 kilograms of CO_2 emissions [34].

Options for improving the planning: Unmet demand for Meppel, Kampen, and Heerde was primarily planned in Zwolle, while unmet demand for Steenwijk was mostly planned in Meppel. While the planned work in Zwolle and Meppel contains a lot of physical work from other locations, most of the physical work in the satellite locations was genuine physical demand for those locations. In

Kampen, for example, this holds for 1013 of the 1140 planned physical hours (88.9%) of the planned hours. For Heerde and Steenwijk, this holds for 82.7% and 80.1% of the physical hours. Although these percentages are high, they may still be increased. Figure 6A illustrates this, graphically displaying the physical demand and planned work from the table. In the figure, bar charts represent planned work by specialty, while reference lines indicate demand. The percentage of planned demand met at each location is highlighted by the part of the bar whose colour corresponds to the location. For example, the green section of the bar for Zwolle shows the portion of planned work that matched Zwolle's actual demand. The department also scheduled telephonic consultations in the satellite locations, specifically 574, 436, and 399 hours, in Kampen, Heerde, and Steenwijk, respectively. These could potentially be replaced with physical demand.

The associated reductions in the number of consultations will be discussed in Section 6.7, since it requires additional information from the following experiments.

Planned work and demand in hours for the entire department (2023)

Location	**Planned work (total)**	**Planned work (physical)**	**Demand**	**Planned demand**	**% of demand planned**
Zwolle	25929 (73.2%)	15895 (71.5%)	8417 (37.8%)	7964	94.5%
Meppel	6281 (17.7%)	4156 (18.7%)	4043 (18.2%)	2360	58.4%
Kampen	1587 (4.5%)	1140 (5.1%)	3992 (17.9%)	1013	25.2%
Heerde	886 (2.5%)	520 (2.3%)	3417 (15.4%)	430	12.5%
Steenwijk	725 (2.1%)	532 (2.4%)	2374 (10.7%)	426	18.7%
Total	35409	22242	22242	12194	54.7%

Table 6a: Table showing demand and how much of the demand was scheduled to the right location

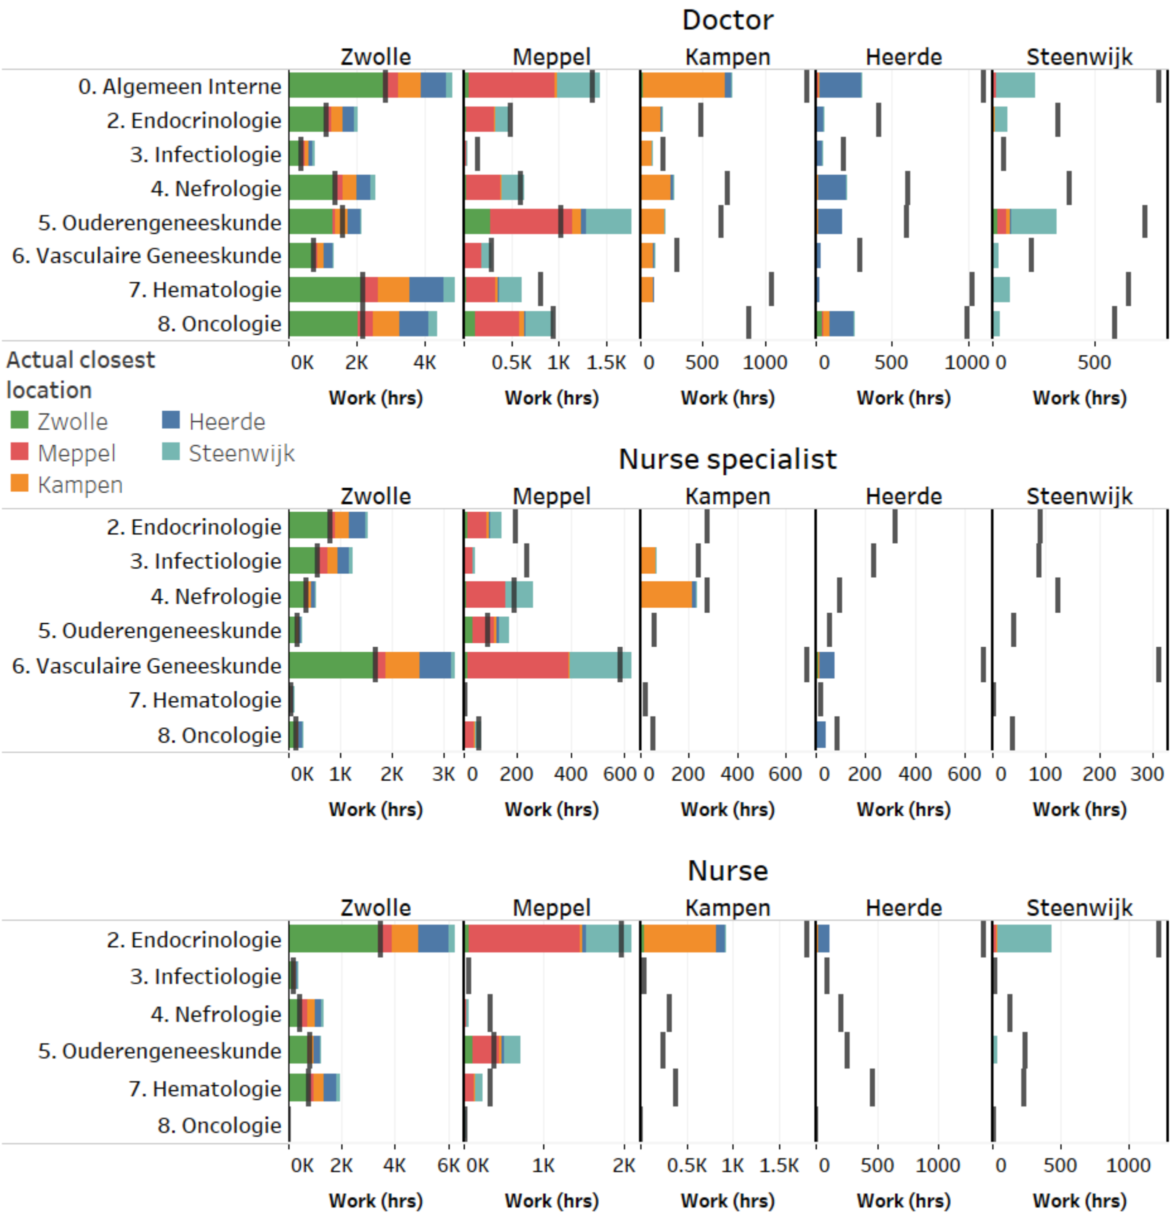

Figure 6A: Planned work (bar) versus demand (reference line) for each specialty per relevant profession

6.2 Experiment 2: New location capacity allocation

In this section we show the maximum demand expressed in hours of physical appointments for each location. From these values, we calculate the number of consultation slots that can be scheduled per week, broken down by location, specialty, and employee.

Doctors' assistants, co-assistants, and physicians' assistants

The hours worked by doctors' assistants, co-assistants, and physicians' assistants are consistently below two hours per week for the satellite locations, regardless of specialty (Table D-XI). Under our criteria from Section 4.2.3, we do not allocate any slots directly to these professions. However, after observing the planned work, we saw that some doctors' assistants already work in Meppel and Kampen. Table D-XII therefore provides the physical demand for these assistants, along with the combined demand for other doctors' assistants. Co-assistants also work in Meppel already, with a total demand of 3.17 hours per week for physical appointments. The yearly hours of unconsidered demand owing to these three professions, across all specialties, are 333.6, 327.7, 207.47, and 168.1 hours for Meppel, Kampen, Heerde, and Steenwijk, respectively.

Doctors, nurses and nurse specialists

The remainder of our analysis focuses solely on doctors, nurses, and nurse specialists. The physical demand in hours per location and profession is shown Table 6b. We also show the telephonic hours for completeness. Table 6c shows the resulting slots for Meppel, Kampen, Heerde, and Steenwijk, following the approach outlined in Section 4.2.3. Table 6d shows the remaining workload for Zwolle, as well as the remaining telephonic slots. *We reiterate that the slots for GIM doctors (specialties 2-6) also contain the general hours.*

Table 6e shows the percentage of general patients for GIM doctors given the slot allocation, which is relevant for the department for the department's decision to plan slots. Table 6c presents the slots as either fractional or whole numbers. The whole numbers indicate slots that contain no telephonic appointments. Fractional numbers represent the physical demand of a slot, with telephonic demand shown in parentheses. For nephrology nurses, the slots in Kampen and Heerde remain fractional even after including telephonic consultations. This is because the telephonic demand is insufficient to complete the slots.

We aggregate across the totals for each profession to obtain the total slots per location. For Meppel, Kampen, Heerde, and Steenwijk, the total slots that can be scheduled given only physical demand are 32, 33, 28, and 19 respectively. The maximum available capacities in those locations are 80, 27 (even) and 29 (odd), 22 (even) and 24 (odd), and 38 (even) and 43 (odd). This means that all consultation slots for both Meppel and Steenwijk can be planned, whereas 6 (even) and 4 (odd) additional slots would be needed in both Kampen and Heerde. The capacity in each location was shown in Table 2a on page 19 (Section 2.1.2).

Plannable weekly hours of demand; all locations and telephonic						
Doctors	**Zwolle**	**Meppel**	**Kampen**	**Heerde**	**Steenwijk**	**Telephonic**
0. Algemeen Interne	20.29	10.70	10.11	7.72	6.06	12.02
2. Endocrinologie	9.04	3.94	3.85	3.03	2.49	10.22
3. Infectiologie	2.53	1.15	1.27	1.26	0.41	1.95
4. Nefrologie	10.62	4.79	5.70	4.62	2.86	4.79
5. Ouderengeneeskunde	13.76	9.50	5.71	5.42	6.83	3.70
6. Vasculaire Geneeskunde	5.68	2.35	2.13	2.02	1.40	2.77
7. Hematologie	16.71	7.02	8.79	7.64	4.70	13.51
8. Oncologie	17.82	7.22	7.06	7.88	4.28	15.27
Total	**96.45**	**46.67**	**44.62**	**39.58**	**29.03**	**64.22**
Nurse Specialists	**Zwolle**	**Meppel**	**Kampen**	**Heerde**	**Steenwijk**	**Telephonic**
2. Endocrinologie	6.07	1.70	2.58	2.93	0.55	6.86
3. Infectiologie	4.15	2.06	1.93	2.09	0.63	2.66
4. Nefrologie	2.35	1.43	1.94	0.61	0.90	0.57
5. Ouderengeneeskunde	2.52	1.35	0.78	0.80	0.63	0.93
6. Vasculaire Geneeskunde	11.49	4.18	4.65	4.71	2.29	5.26
7. Hematologie	0.04		0.01	0.04		0.06
8. Oncologie	1.90	0.88	0.66	1.25	0.46	1.91
Total	**28.51**	**11.60**	**12.55**	**12.42**	**5.45**	**18.24**
Nurses	**Zwolle**	**Meppel**	**Kampen**	**Heerde**	**Steenwijk**	**Telephonic**
2. Endocrinologie	30.94	15.24	15.73	12.15	10.22	25.76
3. Infectiologie	1.54	0.33	0.41	0.72	0.24	0.46
4. Nefrologie	3.31	3.44	4.52	2.20	1.28	0.13
5. Ouderengeneeskunde	7.38	3.77	2.23	2.52	2.31	0.61
7. Hematologie	5.67	3.25	3.21	3.36	1.60	5.05
Total	**48.83**	**26.03**	**26.10**	**20.95**	**15.64**	**32.00**

Table 6b: Summary of Required Hours for Doctors, Nurse Specialists, and Nurses, telephonically and per location

Plannable weekly slots of demand; Meppel and satellite locations				
Doctors	**Meppel**	**Kampen**	**Heerde**	**Steenwijk**
2. Endocrinologie	2	2	2	1
3. Infectiologie	1	1	1	1
4. Nefrologie	3	3.8 (0.2)	2.6 (0.4)	2
5. Ouderengeneeskunde	5	3.8 (0.2)	3	3
6. Vasculaire Geneeskunde	2	1	1	1
7. Hematologie	2.8 (0.2)	3.5 (0.5)	3	1.9 (0.1)
8. Oncologie	2.9 (0.1)	2.8 (0.2)	3	1.7 (0.3)
Subtotal	**19**	**19**	**16**	**12**
Nurse Specialists	**Meppel**	**Kampen**	**Heerde**	**Steenwijk**
2. Endocrinologie	0	1	1	0
3. Infectiologie	0.8 (0.2)	0	0.8 (0.2)	0
4. Nefrologie	0	0	0	0
5. Ouderengeneeskunde	0	0	0	0
6. Vasculaire Geneeskunde	1.7 (0.3)	1.9 (0.1)	1.9 (0.1)	0.9 (0.1)
7. Hematologie	0	0	0	0
8. Oncologie	0	0	0	0
Subtotal	**3**	**3**	**4**	**1**
Nurses	**Meppel**	**Kampen**	**Heerde**	**Steenwijk**
2. Endocrinologie	6.1 (0.9)	6.3 (0.7)	4.9 (0.1)	4.1 (0.9)
3. Infectiologie	0	0	0	0
4. Nefrologie	1	1.8 (0.1)	0.9 (0.0)	0
5. Ouderengeneeskunde	1	0.9 (0.1)	1	0.9 (0.1)
7. Hematologie	1	1	1	0
Subtotal	**10**	**11**	**8**	**6**
Total	**32**	**33**	**28**	**19**

Table 6c: Summary of Required Slots for Doctors, Nurse Specialists, and Nurses by Location

Required weekly slots; Zwolle and telephonic						
	Doctor		**Nurse**		**Nurse Specialist**	
	Zwolle	**Tele-phonic**	**Zwolle**	**Tele-phonic**	**Zwolle**	**Tele-phonic**
2. Endocrinologie	5.0	15.6	12.4	23.1	3.5	6.9
3. Infectiologie	2.0	3.0	1.3	0.5	2.7	2.3
4. Nefrologie	6.0	6.6	2.2	0.0	2.9	0.6
5. Ouderengeneeskunde	8.0	5.0	3.5	0.4	2.4	0.9
6. Vasculaire Geneeskunde	3.7	4.3	0.0	0.0	4.6	4.6
7. Hematologie	6.7	12.7	3.8	5.0	0.0	0.0
8. Oncologie	7.3	14.7	0.0	0.0	2.1	1.9
Total	**102**		**54**		**38**	

Table 6d: Department slot distribution in Zwolle by profession

Doctors	**Zwolle**	**Tele**	**Meppel**	**Kampen**	**Heerde**	**Steenwijk**
2. Endocrinologie	40%	35%	21%	23%	40%	0%
3. Infectiologie	50%	37%	54%	49%	50%	84%
4. Nefrologie	30%	38%	36%	39%	30%	43%
5. Ouderengeneeskunde	31%	30%	39%	24%	28%	10%
6. Vasculaire Geneeskunde	38%	35%	53%	15%	19%	44%

Table 6e: Percentage of general work for GIM doctors across locations

Performance evaluation

Percentage of patients planned to the closest location: We can not directly calculate this percentage, since the department's ability to realize the full allocation may be limited. This is for many reasons, first and foremost being the complexity of planning the right employees to the available capacity at the correct times. Considerations similar to those observed in the room capacity planning problem are involved in this problem as well, such as the external obligations and limited availability for employees. Instead, the determined allocation serves as a guideline that the department can use to inform future employee allocation choices. The expected increase in the number of physical patients treated in the closest location is therefore dependent on the number of slots that the department *is able to* schedule in the locations, and we therefore make our results contingent on this fact.

Assuming that each slot contributes 44 weeks of 2.5 hours of physical patients treated in the closest location, each slot would contribute 110 hours per year. This equals an additional 0.5% of the total physical demand for each additional slot planned, with the maximum number of slots per employee type and specialty shown in Table 6c.

The maximum percentage of patients that could be treated in each location can be determined by evaluating the performance of the allocation determined in Table 6c. Scheduling all hours exactly as in Table 6b would mean that 100% of the patients for doctors, nurses and nurse specialists would be scheduled in the closest location. As we explained in Section 4.2.3, however, some slots will not be scheduled in those four locations. This occurs whenever the number of hours of demand does not suffice based on the limitations discussed in that same section. The hours not allocated to those four locations will instead be allocated to Zwolle. This holds for 144.0, 114.4, 92.4, and 110.0 yearly hours for the four other locations, respectively.

Other hours not considered in our allocation are the demand for doctors' assistants, co-assistants, and physicians' assistants. Given these limitations, we can determine the absolute maximum percentage of physical patients that can be scheduled to each location under perfect conditions and sufficient capacity. We express the percentages as part of the total demand, and as part of the demand for doctors, nurses, and nurse specialists only. This is shown in Table 6f. These numbers represent the maximum number of hours, as well as the percentage, of physical patients that would be treated given the provided allocation. To reiterate, however, we do not expect the department to be able to realize these results. Instead, we expect incremental improvements of 0.5% per slot added, according to the configuration in Table 6c.

CO2 reductions: Based on the same analysis as in Section 6.1, every 2.5 hours of physical demand moved to the satellite locations reduces the CO_2 emissions of the department by an estimated 5.3 kg [34].

Location	**Number of hours planned**	**% of doctors, nurses, and nurse specialists**	**% of total**
Zwolle	8417	100%	100%
Meppel	3555	95.8%	87.9%
Kampen	3550	96.8%	88.9%
Heerde	3117	97.1%	91.2%
Steenwijk	2096	95.0%	88.3%
Total	**20705**	**97.6%**	**93.1%**

Table 6f: Percentage of closest patients treated per location given perfect use of the new allocation

6.3 Experiment 3: Current planning

We show the performance of the current planning for the weeks discussed in Section 5.3 below. The number of slots and number of windowless slots are shown in Figures 6B and 6C for both sub-departments. The red graph shows the total number of slots, and the blue graph shows the number of windowless slots. The black graph shows the available *windowed* capacity for each sub-department. For GIM we show two additional graphs. The green graph shows the red graph plus the number of administration slots that were scheduled into windowed rooms. The orange graph shows the total number of slots planned within the rooms, including slots for management and other non-outpatient purposes. For H&O we show one additional orange graph, which also contains slots that were scheduled at home or other locations in Isala. The hollow dots in the graphs represent reduction weeks. These fell on weeks 2023-1, 9, 17-18, 30-35, 43, and 52 for GIM. For H&O they fell on 2023-32 to 2023-35, 2023-44, 2023-52, 2024-1, and 2024-9.

Performance evaluation

Number of consultations planned: Looking at Figure 6B, we see that the number of planned slots exceeded the room capacity (230 slots) in weeks 12 and 47 for the red graph. This means that use of windowless rooms was required to be able to schedule all consultations during these weeks. The maximum number of consultation slots for GIM was 235, with an average of 197.2. Looking at Figure 6C, we see similar results for H&O as for GIM. In most weeks the number of slots is lower than the available capacity (110 slots). Weeks for which the capacity was instead exceeded were 6, 7, 9, 12, 13, 14, and 15 of 2024. The maximum number of consultation slots for H&O was 118, with an average of 95.6.

Number of windowless consultations planned: The number of consultations in windowless slots (blue graph) averages around 36.2 for GIM, whereas there are only two weeks where the red graph exceeded the capacity (by 5 and 3 respectively). This means that the department used windowless slots even though windowed slots were still available. The number of windowless slots never reaches zero for GIM, even for the minimum number of slots in week 52 (108 slots). The maximum number of windowless slots for GIM was 47, with an average of 36.2. The number of consultations in windowless slots (blue graph) averages around 5.2 for H&O, whereas the average number of slots in excess of capacity is only 1.1. This again means that windowless slots were used even though windowed slots were still available. H&O can achieve zero windowless slots in several reduction weeks, indicating that a low enough number of slots may prevent the use of windowless rooms in the future, even given the current planning methods. The maximum number of windowless slots for H&O was 14, with an average of 5.2.

Room switches & counts of room sharing For week 7 of 2023 for GIM, the number of switches was 47 and the number of shares was 59. For week 43 of 2023 for H&O, the number of switches was 13 and the number of shares was 21.

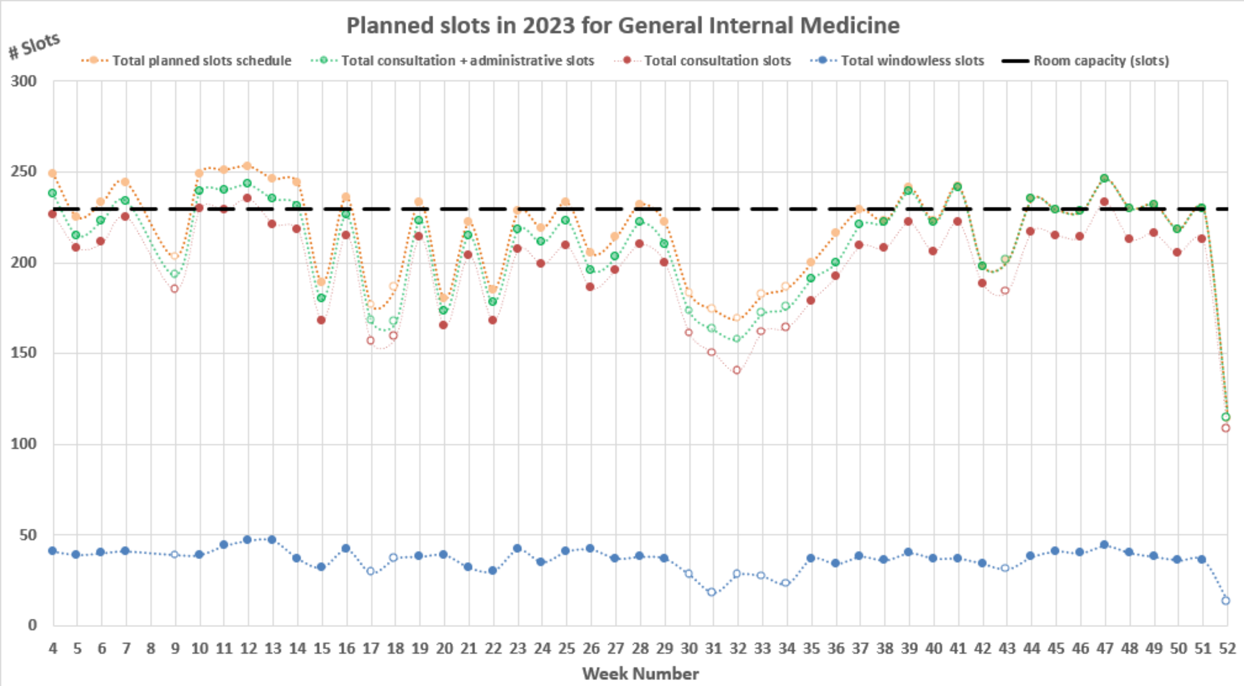

Figure 6B: Number of slots planned within 2023 for GIM

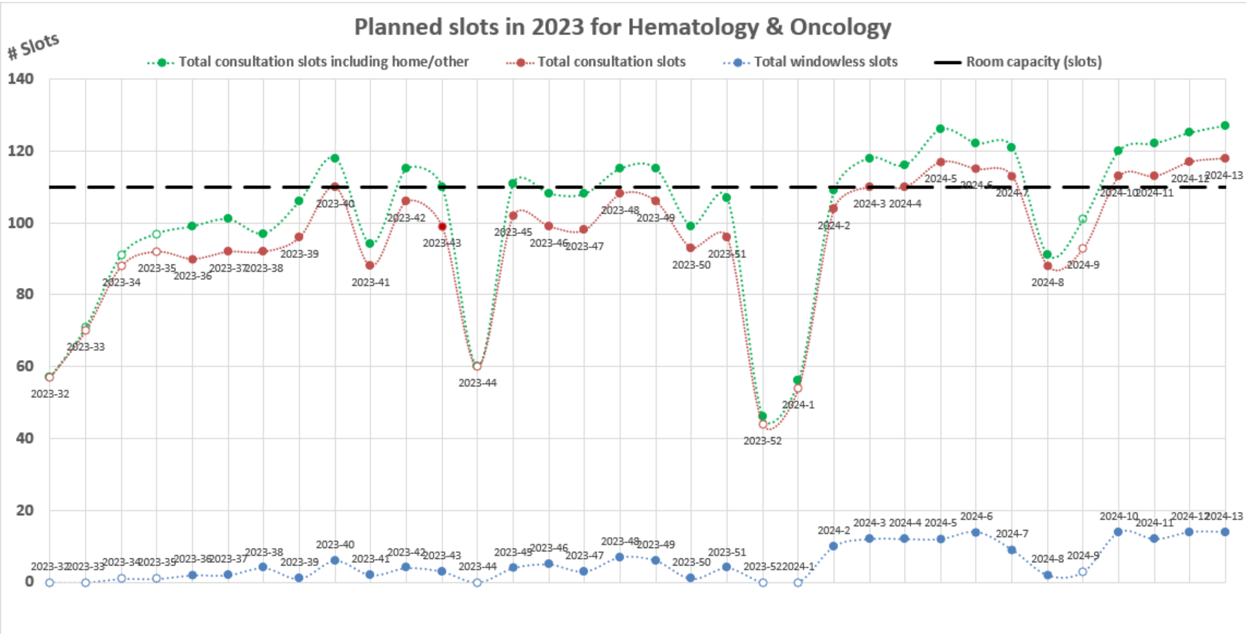

Figure 6C: Number of slots planned from 2023-32 to 2024-13 for H&O

Additional slots that were scheduled: Several administrative and other non-consultation slots were scheduled almost every week of the year within GIM. The number of administrative slots averaged 12, and the number of additional slots averaged 10 from week 1 to week 37, and 0.33 in the remaining weeks. The green and orange graphs exceed the capacity in additional weeks, whereas the red graph does not. This further complicates the planning. For H&O, several slots were scheduled at home or in other locations in most weeks. This is actually beneficial to the planning, since it frees up room slots in the consultation rooms in Zwolle. The average number of home/other slots was 6.7.

Average consultation slot duration: From the available data, we can determine the total number of hours that were planned during each week. If we divide this number by the total number of slots planned in that week, we can determine the average duration of a slot in that week. For GIM, the mean average duration throughout the year is 1 hour and 45 minutes, with the 25th percentile at 1 hour 42 minutes and the 75h percentile at 1 hour and 50 minutes. For hematology these numbers are a mean of 1 hour and 36 minutes, 25th percentile at 1 hour and 25 minutes, and 75th percentile at 1 hour and 36 minutes.

Errors in planning: Several planning errors were observed for week 7 of GIM. We noticed three instances of employees being planned into two different rooms on the same day, in addition to many spelling errors and variations in the formatting of names.

Relationship between total slots and windowless slots

The last step for this experiment is to determine whether the current planning method might be able to prevent inadequate rooms if we manage to decrease the total number of slots. To accomplish this, we can compare the number of windowless slots to the *total* number of slots in the room schedules, and determine the relationship between the two. The results can be seen in Figures 6D and 6E. The equations in the graphs can be used to calculate the expected number of inadequate slots in the analyzed room schedules, given the number of total slots. These allow us to estimate the number of windowless slots from the total number of slots given the current planning approach.

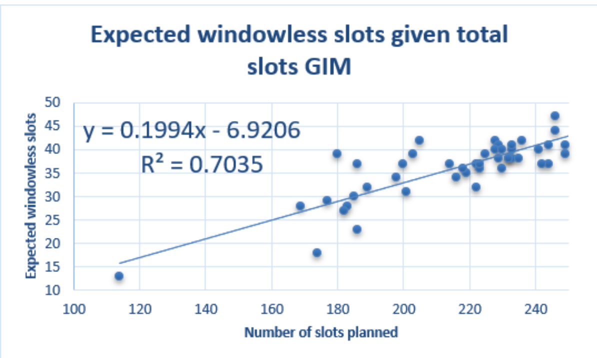

Figure 6D: Relation between the number of slots and the number of windowless slots in GIM

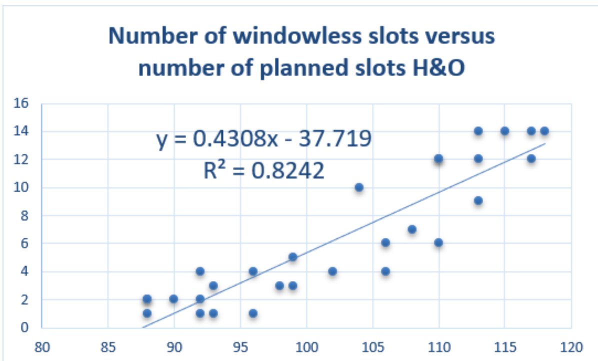

Figure 6E: Relation between the number of slots and the number of windowless slots in H&O

6.4 Experiment 4: New consultation durations with current planning

We recalculated the number of slots that would have been scheduled in each week for every employee given the established standard durations (Section 4.3). The results following this calculation are shown in Figure 6F and 6G. The boxes in the figures show the interquartile range (IQR) for all weeks of 2023, while the cross shows the mean. The first box shows the results for the current planning, whereas the other boxes show the results for various values of the standard duration. The line shows the capacity in terms of slots.

Performance evaluation

Number of consultations planned: The averages and maxima for both sub-departments on the number of consultations planned are shown in Table 6g. For GIM the inclusion of standard consultation durations of least two hours and fifteen minutes will decrease the expected total slots, while the same behaviour occurs for H&O from a standard duration of two hours.

Number of windowless consultations planned: We estimate the number of windowless slots using the calculated total number of slots and the relationships shown in Figures 6D and 6E. The ranges are shown in Figures 6H and 6I. We see that the expected number of windowless slots decreases for both sub-departments for a standard consultation duration of two hours or more. The averages and maxima for both sub-departments are also shown in Table 6g.

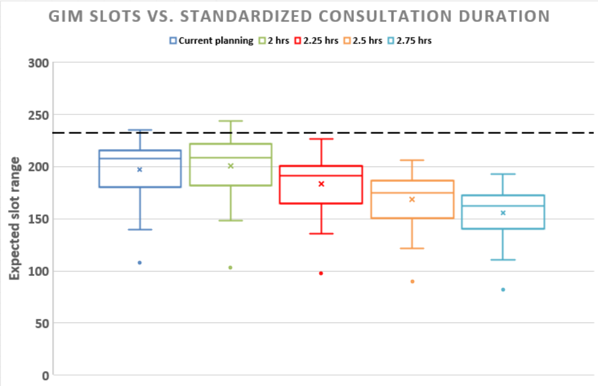

Figure 6F: Expected range of number of slots given consultation durations for GIM in 2023

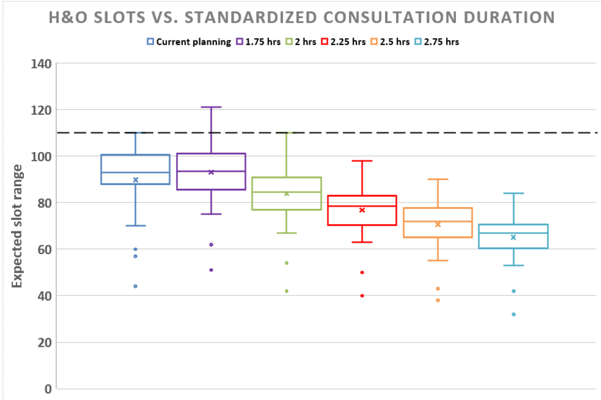

Figure 6G: Expected range of number of slots given consultation durations for H&O in 2023

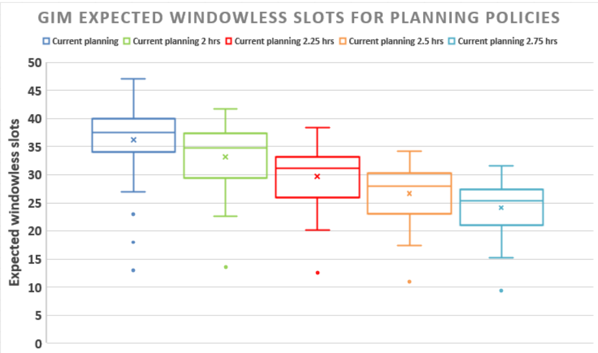

Figure 6H: Expected windowless slots for GIM given standard consultation durations

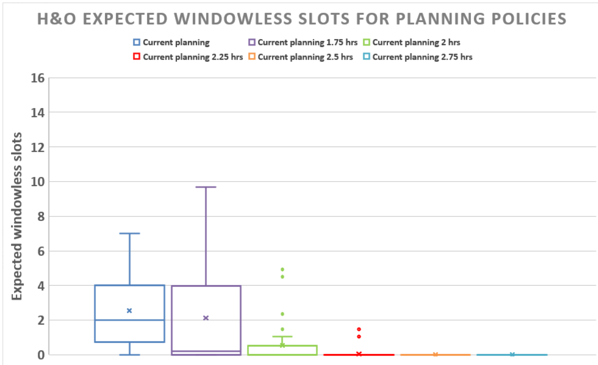

Figure 6I: Expected windowless slots for H&O given standard consultation durations

GIM					
Standard duration (hours)	**-**	**2**	**2.25**	**2.5**	**2.75**
Avg. slots	-	201.0	183.5	168.5	155.8
Max slots	-	244.0	227.0	206.0	193.0
Est. avg. windowless slots	-	33.2	29.7	26.7	24.2
Est. max. windowless slots	-	41.7	38.3	34.2	31.6
H&O					
Standard duration (hours)	**1.75**	**2**	**2.25**	**2.5**	**2.75**
Avg. slots	93.1	83.9	76.8	70.6	65.2
Max slots	121.0	110.0	98.0	90.0	84.0
Est. avg. windowless slots	3.6	1.2	0.3	0.0	0.0
Est. max windowless slots	14.4	9.7	4.5	1.1	0.0

Table 6g: Expected average and maximum total/windowless consultations for GIM and H&O given the total number of slots and current room planning regression

6.5 Experiment 5: new planning using SA

The total slots to plan for each employee are taken from the current room schedules. We use this as input to create a new planning for each week. Figures 6J and 6K show the number of windowless slots in for each of the weeks we planned. The values for 2024-1 to 2024-13 of H&O are estimates, indicated by the x markers. We explain the reasons below.

Performance evaluation

Number of windowless consultations planned: We now see that the number of windowless slots is either zero, or exactly equal to the number of planned slots minus the capacity for both GIM and H&O across every week of 2023. This means that the heuristic is able to use the minimum number of windowless slots. This minimum number is zero in most weeks for GIM, except weeks 12 and 47, where it is 5 and 3 respectively. From this we can determine that the average number of windowless consultations is 0.17 over the 48 evaluated weeks. The maximum is 5.

For H&O, the number of windowless slots is zero throughout 2023. From these results we predict similar results for the remaining weeks in H&O in 2024, which suggests that the number of windowless slots will be between 3 and 8 for weeks 5, 6, 7, and 10-13 of 2024. The average number of windowless slots in 2023 is zero, the average number of windowless slots across all weeks is 1.72. The maximum is 8.

Room switches & counts of room sharing When rescheduling week 7 of 2023 for GIM, we see that the number of room switches decreases to 4, and the number of room shares decreased to 44. This is a decrease of 91.5% and 25.4% respectively with respect to the current schedule in this week. When rescheduling week 43 of 2023 for H&O, we see that the number of room switches decreases to 1, and the number of room shares decreased to 11. This is a decrease of 92.3% and 47.6% respectively with respect to the current schedule in this week. In week 16 for GIM and week 50 for H&O the numbers are similarly low, while the number of switches and shares does increase when the number of slots is either at or over available capacity in week 12 for GIM and week 40 for H&O.

GIM	**Week 7**	**Week 12**	**Week 16**
Number of slots	225	235	215
Number of times any doctor switches	4	14	4
Number of shared rooms	44	50	35
Number of windowless slots	0	5	0
H&O	**Week 40**	**Week 43**	**Week 50**
Number of slots	110	99	93
Number of times any doctor switches	12	1	0
Number of shared rooms	18	11	8
Number of windowless slots	0	0	0

Table 6h: Weekly Planning Metrics for Room Utilization and Staff Allocation

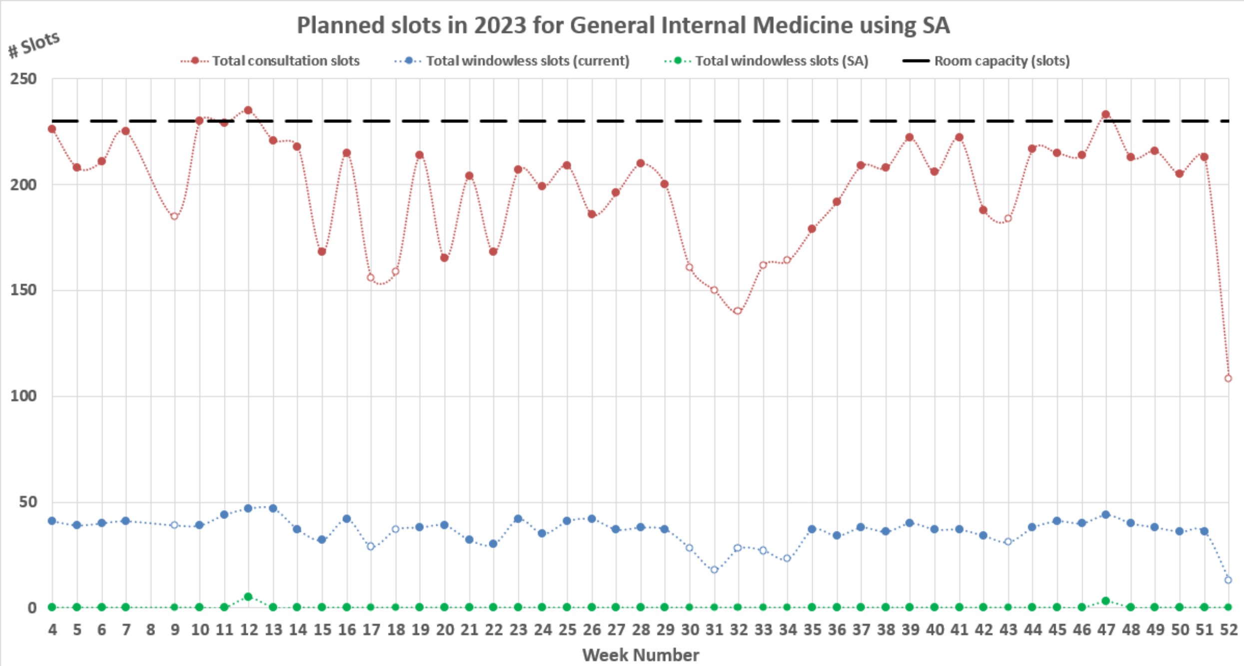

Figure 6J: Number of slots planned within 2023 for GIM including SA data

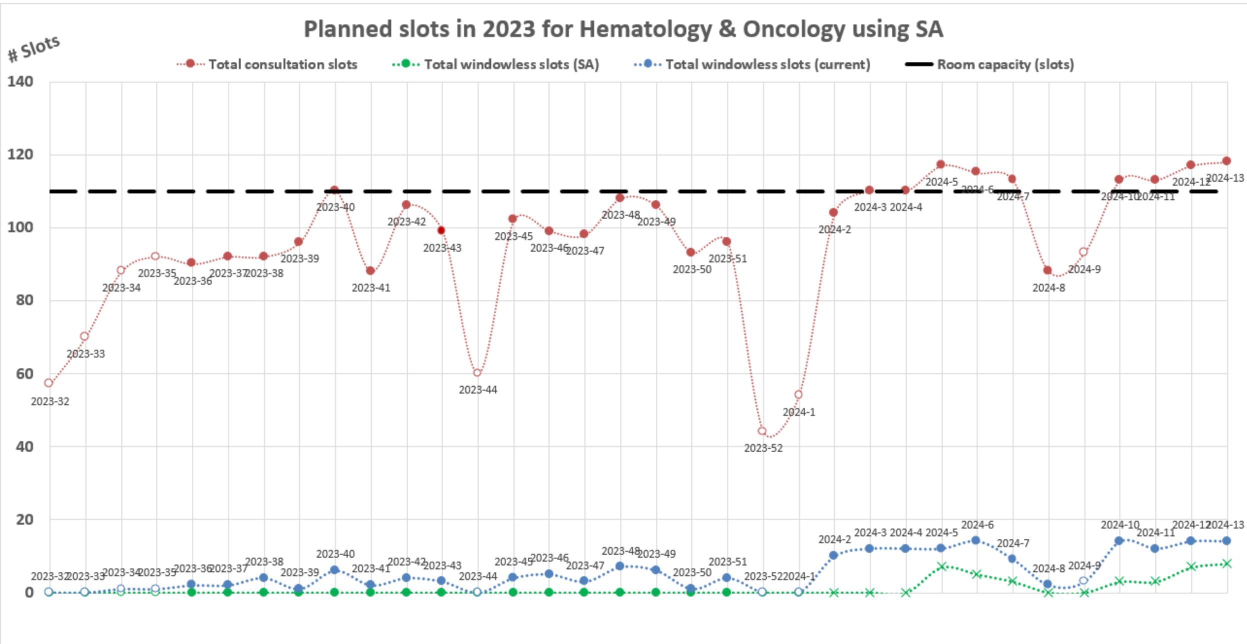

Figure 6K: Number of slots planned from 2023-32 to 2024-13 for H&O including SA data

6.6 Experiment 6: Combination of policies and SA scheduling

In the previous two experiments we identified the individual effects of SA planning and standard consultation durations. This section shows the performance when combining both approaches.

Performance evaluation

Number of windowless consultations planned: Given a standard duration of two hours, the only weeks with slots in windowless rooms are those with a higher number of slots than the available capacity. This indicates that any time the number of slots is below 230 for GIM and below 110 for H&O, the number of windowless slots will be zero for longer standard durations as well. From this we can make an estimate for the average and maximum number of windowless slots for each standard consultation duration, shown in Table 6i.

Room switches & counts of room sharing The counts of room switching and sharing for this experiment can be seen in Table 6j. In week 7, we obtained 4 switches and 46 counts of sharing for a standard duration of two hours for GIM, where results improve slightly for longer standard durations, to a minimum of zero and 42 for a duration of two hours and 45 minutes. In week 43, we obtained zero switches and six counts of sharing using a standard duration of two hours or more.

	GIM			
Standard duration (hours)	**2**	**2.25**	**2.5**	**2.75**
Avg. slots	201.0	183.5	168.5	155.8
Max slots	244	227	206	193
Est. avg. windowless slots	0.8	0.0	0.0	0.0
Est. max windowless slots	14	0	0	0
	H&O			
Standard duration (hours)	**2**	**2.25**	**2.5**	**2.75**
Avg. slots	83.9	76.8	70.6	65.2
Max slots	110	98	90	84
Est. avg. windowless slots	0.0	0.0	0.0	0.0
Est. max windowless slots	0	0	0	0

Table 6i: Expected average and maximum windowless slots for GIM given the total number of slots and current room planning regression

GIM Week 7				
Standard duration (hours)	**2**	**2.25**	**2.5**	**2.75**
Number of slots	225	205	189	178
Number of windowless slots	0	0	0	0
Number of times any doctor switches	4	1	1	0
Number of shared rooms	46	43	43	42
H&O Week 43				
Standard duration (hours)	**2**	**2.25**	**2.5**	**2.75**
Number of slots	70	63	59	54
Number of windowless slots	0	0	0	0
Number of times any doctor switches	0	0	0	0
Number of shared rooms	6	6	6	6

Table 6j: Weekly Planning Metrics for Room Utilization and Staff Allocation

6.7 Results analysis

In this section we discuss the performance on each evaluation criterion on the basis of all the various experiments, and extract insights from the performance variation. All results for the room schedules are summarized in Table 6l, and the results for the suggested changes in the capacity allocation over the locations are shown in Table 6f.

Percentage of patients treated in the closest location

The percentage of patients treated in the closest location are for each individual location (Zwolle, Meppel, Kampen, Heerde, and Steenwijk) 97.4%, 58.4%, 25.2%, 12.5%, and 18.7%. These numbers represent the hours of physical demand scheduled to those locations. In addition to this physical demand, however, the department is also scheduling consultations that are demand for other locations, as well as telephonic consultations, to each location. The yearly hours of physical demand from other locations are, respectively, 1796, 127, 90, and 96 hours. The yearly hours of telephonic consultations are 2125, 447, 366, and 193 for Meppel, Kampen, Heerde and Steenwijk respectively. Finally, the average duration of consultations in the satellite locations is lower than it could be. Based on these observations, we analyze the effects of mediating these issues. These approaches can all be executed without allocating any new slots to the satellite locations. The options are:

- to increase the percentage of physical demand among the currently planned physical appointments (if planned work all closest);
- to replace telephonic appointments with physical demand (if no telephonic);
- to increase the average duration of the consultations with physical demand, e.g. to 2.5 hrs (if avg. duration 2.5hrs).

These effects are expressed in terms of the maximum expected change, yet there may be many reasons why this can not be realized. We discussed before, for example, that some patients may not want to travel to Zwolle even though it is the closest location, or vice versa. We also discussed the fact that employees prefer doing at least a small number of telephonic appointments during a consultation, which the department may allow. An average duration of 2.5 hours may also not be achievable due to limited available demand for the specialties and professions in the current capacity allocation. Regardless, improvements can certainly be expected.

Further improvements can be generated by planning additional slots according to the slot configuration shown in Table 6c. Each additional slot planned following this configuration will increase the percentage of correctly treated patients within the department by 0.5%, if the appointments in the consultation add up to 2.5 hours of physical demand for that location.

All potential decisions that the department can make and the maximum expected effects have been shown in 6k. The second to last column shows the maximum achievable percentage given the current capacity allocation. This can be achieved without planning any additional slots. The maximum achievable percentage given the new capacity allocation is shown in the last column, which can be attained by planning additional slots.

Location	**Total demand**	**% currently treated**	**If planned work all closest**	**If no tele-phonic**	**If avg. duration 2.5hrs**	**Max current alloc.**	**Max new alloc.**
Zwolle	8417	94.5%	+5.5%	-	-	100%	100%
Meppel	4043	58.4%	+41.6%	-	-	100%	100%
Of total	**12460/22242**	**46.4%**	**+9.6%**	-	-	**56.0%**	**56.0%**
Kampen	3992	25.2%	+3.2%	+11.2%	+32.0%	71.6%	88.9%
Heerde	3417	12.5%	+2.6%	+10.7%	+18.0%	43.8%	91.2%
Steenwijk	2374	18.7%	+4.0%	+8.1%	+18.5%	49.3%	88.3%
Of total	**9782/22242**	**8.4%**	**+1.4%**	**+4.5%**	**+10.5%**	**24.8%**	**39.4%**
Total	**22242**	**54.7%**	**+11.0%**	**+4.5%**	**+10.5%**	**80.9%**	**95.3%**

Table 6k: Total percentage of physical demand that can be planned in each location given different policies

Total number of slots planned

The total number of slots planned for the current room schedules average 197.2 and 95.6 for GIM and H&O. We expect this number to decrease when policies are implemented that increase the average duration of consultations. To accomplish this, the department can look into policies such as the one outlined in the experiment of Sections 5.4 and 6.4. Using this policy with a standard duration of two hours and thirty minutes will decrease the average number of consultations in GIM from 197.2 to 168.5, and decrease the maximum number of consultations from 235 to 206. For H&O the effects of the same duration are to decrease the average from 95.6 to 70.6, and the maximum from 118 to 90.

Although this policy will decrease the total number of consultations, it will actually increase it for employees who currently work more than two hours and thirty minutes on average per consultation. 5.1 hours, for example, will allocate three slots to an employee. Because of this, we recommend that the department discuss the number of consultations with all employees to find a feasible number of consultations per individual.

Treating more patients in the closest location instead of in Zwolle will also decrease the number of consultations in Zwolle. These consequences are simpler; every slot planned in a different location is one fewer weekly slot in Zwolle. The maximum number of weekly slots that can be scheduled to the satellite locations while maintaining the total workload in Meppel is 38.5, which represents an average decrease of the number of slots in Zwolle of, again, 38.5.

Total number of windowless slots planned

The total number of windowless slots depends on two factors. The first of these factors is the total number of slots, and the second the planning method. Under the old planning method, we expect the number of windowless slots to decrease with the total number of slots according to the equations in Figures 6D and 6E. These suggest that each decreased total slot decreases the windowless slots in GIM by 0.2, and in H&O by 0.43. From these equations we expect the number of windowless slots to be zero for GIM from 34 total consultations, and 87 total consultations for H&O.

For H&O, windowless consultations can be avoided by implementing policies that increase the average consultation duration. Using a standard duration of 2.75 hours for H&O, for example, would increase the average consultation duration to 2.34 hours and decrease the number of windowless slots to zero. The equation for GIM predicts zero windowless slots from 34 total slots, which is infeasible given the required care. However, we should consider that the current schedules are created using a blueprint based on the current number of consultation slots. Decreasing the number of slots as a whole would also allow the creation of a new blueprint that is based on fewer slots, which could further improve the schedules.

New planning methods can decrease the number of windowless slots even further. Given the current number of slots and SA planning, the average number of windowless slots in 2023 for GIM and H&O reduced to 0.17 and 1.72 across the evaluated weeks, with respective maxima of 5 and 8. Combining new slots with new planning methods, we expect the number of windowless slots to decrease to zero with policies that generate even small increases in the average slot duration. Using SA in combination with a standard duration of two hours and fifteen minutes for GIM and two hours for H&O, which would increase the average consultation durations to 1.89 hours and 1.82 hours respectively, would already generate zero windowless slots for both sub-departments.

The results suggest that planning is a major factor in generating windowless slots. We discussed before that this may largely be due to the complexities of the scheduling process. This complexity is inevitably higher for GIM than for H&O, given the significantly larger number of employees typically involved in the planning (39-67 for GIM and 13-24 for H&O), as well as the larger number of rooms involved (23 windowed rooms for GIM and 11 for H&O). This suggests that looking into better planning rules for both departments might significantly benefit the room schedule on windowless room scheduling. We also suggest that the GIM sub-department discuss their planning methods with the H&O sub-department to see if they might have valuable insights, given the lower number of windowless slots in H&O.

Scheduling additional consultations to the satellite locations may also decrease the number of windowless slots through a reduction in total slots.

Room switches & counts of room sharing

The total number of room switches in week 7 of 2023 for GIM was 47 for the current planning, 4 when creating a new planning with SA, and could be decreased further to 1, 1 and 0 when using standard durations of 2.25, 2.5, and 2.75 hours in combination with SA. For H&O, the number of room switches in week 43 of 2023 was 13 for the current planning, 1 when creating a new planning with SA, and 0 when combining SA planning with standard durations of 2 hours or higher.

The total number of room shares in week 7 of 2023 for GIM was 59 for the current planning, 44 when creating a new planning with SA, and could be decreased further to 43, 43 and 42 when using standard durations of 2.25, 2.5, and 2.75 hours in combination with SA. For H&O, the number of room shares in week 43 of 2023 was 21 for the current planning, 11 when creating a new planning with SA, and 6 when combining SA planning with standard durations of 2 hours or higher.

We see that with a different planning method, both room switching and room sharing decrease significantly, regardless of a standard consultation duration. This suggests that room switches and shares for both sub-departments are mainly caused by the planning complexity, and can be decreased significantly with better planning methods.

Summary all room policies

All policies that can be implemented for the room schedules have been summarized in Table 6l.

Evaluation criterion	**GIM**	**H&O**
Number of slots (avg)		
Current schedule	197.2	95.6
Standard slot duration (2.5hrs)	168.5	70.6
SA rescheduling	-	-
Standard slot duration (2.5hrs) + SA	-	-
Number of slots (max)		
Current schedule	235	118
Standard slot duration (2.5hrs)	206	90
SA rescheduling	-	-
Standard slot duration (2.5hrs) + SA	-	-
Number of windowless slots (avg)		
Current schedule	36.2	5.2
Standard slot duration (2.5hrs)	26.7	0.0
SA rescheduling	0.17	1.71
Standard slot duration (2.5hrs) + SA	0.0	0.0
Number of windowless slots (max)		
Current schedule	47	14
Standard slot duration (2.5hrs)	34.2	1.1
SA rescheduling	5	8
Standard slot duration (2.5hrs) + SA	0	0
Number of switches		
Current schedule	47 (week 7)	13 (week 43)
Standard slot duration (2.5hrs)	-	-
SA rescheduling	4 (week 7)	1 (week 43)
Standard slot duration (2.5hrs) + SA	1 (week 7)	0 (week 43)
Number of shares		
Current schedule	59 (week 7)	21 (week 43)
Standard slot duration (2.5hrs)	-	-
SA rescheduling	44 (week 7)	11 (week 43)
Standard slot duration (2.5hrs) + SA	43 (week 7)	6 (week 43)

Table 6l: Performance of each planning approach on each evaluation criterion

7 Conclusions, recommendations, discussion, and future research

This chapter forms the culmination of our research. We summarize our results in the conclusions of Section 7.1, based on which we make recommendations in Section 7.2. In Section 7.3 we discuss potential limitations to our research and results, and if possible provide counterarguments and suggestions for mitigation of the limitations. In Section 7.4 we describe the broader applicability of our research, and make suggestions for future research.

7.1 Conclusions

In this thesis, we investigated two aspects of the department's planning. The first goal was to evaluate the department's current performance with respect to scheduling appointments to the closest locations. We accomplished this by analyzing the department appointment data from 2023, and by developing a method to determine the closest location for each patient based on travel distance. The results show that most of the patients living close to Zwolle are treated in Zwolle (94.5%), while this holds for only 58.4% of the patients living closest to Meppel, and for 25.2%, 12.5%, and 18.7% for Kampen, Heerde, and Steenwijk, respectively. The total percentage of patients treated close to home for the department is 54.7%.

Our analysis shows that the department can increase this percentage in multiple ways. First, given the available capacity, the department can increase the percentage by up to an additional 9.6% by more accurately planning the right appointments in Zwolle and Meppel. For the remaining three locations, the department can also more accurately plan the right appointments (+1.4%), replace telephonic appointments with physical demand (+5.5%), or increase the duration of consultations (+10.5%). Without changing the current capacity, the department can expect to handle at most 80.9% of their appointments in the closest locations. Up to another 14.4% can be gained by planning more consultations in the satellite locations according to the configuration of Table 6c. The maximum percentage given all these changes is 95.3%. Each slot added according to the Table increases the percentage of patients treated in the closest location by an expected 0.5%, assuming 2.5 hours of physical patient care per slot. All of these values have been summarized in Table 6k.

The second goal of our research was to investigate methods for improving room planning in Zwolle in terms of windowless room use, room switching, and room sharing. The number of windowless slots used in current weeks averages 36.2 (GIM) and 5.2 (H&O) over the evaluated weeks. Room switching and counts of room sharing were investigated in week 7 of GIM, and week 43 of H&O. There were 47 and 13 counts of room switches, and 59 and 21 counts of room shares for the sub-departments, respectively. The current room schedule for GIM also showed errors in terms of employees being planned in multiple rooms for a single timeslot. Furthermore, the number of windowless consultations was shown to decrease with a respective 0.2 and 0.43 with each point decrease in the total consultations.

We investigated the effects of two approaches. The first was to decrease the total number of consultations by implementing a standard duration for consultations. We demonstrate that implementing a standard duration of at least two hours and fifteen minutes for GIM and two hours for H&O decreases the number of consultations in the respective busiest weeks to below the capacity limits. A standard duration of two hours and thirty minutes for the respective sub-departments decreases the average number of windowless slots by an estimated 9.5 and 5.2, to respectively 26.7 and 0.0. and the maxima by 12.8 and 12.9, to 34.2 and 1.1.

The second approach uses Simulated Annealing (SA) to reorganize the available planning, without decreasing the total slots. This approach reduced the average number of windowless slots to 0.17 and 1.71, with maxima of 5 and 8. The approach also reduced the number of room switches by 43 and 12, to 4 and 1. The count of room sharing was decreased by 15 and 10, to 44 and 11. Combining both approaches reduced windowless slots even further, to zero for both sub-departments. Room switches were also reduced further, to one and zero. The count of room shares was reduced by an additional 1 and 5, to 43 and 6.

7.2 Recommendations

How should employee capacity be allocated across the Isala locations to ensure that patients receive treatment as close to their homes as possible?

Based on our findings surrounding the first goal, we *recommend* that the department investigate the applicability of the policies in Table 6k under the current capacity allocation. We *suggest* investigating whether the physical appointments that are not in the determined demand for the satellite locations are being scheduled to those locations for viable reasons. We *recommend* discussing the possibility of reducing the number of telephonic appointments in the satellite locations with employees. Additionally, we *recommend* that the department looks into setting a minimal requirement on the duration of consultations in the satellite locations; we recommend a minimum of two hours. Lastly, we *recommend* that the department uses the configuration shown in Table 6c to guide future decisions on allocating employees to the different locations. We recommend that the department remains critical whenever scheduling these consultations. Firstly, the department should evaluate the allocation to observe whether the calculated demand truly manifests, and keep track of yearly changes in patient demand to determine the retained validity of the allocation. We also recommend that the department incorporate additional data dimensions into its system, specifically the closest appointment location and the appointment specialty. For support, the department can refer to Section 4.2.2 and Appendix A. This will directly allow the analysis of this thesis using their PowerBI tools.

What are the strategies that the Internal Medicine department should consider in order to improve the room planning in Zwolle?

Based on our findings surrounding the second goal, we *recommend* that the department implements policies to avoid consultations that last for shorter than two hours for both sub-departments, while facilitating even longer durations. We also *recommend* that the department looks into new planning policies in three stages. The first stage is to encourage a discussion between the planning team of both sub-departments to brainstorm ways for improving the planning. The second stage is to evaluate policy-based planning approaches, such as sequentially planning employees in the first available slot of the week according to availability, starting with employees with the highest number of required consultations. This algorithm has been outlined in Section 4.4.4. The third stage is to consult with the capacity team to implement a planning method comparable to the SA planning outlined in this thesis. The only requirements to enable this planning approach are knowledge of all employee availability restrictions, and access to required software.

The department will also be able to decrease the total number of slots by scheduling more consultations to the satellite locations. The department may also consider reducing the number of slots scheduled in consultation rooms by scheduling some consultations at home. Nevertheless, we do not suggest either of these as a method that the department should prioritize to improve the planning specifically.

7.3 Research limitations

Now that we have presented our results, we will discuss factors that might limit their veracity. For each of them, we discuss why they may or may not have caused inaccuracies in our results, and if necessary how the department should consider them when implementing our recommendations.

Capacity allocation Hematology The most notable limitation to our research is in the calculated slots for Hematology. These slots have been calculated assuming that all considered appointments can be scheduled in the satellite locations, but due to limitations material limitations surrounding blood testing and certain therapies this is not the case for all appointments. The exact percentages could not be retrieved from the data, since these appointments are not tracked in the HiX data. Rough estimates by the H&O sub-department are that this limitation holds for approximately 15-20% of the appointments, although it may vary significantly per profession.

We have chosen to leave this consideration out of all of our calculations, because this allows the department to scale the numbers according to the exact percentages in the future. Regardless, the department should keep this fact in mind when decisions are made to allocate hematology slots to any of the satellite locations.

Data validation The data analysis is based directly on appointment data from the department and all steps in the data processing have been thoroughly researched and validated in consultation with the department. We describe all steps that we have taken, as well as all assumptions made, in Appendix A.

Inaccurate location selection In Section 4.2.2 we discuss how we determine the closest location for each appointment. We mention there that we determine the ZIP codes of patients only on the basis of the numbers, which sometimes covers quite a significant area. Some districts might be entirely allocated to one location, whereas a percentage of them may live closer to a different location. Nevertheless, taking the center of a district will never wrongfully allocate more than half of a district. We suspect that there may be very minor inaccuracies in the demand calculations per locations, although we altogether think most inaccuracies average out across the locations.

Limited evaluation diagnosis In Section 2.1.3 we discussed the distribution of diagnoses across specialties based on the study by van der Wel, (2023). This study only covered the 90% most frequent diagnoses, and the remaining 10% may therefore mean that the demand per specialty is represented slightly inaccurately. Nevertheless, less than half of the 90% most frequent diagnoses could be treated by multiple specialties, and this situation was mainly limited to the *most* common diagnoses. From the relative rarity of the remaining diagnoses, we hypothesize that they are more specialized and typically treated by the right specialty as represented in the appointment data. In the worst case scenario, at most half of 10% of the demand of specialties may be wrongfully allocated. Nevertheless, inaccuracies are only limited to how workload is distributed across the *specialties*, and not across the *locations*. Given the fact that the distribution of care across the specialties does not currently match the desired distribution, the effects of this inaccuracy will be insignificant in the short term.

Yearly demand changes The results of our thesis have been based on the demand in 2023, and for them to remain applicable the demand within the department should not change significantly. In practice we do expect the demand to be stable, since national policy dictates that hospitals can not grow their total care output in terms of the number of DBC's treated. Nevertheless, the department may increase or decrease time spent on specific DBC's, the distribution between telephonic and physical appointments may change, and the distribution of care across the specialties may change. In the available data, for instance, we do see an increase in physical demand from 2021 to 2023, as can be seen in Figure 7A. We recommend the department to monitor the changes in their workload, and to scale the number of hours shown in Table 6b when significant changes emerge with respect to 2023.

Unpredictable patient behaviour The foundation for the capacity allocation to locations is built upon the ideas of the IZA, which proposes that patients should be treated as close to home as

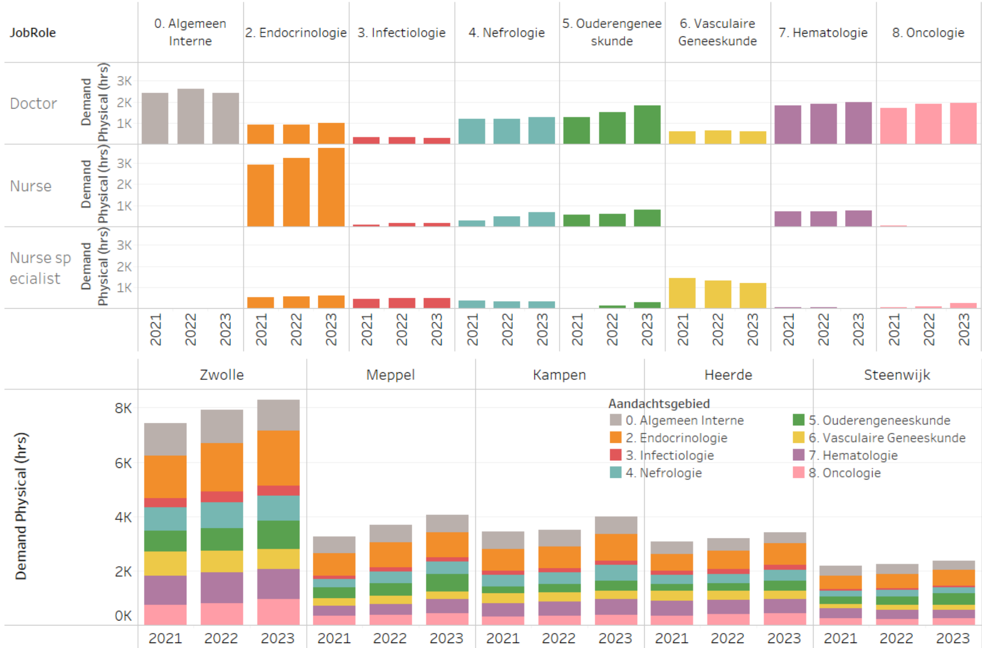

Figure 7A: How physical demand changed from 2021-2023

possible. We have also investigated whether the scientific literature suggests benefits for patients that can access care closer to home. Both of these suggest broad-scale benefits for treating patients closer to home, but individual preferences may still vary especially in specific hospitals or locations. It might be, for instance, that the environment in certain locations is much more, or less, desirable. There are two main ways to really investigate the applicability of this idea to Isala itself. The first is an extensive patient inquiry that compares the closest location for patients to the location where they would like to be treated. The second is to base ourselves in the strategic goals from IZA without accounting for this. This means incrementally implementing the complete allocation as we have calculated it, and evaluate its performance operationally. The latter is what we suggest the department does.

Variability in weekly demand In order to calculate the required number of hours in Table 6b, we have assumed that demand is evenly spread out throughout the weeks. Demand fluctuates, however, and we discuss the sources of this demand fluctuation below. The first is the reduction weeks, as discussed in Section 2.1.6. During these weeks, fewer patients can be scheduled due to employee absence. The second of these is doctor absence outside of reduction weeks due to conferences, additional holidays, or personal circumstances. We can deal with both of these sources of variability using effective communication with the employees. When employees know that they will be absent, they can communicate this with the planning department, and their slots can be removed from the schedule. If a doctor that works one slot in Kampen is absent, we require one fewer slot in Kampen for that week.

Another source of variability is inherent planning variability. In absence of clear ranges for the number of patients to schedule in a week, a given week can have many more appointments scheduled for a specialty than another. This can create inconsistent workloads, especially when this is done for new patients. Any peaks or dips in new patient scheduling will affect variability in future return patients through the bullwhip effect[35]. The last source of variability is due to inherent patient demand fluctuations.

We suggest that the department address the latter two at the same time. The variability in patient demand can be offset through planning patients from weeks with higher demand to weeks with lower demand instead. None of this will completely avoid variability, but doing so will significantly reduce it. We recommend that any remaining dips in the physical demand are balanced out by scheduling additional telephonic appointments.

To provide insight into the current effects of variability, we experimented with a hypothetical planning using the calculated number of hours from Table 6b. If these hours had been scheduled to each location in 2023, we would have treated between 90% and 95% of patients in each week depending on the specialty. This is a lower-bound estimate, since the planning department can account for the number of patients planned per week if capacity is known, as discussed above.

Increased physical demand as a result of proposed changes The choice of either a telephonic or a physical appointment is largely made through a discussion between the doctor and patient, where patients who specifically request physical appointments will likely be granted one. Facilitating care closer to home may influence the likelihood of patients requesting physical appointments, and thereby increase the number of physical appointments. This may affect physical demand in locations in the future, which the department should monitor.

Excluded professions In our location capacity allocation results specifically, Section 6, we excluded the work done by doctors' assistants, co-assistants, and physician's assistants. We have directly accounted for this in our results. We have not accounted for the hours by psychologists, psycho-diagnostic employees, research nurses, and dietitians. The number of hours for these employees is limited, but they should be considered upon real-life implementation.

Specific unavailability of employees In the creation of the room schedules, we do not directly account for employee preferences and unavailability. Although we have built the program such that this *can* be included, the exact preferences and unavailability are not explicitly available in a central unifying document. Instead, this knowledge is spread over the department drive and partially known to the department. As such, the complex preferences and unavailability of employees might slightly decrease the quality of the SA schedule if properly incorporated. To estimate the effects of this, we implemented a procedure within our SA program that would randomly allocate up to three slots of unavailability to employees. We found only minor differences in room schedule quality regarding room sharing and switching, and no differences in the number of windowless slots.

Windowless slots estimated The calculated number of windowless slots in Experiment 4, where we use policies to reduce the total number of slots, has been estimated using the relationship determined in Figures 6D and 6E using the current schedules. Although these relationships reflect the relationship between the total number of slots and the total number of windowless slots in the available schedules, their applicability to future schedules may be limited. The main reason for this is that the number of windowless slots is also dependent on the blueprint used to create the schedules, and if this blueprint changes due to new policies, then performance may be different. We do expect the performance to be better, however, as a blueprint with fewer consultation slots should be less complicated to plan. This, indeed, has been one of the goals of the research.

Additionally, we estimated the number of windowless slots by combining standard durations and the SA scheduling approach. This was based on a large number of runs that showed specific behaviour, but there is a minor chance that certain weeks could have resulted in more windowless slots than currently shown. We do consider this very unlikely, given the observed behaviour of the approach.

7.4 Broader applicability and future research

The aims of our research draw on the broader strategic goals of the Integral Care Agreement (IZA). Effective operations management has the potential to improve many different aspects of healthcare delivery, such as reduced access times and improved quality of care. Our research that incorporating patient proximity considerations into the management for a single department may significantly reduce travel times for patients, with established effects on both patient health outcomes and the likelihood of patients accessing care. Expanding this analysis to the full hospital could allow optimal use of the various hospital locations, and could have additional benefits resulting from the ensuing standardization and simplification.

On a national scale, basing the capacity of each hospital in the Netherlands on the localized demand could have even more significant benefits. We hypothesize that specialists from one location may be better suited working in a different location. We see this reflected in the internal medicine department, where the infectious disease workload far exceeds the specialty's capacity, such that those specialists mainly treat patients from other specialties. Such an endeavour on a national scale requires effective planning in addition to demand calculations, for which ideas can be derived from our room planning approach. Simulated Annealing is particularly effective at multi-objective optimization, and may be able to balance the need for care close to home with specialist availability and other hospital-specific requirements.

We recommend that future research continues by developing a planning approach that may ensure the calculated allocation of slots across the locations within the department. Given operationality, this may even be expanded to include the entire hospital. Doing so would allow optimal use of the satellite locations, since it could assess which departments could make the best use of the available capacity. We see this thesis as a proof of concept for expanding the objectives of hospitals to include patient proximity directly, and to show the utility of doing so. We see the room planning approach furthermore as a stepping stone towards future implementation.

References

[1] Centraal Bureau voor de Statistiek. *Regionale bevolkings- en huishoudensprognose 2022–2050: Steden en randgemeenten groeien verder*. Prepress: Textcetera, Den Haag en CCN Creatie, Den Haag; Ontwerp: Edenspiekermann. Den Haag, Heerlen, Bonaire, 2022. URL: `https://www.cbs.nl/infoservice`.

[2] Ministry of VWS. *Integraal Zorg Akkoord*. Accessed: 2024-08-14. 2022. URL: `https://www.rijksoverheid.nl/binaries/rijksoverheid/documenten/rapporten/2022/09/16/integraal-zorgakkoord-samen-werken-aan-gezonde-zorg/integraal-zorg-akkoord.pdf`.

[3] E.P. Mseke et al. "Impact of distance and/or travel time on healthcare service access in rural and remote areas: A scoping review". In: *Journal of Transport & Health* 37 (2024), p. 101819. ISSN: 2214-1405. DOI: `https://doi.org/10.1016/j.jth.2024.101819`. URL: `https://www.sciencedirect.com/science/article/pii/S2214140524000653`.

[4] Charlotte Kelly et al. "Are differences in travel time or distance to healthcare for adults in global north countries associated with an impact on health outcomes? A systematic review". In: *BMJ Open* 6.11 (2016). ISSN: 2044-6055. DOI: `10.1136/bmjopen-2016-013059`. eprint: `https://bmjopen.bmj.com/content/6/11/e013059.full.pdf`. URL: `https://bmjopen.bmj.com/content/6/11/e013059`.

[5] Kevin Rui-Han Teoh et al. "Doctors' working conditions, wellbeing and hospital quality of care: A multilevel analysis". In: *Safety Science* 135 (2021), p. 105115. ISSN: 0925-7535. DOI: `https://doi.org/10.1016/j.ssci.2020.105115`. URL: `https://www.sciencedirect.com/science/article/pii/S0925753520305129`.

[6] Roger S. Ulrich. "Biophilic theory and research for healthcare design". In: *Biophilic design: The theory, science, and practice of bringing buildings to life*. Ed. by Stephen R. Kellert et al. Hoboken, NJ: Wiley, 2008.

[7] E. R. C. M. Huisman et al. "Healing environment: A review of the impact of physical environmental factors on users". In: *Building and Environment* 58 (2012), pp. 70–80. DOI: `10.1016/j.buildenv.2012.06.016`. URL: `http://dx.doi.org/10.1016/j.buildenv.2012.06.016`.

[8] Mustafa K. Alimoglu et al. "Daylight Exposure and the Other Predictors of Burnout Among Nurses in a University Hospital". In: *International Journal of Nursing Studies* 42.5 (2005), pp. 549–555. DOI: `10.1016/j.ijnurstu.2004.09.001`.

[9] Mohamed Boubekri et al. "Impact of Windows and Daylight Exposure on Overall Health and Sleep Quality of Office Workers: A Case-Control Pilot Study". In: *Journal of Clinical Sleep Medicine* 10.6 (2014), pp. 603–611. DOI: `10.5664/jcsm.3780`.

[10] John E. Ware et al. "A 12-Item Short-Form Health Survey: Construction of scales and preliminary tests of reliability and validity". In: *Medical Care* 34.3 (1996), pp. 220–233. DOI: 10.1097/00005650-199603000-00003.

[11] Frank Burger. *Documentation on designing the Isala building.* Originally written in English for the International Health & Design Congress, Brisbane, Australia, July 2013; later translated into Dutch by the author and architect Frank Burger from AMI. 2013. URL: https://daf9627eib4jq.cloudfront.net/app/uploads/2017/01/attachment-ctdocumentatie-augustus-2013.pdf.

[12] H. Heerkens et al. *Solving Managerial Problems Systematically.* Publisher Name, 2021. Chap. 1.

[13] Stan van der Wel. *Determining the balance between general and specialist care of Internal Medicine practitioners based on patient demand.* 2023. URL: http://essay.utwente.nl/97735/.

[14] Symerio. *pgeocode.* Python library for high performance off-line querying of GPS coordinates. 2023. URL: https://github.com/symerio/pgeocode.

[15] Google. *Google Maps Platform: Distance Matrix API.* Accessed: 2024-11-15. Google. 2024. URL: https://developers.google.com/maps/documentation/distance-matrix.

[16] Brecht Cardoen et al. "Operating room planning and scheduling: A literature review". In: *European Journal of Operational Research* 201.3 (2010), pp. 921–932.

[17] Peter JH Hulshof et al. "Tactical resource allocation and elective patient admission planning in care pathways". In: *Artificial Intelligence in Medicine* 55.3 (2012), pp. 193–207.

[18] R. L. Graham et al. "Optimization and approximation in deterministic sequencing and scheduling: a survey". In: *Annals of Discrete Mathematics* 5 (1979), pp. 287–326.

[19] Ilyes Kacem et al. "Pareto-optimality approach for flexible job-shop scheduling problems: Hybridization of evolutionary algorithms and fuzzy logic". In: *Mathematics and Computers in Simulation* 60.3-5 (2002), pp. 245–276.

[20] Jully Jeunet et al. "Optimizing temporary work and overtime in the Time Cost Quality Trade-off Problem". In: *European Journal of Operational Research* 284.2 (2020), pp. 743–761. ISSN: 0377-2217. DOI: https://doi.org/10.1016/j.ejor.2020.01.013. URL: https://www.sciencedirect.com/science/article/pii/S0377221720300345.

[21] Andreas T. Ernst et al. "Staff scheduling and rostering: A review of applications, methods and models". In: *European Journal of Operational Research* 153.1 (2004), pp. 3–27.

[22] Edmund Burke et al. "Metaheuristics for Handling Time Interval Coverage Constraints in Nurse Scheduling." In: *Applied Artificial Intelligence* 20 (Dec. 2006), pp. 743–766. DOI: 10.1080/08839510600903841.

[23] Jacek Błażewicz et al. *Handbook on scheduling: From theory to applications.* Springer, 2007.

[24] Silvano Martello et al. "An algorithm for the generalized assignment problem". In: *Operations Research* 29.5 (1981), pp. 761–774.

[25] S.A. Kravchenko et al. "Parallel machine scheduling problems with a single server". In: *Mathematical and Computer Modelling* 26.12 (1997), pp. 1–11. ISSN: 0895-7177. DOI: https://doi.org/10.1016/S0895-7177(97)00236-7. URL: https://www.sciencedirect.com/science/article/pii/S0895717797002367.

[26] Héctor G.-de-Alba et al. "A mixed integer formulation and an efficient metaheuristic for the unrelated parallel machine scheduling problem: Total tardiness minimization". In: *EURO Journal on Computational Optimization* 10 (2022), p. 100034. ISSN: 2192-4406. DOI: https://doi.org/10.1016/j.ejco.2022.100034. URL: https://www.sciencedirect.com/science/article/pii/S2192440622000107.

[27] Jeroen Van den Bergh et al. "A literature review on planning and scheduling in the healthcare sector". In: *European Journal of Operational Research* 258.1 (2013), pp. 3–17.

[28] Amy Cohn et al. "Scheduling Medical Residents at Boston University School of Medicine". In: *Interfaces* 39 (June 2009), pp. 186–195. DOI: 10.1287/inte.1080.0369.

[29] B. Naderi et al. "Scheduling open shops with parallel machines to minimize total completion time". In: *Journal of Computational and Applied Mathematics* 235.5 (2011), pp. 1275–1287. ISSN: 0377-0427. DOI: https://doi.org/10.1016/j.cam.2010.08.013. URL: https://www.sciencedirect.com/science/article/pii/S0377042710004516.

[30] Jiaxing Li et al. *A Simulated Annealing Approach to Identical Parallel Machine Scheduling.* 2024. URL: https://arxiv.org/abs/2410.07987.

[31] Peter Hulshof et al. "Tactical resource allocation and elective patient admission planning in care processes". In: *Health Care Management Science* 16 (Jan. 2013). DOI: 10.1007/s10729-012-9219-6.

[32] E.W. Hans et al. "Robust surgery loading". In: *European Journal of Operational Research* 185.3 (2008), pp. 1038–1050. ISSN: 0377-2217. DOI: https://doi.org/10.1016/j.ejor.2006.08.022. URL: https://www.sciencedirect.com/science/article/pii/S0377221706005844.

[33] Scott Kirkpatrick et al. "Optimization by simulated annealing". In: *Science* 220.4598 (1983), pp. 671–680.

[34] Compendium voor de Leefomgeving. *CO2-emissie per voertuigkilometer van nieuwe personenauto's, 1998-2017.* Accessed: 2024-11-20. 2017. URL: https://www.clo.nl/indicatoren/nl013415-co2-emissie-per-voertuigkilometer-van-nieuwe-personenautos-1998-2017.

[35] F. Costantino et al. "Exploring the Bullwhip Effect and Inventory Stability in a Seasonal Supply Chain". In: *International Journal of Engineering Business Management* 5 (2013). DOI: 10.5772/56833.

[36] ANWB. *Routeplanner*. `https://www.anwb.nl/verkeer/routeplanner`. Accessed: 2024-11-08. 2024.

A Data analysis and structuring

We start this chapter by explaining the datasets available for our analysis in Section A.1. In this same section we also show the information required for our analysis, and how we structured this information. In Section A.3 we explain the steps we undertook to make sure that our final dataset would only contain relevant appointments. Any steps undertaken up to the end of Section A.3 hold for both the *location allocation*, and the *room planning*. For the first, the dataset will be used to determine the required capacity in each location. For the second, we use the dataset to determine the work that was performed in each week, and create new schedules based on this work. We explain additional considerations that only hold for the location allocation in Section A.4. Lastly, we describe any explicit or implicit assumptions we made in the data processing in Section A.5.

A.1 Available datasets

For our data analysis we rely on the appointment dataset exported from Isala's HiX system. This dataset includes all appointments carried out by the department from 2021-2023. Not all required information was explicitly present in this initial dataset, and was therefore included through merging information from other datasets. Each of the used datasets can be seen in Figure A-I.

Dataset	Information contained within
Main HiX dataset	Contains all information on the appointment. Relevant information are the appointment code, planned location, all information on the date/time of the appointment, the attending employee, the diagnosis label and the duration.
Employee Data	Contains information on all employees working within the full internal medicine department. It includes their role within the hospital, their specialty, and whether or not they hold consultations.
Appointment Code Data	Contains a mapping for each of the original appointment codes to the simpler unifying codes of CP, NP, TC and NPTC.
Patient Referral Data	Includes all patient data, the only relevant data being the patients' ZIP codes.
DBC Data	Contains information on the 90% most common diagnoses for GIM, showing which specialties can treat the diagnosis.

Table A-I: Explanation of available datasets

The HiX dataset and patient referral dataset were generated directly from the database of Isala. The employee dataset was obtained by first extracting all unique employee names contained within the HiX dataset, and subsequently inquiring with all the different specialties about the role of each of the extracted employees. We informed, at the same time, about whether each considered employee ever holds consultations. The DBC dataset was created by the doctors of the GIM sub-department, and derived from results of a study by van der Wel, (2023).

We created the appointment code dataset by manually looking at all the appointment codes. Most contain control patient (CP), new patient (NP), or telephonic consultation (TC) directly within the code, and could therefore immediately be classified as such. To achieve the classification of the few remaining codes, we used another column in the dataset. This column contained a lot of missing values, but for some of the appointments contained a classification showing either 'telephonic', 'return', or 'first time'. By summarizing the frequency of these classifications in Excel, we could identify which of CP, NP, TC and NPTC was appropriate.

Required information

The information required to answer the questions within this research is explained in Table A-II. In this table we also show a variable name used to show the structure of our final dataset. (Figure A-1)

Explanation	Variable Name
The duration of the appointment.	AppointmentDuration
The location where the appointment took place.	PlannedLocation
Information on the time of the appointment (respectively year, week, day, morning/afternoon).	Year, WeekNr, Day, Morning/Afternoon
The name of the employee that attended the appointment.	AttendingEmployeeName
The profession of the employee that attended the appointment.	AttendingEmployeeProfession
The specialty of the employee that attended the appointment.	AttendingEmployeeSpecialty
The specialty under which the appointment is categorized.	AppointmentSpecialty
The location closest to the patient's residence for the appointment.	ClosestLocation
An identifier to determine whether the appointment involved a new patient or a return patient, and whether it was a physical or telephonic appointment.	AppointmentType

Table A-II: Appointment data required information and variable names

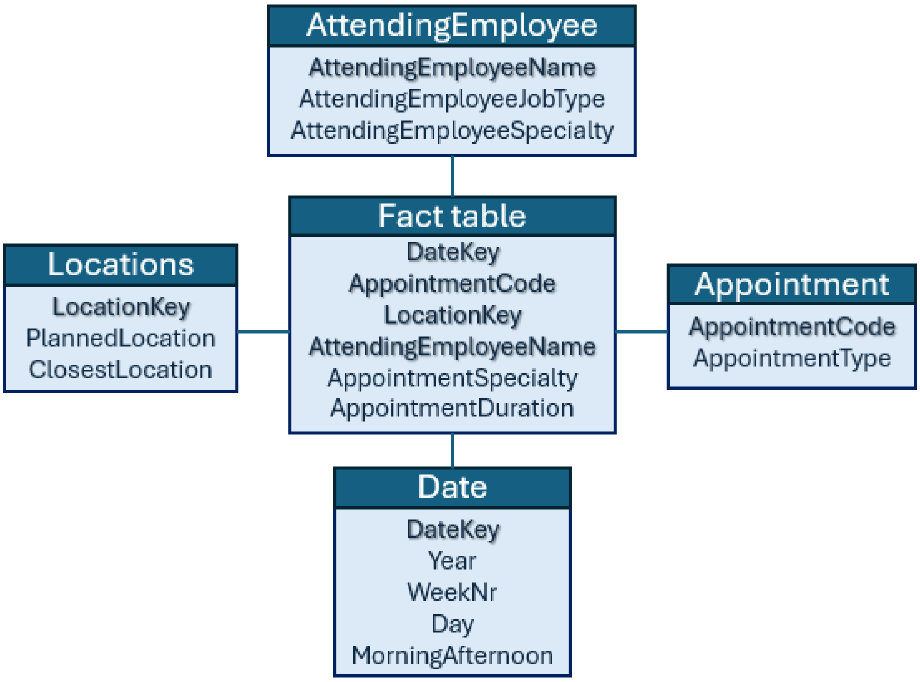

Figure A-1: Datastructure pivot table

A.2 Including previous relevant Isala studies

In Section 2.1.3 we discussed how the distribution of workload across specialties should be based on the specialties. Before we can do this, however, we need to couple the datasets, which is complicated by an issue within the data. Specifically, some DBC numbers are attached to multiple unique diagnoses. We select the most frequent of the diagnoses, as we know we are looking for those. Since the 90% DBC's are covered by only 55 codes, we can validate by eye whether the diagnosis of the DBC in the *HiX dataset* matches the diagnosis attached to the same DBC number in the *DBC dataset* from van der Wel, (2023).

Having done this correctly, we can redistribute the DBC's. To accomplish this, we copy the relevant appointments for every specialty involved, and scale it to each specialties' involvement. For instance, if two specialties are expected to handle 60% and 40% of appointments for a particular DBC, the original appointment is copied twice, and each copy is assigned a proportionate share of the total duration - 60% of the duration goes to specialty 1, and 40% to specialty 2. The original is deleted.

A.3 Determining relevant appointments

Appointments that did not involve a consultation room

The care we are optimizing for within our scope is care that unequivocally requires a consultation room. In accordance with this criterion, we can immediately exclude some appointments on the basis of specific knowledge from the department. First, although our dataset is an 'appointment' dataset, it also contains entries that are known to not directly require a consultation room. This is due to the fact that employees track all of their patient-related activities in the HiX system. This means that the data also includes slots that involved administrative work, measurements, analysis of measurements, emails, et cetera. These types of appointments should not be present in the appointment data, since we are only interested in appointments in which the patient was directly involved. We consulted hospital management and the different specialties to generate a list of codes and names that should be excluded from the data. These can be found in Appendix A.6. All appointments containing any of these codes or names were excluded (16.03%). We also excluded any appointments marked as 'Niet Voldaan' (not executed - 0.62%), 'Onbekend' (unknown - 0.004%) or 'No-show' (1.04%).

Weekends and nights

Another type of appointment that should be excluded when looking at the daily workload is any appointment logged during the weekends (0.13%) or at night during the week (0.09%). These unscheduled 'appointments' are unavoidable, since there will always be a number of occurrences that a healthcare professional is required outside of working hours. As these appointments, therefore, can not be moved to the day, it can be seen as a separate workload whenever we analyze the daily demand. We should specifically mention that work in the evenings is included in our analysis, since this work *is* usually scheduled.

Minimum workload

The consultation schedule operates on a cyclic weekly basis, meaning that any employee assigned a recurring slot will be allocated a room-slot every week. To ensure efficient use of this capacity, these slots should only be given to professions with a sufficiently high workload. We mention professions specifically because, within each specialty, the workload is typically shared; any patient seen by one doctor of a certain specialty can usually be treated by another doctor of that same specialty. A precondition for efficient capacity use for a profession, therefore, is that there is enough work for that profession on a weekly basis. We consulted with the hospital and decided on a minimum weekly workload of one hour and thirty minutes.

We can determine the number of hours worked per profession in previous years to see if there are any roles that should be excluded on this criterion (Table A-III on page 87). This holds for five, specifically: Nurse-research, Nurse-Noorderboog, Psychologist, Psycho-diagnostic workers, and Medtronic. Through consulting the hospital, we also know that Nurse - Noorderboog is a different party, and should not be considered for Isala's capacity on that basis either. Intermezzo and Medtronic additionally do not require a consultation room for the work they do, and should also be excluded on that basis. The percentage of the total workload is shown for the excluded professions in Table A-III as well. For the percentages of the remaining professions after this exclusion we refer the reader back to Figure D-2 on page 103.

Profession	Weekly hrs. 2021	Weekly hrs. 2022	Weekly hrs. 2023	Included?
Doctor	411.7	422.4	412.8	Yes
Nurse	193.9	206.6	219.4	Yes
Nurse specialist	109.3	111.7	115.7	Yes
Doc. assistant	38.0	31.6	35.5	Yes
Intermezzo	11.8	12.1	13.2 (1.61%)	No
Co-assistant	8.7	9.0	10.7	Yes
Phys. assistant	0.7	0.1	5.5	Yes
Nurse - research	0.3	0.7	1.0 (0.13%)	No
Nurse - Noorderboog		0.2	1.3 (0.16%)	No
Psychologist		0.2	1.1 (0.14%)	No
Psycho-diagnostic worker		0.0	1.1 (0.13%)	No
Medtronic			0.3 (0.03%)	No

Table A-III: Weekly hours worked per profession in each year and whether or not the role is included

A.4 Additional considerations for location allocation

Material limitations

The last fact that we should consider, relevant only for the location allocation, is that certain appointments can not be moved to different locations, as we alluded to in Section 2.1.1. This holds specifically for all appointments from the Acute care specialty (0.47%), and appointments classified as palliative care (0.003%). It also involves Fibrosis-4 measurements, several treatments and diagnoses under Hematology, and any appointment in H&O that is preceded by a blood test. (percentages shown in Table A-IV) The reason for the exclusion of the first two appointment types is that the relevant equipment is only in available in Zwolle. The reason for the exclusion of the last is that blood tests can only be analyzed in Zwolle. This means that follow-up appointments based on the results of the blood test can not be scheduled on the same day, and will therefore have to be scheduled on two separate days, increasing total travel time.

For the blood tests and chemotherapy the exact percentage can currently not be determined, since those two appointments have not been logged in the 2021-2023 datasets. We have made an estimate in consultation with the H&O sub-department, and will exclude a percentage of the work from the calculations at the satellite locations. The other two exclusions have been incorporated within the dataset.

Specialty	**Type of appointment**	**Percentage** (*specialty*)	**Percentage** (*total*)
Infectious disease	Fibrosis-4 measurements	3.73%	0.24%
Hematology	Various treatments and diagnostics	Doctors: 5.24% Nurses: 45.65% Nurse spec.: 7.14%	2.37%
Hematology and oncology	Appointments on the same day as a blood test or chemotherapy	15% (estimate)	2.23% (estimate)

Table A-IV: Appointments that can not be executed in the satellite locations

A.5 Assumptions

There was a small number of missing values in the initial datasets. We tracked these values throughout the data processing and determined that most of them were excluded by one or other of our exclusion criteria. The missing values remaining after the process were then identified, and either excluded or completed through inside knowledge.

The most heavily affected variable was AppointmentSpecialty, with 0.08% missing values for doctors of the GIM sub-department. As most doctors' appointments are handled by an employee from the correct specialty, we assume that the AppointmentSpecialty would have corresponded to the

Assumptions methods	Explanation
Any work currently done by one employee can not be done by any other employee	This assumption simplifies the analysis. Without this, we would also have to determine all the appointments within the dataset that each profession would be eligible for.
The resources and demand in the near future can be assumed equivalent and unchanging with respect to 2023	This assumption is based on the fact that the workload of the department is not allowed to increase year over year by national policy. The distribution of work may change slightly, however.
Patients want to be treated in the closest location all else being equal.	This assumption is very likely, but even if preferences for different locations exist, national policy according to IZA is to treat patients in the closest location.
Psychologists, Nurse-research, and psycho-diagnostic employees do not require explicit consideration in the room schedule	The number of hours worked by these employees is sufficiently low that we de-prioritize providing them with a windowed room. They can be allocated an available windowed slot after creation of the schedule, or choose a windowless slot at any time. This would mean they work less than one hour and thirty minutes in windowless rooms.
The general pool work will be divided accurately among employees in the satellite locations	The department has to communicate this effectively to the planning department in order to achieve the right percentage of general patients that follows from our results. We reiterate this in the conclusions.
Each room can be used for two slots per day	To maintain clarity and consistency of Zwolle with other locations, we adhere to the hospital policy of two consultation slots per day.
Patients to the east of the river Ijssel are allocated to Zwolle instead of Heerde	We assume that the inconvenience and cost of the ferry is enough to motivate these patients to go to Zwolle instead of Heerde.

Table A-V: Assumptions and explanation

attending doctors' specialty. Since 0.08% only corresponds to an average 0.2 hours per week that might be placed with the wrong specialty, and our approach has a high likelihood of finding the correct specialty, we assume this will not have a significant effect on our results.

Next is appointments without a ZIP code (0.01%), and appointments without information on the location that they were planned in (0.008%). For both of these, we assume that they would be treated in Zwolle. Since this, again, affects only a small number of patients, and Zwolle is the most likely hospital to be treated in, we again assume insignificant effects.

Last are a few missing appointment codes for a specific doctor in 2021 (0.003%). Since our calculations are focused on 2023, and this doctor no longer works here, we just exclude these appointments. This has no effect on our analysis of 2023.

A.6 Excluded codes and employee names

Excluded Employees	**Excluded codes**
24 UURS METING KAMPEN	(blank)
BEENMERGPUNCTIE	*ADM
BLOEDDRUK/ECG STEENWIJK	*CPAFD
BLOEDDRUKAGENDA KAMPEN 1	*CPSCHRF
BLOEDDRUKAGENDA KAMPEN 2	*CRISTKL
BLOEDDRUKMETER	*ECONSUL
BLOEDDRUKMETING HEERDE	*KLBEZ
BLOEDDRUKMETING MEPPEL	*LP
DAGCOORDINATOR RVO MEPPEL	*LUMBP
ECG/BLOEDDRUK MEPPEL	*NPSCHRF
Klinisch	*RESERVC
LUMBAALPUNCTIE	*TCVKUIT
ONCOLOGISCH SECRETARIAAT	AB
ONDERZOEKEN	ABNEURO
ONDERZOEKEN MEPPEL	ABSTART
SECRETARIAAT ZWOLLE	ABSTOP
UITLEZEN	ADM
	ADMCOACH
	ADMINSAN
	BA
	BUIKVETB
	ECG
	EMAIL
	EMAILVK
	GLSENAAN
	GLSENAF
	KL
	KLINISCH
	MENW
	ORTHOMET
	OVLBESP
	PDHUISBZ
	TCUITLEZ
	TRIPEECG
	UITLEZEN
	VACCIN
	VPHUIS

A.7 Included codes & classification

Code	*Type*	*Code*	*Type*	*Code*	*Type*	*Code*	*Type*
***NAGESPR**	CP	**CPNDVTVK**	CP	**NPFIBVS**	NP	**PDPOLI**	CP
***NPONCO**	NP	**CPNDVTVS**	CP	**NPGEH**	NP	**PDPOSTOK**	CP
***TCNAGES**	TC	**CPOG**	CP	**NPGEHCMD**	NP	**PDTRAIN**	CP
BA	NP	**CPOGGH**	CP	**NPGRAVPA**	NP	**POLIPO**	CP
BEGELEID	CP	**CPOGVK**	CP	**NPGRAVVK**	NP	**PORINSPA**	CP
CAPD	CP	**CPOGVS**	CP	**NPGRAVVS**	NP	**PORINSVK**	CP
COCHEVPK	CP	**CPOI**	CP	**NPGRVK**	NP	**POVER2**	CP
COPOMP	CP	**CPOIPA**	CP	**NPGRVTCV**	NPTC	**POVERLP1**	CP
COPOMPVS	CP	**CPOIVS**	CP	**NPH**	NP	**POVERLPA**	CP
COPOPNAM	CP	**CPPA**	CP	**NPHEM**	NP	**POVERLVK**	CP
COSPOED	CP	**CPRR**	CP	**NPHVK**	NP	**POVULPA**	CP
COVUL	CP	**CPSANA**	CP	**NPHVS**	NP	**POVULVK**	CP
CP	CP	**CPSPINA**	CP	**NPKDO**	NP	**SENAANVK**	CP
CP15	CP	**CPSPOED**	CP	**NPKDOVK**	NP	**SENSAAN**	CP
CP30	CP	**CPTECHSU**	CP	**NPMDO**	NP	**SENSAFVK**	CP
CPALLO	CP	**CPTHD**	CP	**NPMOB**	NP	**TC**	TC
CPBBVK	TC	**CPTHY**	CP	**NPMOBVK**	NP	**TCBC**	TC
CPBEELDB	TC	**CPTHYVK**	CP	**NPMOBVS**	NP	**TCBCCHEM**	TC
CPCHEMO	CP	**CPTRANS**	CP	**NPNOD**	NP	**TCBCKUUR**	TC
CPCHEVPK	CP	**CPTVS**	CP	**NPNODVK**	NP	**TCBCPACT**	TC
CPCHEVS	CP	**CPV2.3**	CP	**NPOG**	NP	**TCBCSANA**	TC
CPCVS	CP	**CPV3.3**	CP	**NPOGGH**	NP	**TCDVTVK**	TC
CPDIALYS	CP	**CPVAATVK**	CP	**NPOGGHVK**	NP	**TCDVTVS**	TC
CPDVTVK	CP	**CPVAATVS**	CP	**NPOGVK**	NP	**TCGRAVVS**	TC
CPDVTVS	CP	**CPVK**	CP	**NPOGVS**	NP	**TCKUURVS**	TC
CPGEHEUG	CP	**CPVPK**	CP	**NPOIPA**	NP	**TCLANG**	TC
CPGRAVK	CP	**CPVS**	CP	**NPOIVS**	NP	**TCLANGVK**	TC
CPGRAVS	CP	**CPVSDIET**	CP	**NPON**	NP	**TCLANGVP**	TC
CPGRAVVK	TC	**CPVSIO**	CP	**NPONC**	NP	**TCLANGVS**	TC
CPGROEI	CP	**CRISTAB**	CP	**NPPA**	NP	**TCNZ**	TC
CPGRVK	CP	**FIB**	CP	**NPPKOVS**	NP	**TCOIPA**	TC
CPH	CP	**GESP**	CP	**NPPOL**	NP	**TCOSTPA**	TC
CPHEP	CP	**GESPREK**	CP	**NPTC**	NPTC	**TCOSTVS**	TC
CPHEPVK	CP	**GESPRIZ**	CP	**NPTCCOV**	NP	**TCSANA**	TC
CPHEPVS	CP	**GESPVPK**	CP	**NPTCDVTK**	NPTC	**TCTERUG**	TC
CPHVK	CP	**GESPVS**	CP	**NPTCDVTS**	NPTC	**TCVK**	TC
CPHVS	CP	**INFO30VP**	CP	**NPTCVK**	NPTC	**TCVKNZ**	TC
CPIDE	CP	**INFO30VS**	CP	**NPTHY**	NP	**TCVPK**	TC
CPIDEFUN	CP	**INFO60VP**	CP	**NPTHYVK**	NP	**TCVPK15**	TC
CPIDERR	CP	**INFO60VS**	CP	**NPVAATVK**	NP	**TCVPK30**	TC
CPIDEVK	CP	**INFO90VP**	CP	**NPVAATVS**	NP	**TCVPK60**	TC
CPKUUR	CP	**NAGESPR**	CP	**NPVIOS**	NP	**VOORLVP**	CP
CPKUURVP	CP	**NP**	NP	**NPVK**	NP	**CPNACHT**	CP
CPKUURVS	CP	**NPBBVK**	NP	**NPVPK**	NP	**NPDVTVS**	NP
CPMDO	CP	**NPBEELDB**	NPTC	**NPVS**	NP	**PDPETTST**	CP
CPMOB	CP	**NPCO**	NP	**NWP+CO**	NP	**CPMVS**	CP
CPMOBVK	CP	**NPCONTRA**	NP	**NWPOI**	CP	**NPDVTVK**	NP
CPMOBVS	CP	**NPDVNTR**	NP	**PDBEELDB**	TC	**PDOVERIG**	CP

A.8 Job roles considered for reallocation

Job Role	Work considered?	Employee considered?
Doctor	Yes	Yes
Nurse	Yes	Yes
Nurse specialist	Yes	Yes
Co-assistant	Yes	No
Doctors' assistant	Yes	No
Physician Assistant	Yes	No
Diagnostics	No	No
Intermezzo	No	No
Medtronic	No	No
Nurse - Noorderboog	No	No
Nurse - Research	No	No
Palliative care	No	No
Psycho-diagnostic worker	No	No
Psychologist	No	No
Secretariat	No	No

Table A-VI: Job roles and whether they are considered for reallocation

A.9 Example location

Below we show an example of a hypothetical patient living at Mullegenweg 1, Oldebroek. This street is one of the places in the adherence area of Isala that most clearly has similar distances to multiple locations. Using a combination of straight-line distance between coordinates and the ANWB routeplanner[36], we show the distance from Mullegenweg 1 to Zwolle, Kampen, and Heerde in Table A-VII.

Possible location	Distance (straight)	Distance (travel)	Travel time
Zwolle	15.7 km	21.1 km	21 minutes
Kampen	12.7 km	16.9 km	17 minutes
Heerde	9.0 km	15.6 km	18 minutes

Table A-VII: Distance for Mullegenweg 1 to Isala locations in Zwolle, Kampen and Heerde

B Description of code

Class structure

Class CWeek

- **Attributes:**
  - `WeekNr`: `Integer` Represents the week number in the calendar.
  - `OV, CurrentOV, BestOV`: `Integer`
    Objective value of the solution for this week. `Current` and `Best` are used during optimization to track intermediate and best results.
  - `TotalNumExcessRooms, CurrentTotalNumExcessRooms, BestTotalNumExcessRooms`: `Integer`
    Total number of times an employee uses more than one room in this week. Used in calculating the objective value. `Current` and `Best` are used during optimization.
  - `TotalNumExcessEmployees, CurrentTotalNumExcessEmployees, BestTotalNumExcessEmployees`: `Integer`
    Total number of times a room is occupied by more than one employee in this week. Used in calculating the objective value. `Current` and `Best` are used during optimization.
  - `TotalNumRoomSwitches, CurrentTotalNumRoomSwitches, BestTotalNumRoomSwitches`: `Integer`
    Total number of times an employee switches rooms during this week. Used in calculating the objective value. `Current` and `Best` are used during optimization.
  - `TotalNumInadequateSlots, CurrentTotalNumInadequateSlots, BestTotalNumInadequateSlots`: `Integer`
    Total number of slots where assignments are inadequate (e.g., mismatch in workload or room requirements). `Current` and `Best` are used during optimization.
  - `Rooms`: `Array of CRoom`
    Array of `CRoom` objects associated with this week.
  - `Employees`: `Array of CEmployee`
    Array of `CEmployee` objects associated with this week.

Class CRoom

- **Attributes:**
  - `RoomNr`: `Integer` Identifier for the room.
  - `Quality`: `String`
    Describes the quality level of the room (e.g., equipment, size).

  - `PlannedEmployeeIDs, CurrentPlannedEmployeeIDs, BestPlannedEmployeeIDs`: `Array[0..NumDa`
    `of Array[0..NumSlotsPerDay-1] of Integer`
    Matrix tracking which employees are scheduled in the room for each timeslot. `Current` and `Best` are used during optimization.
  - `NumExcessEmployees, CurrentNumExcessEmployees, BestNumExcessEmployees`: `Integer`
    Tracks the number of times the room is over-occupied (more than one employee assigned per timeslot). `Current` and `Best` are used during optimization.
  - `InadequateSlots, CurrentInadequateSlots, BestInadequateSlots`: `Integer`
    Tracks the number of slots where room assignments fail to meet adequacy criteria. `Current` and `Best` are used during optimization.

- **Methods:**
  - `procedure InitializePlannedEmployeeIDs`
    Sets all slots to an initial unassigned state (e.g., -1 for free slots).
  - `procedure InitRoom(const r: Integer)`
    Calls `InitializePlannedEmployeeIDs` and sets all other attributes of the room to their initial values.

Class CEmployee

- **Attributes:**
  - `Name`: `String` Name of the employee.
  - `ID`: `Integer`
    Unique identifier for the employee.
  - `RequiredWorkload`: `Double`
    Full-time equivalent workload required for the employee during the week.
  - `SpecialtyID`: `Integer`
    Identifier for the employee's specialty.
  - `WorkSchedule, CurrentWorkSchedule, BestWorkSchedule`: `Array[0..NumDaysPerWeek-1]`
    `of Array[0..NumSlotsPerDay-1] of Integer`
    Matrix representing the room assignments for the employee in each timeslot. `Current` and `Best` are used during optimization.
  - `NumExcessRooms, CurrentNumExcessRooms, BestNumExcessRooms`: `Integer`
    Tracks the number of times the employee is assigned to more than one room simultaneously. `Current` and `Best` are used during optimization.

  - `NumSwitches, CurrentNumSwitches, BestNumSwitches`: `Integer`
    Tracks the number of times the employee switches rooms within the week. `Current` and `Best` are used during optimization.
  - `PlannedWorkloadWeek, CurrentPlannedWorkloadWeek, BestPlannedWorkloadWeek`: `Double`
    Tracks the total workload assigned to the employee for the week. `Current` and `Best` are used during optimization.

- **Methods:**
  - `procedure InitializeWorkSchedule`
    Sets all elements in the `WorkSchedule` to an unassigned state (e.g., -1).
  - `procedure InitializeScheduleConstraints(const e:  Integer)`
    Marks unavailable timeslots in the `WorkSchedule` (e.g., -2 for unavailable slots).
  - `procedure InitEmployee(const empNr:  Integer)`
    Calls `InitializeWorkSchedule` and `InitializeScheduleConstraints`, setting other attributes to their initial values.

Algorithm 8 Sequential Construction

```
procedure ScheduleAllEmployeesSequentiallyRoomsFirst
variables: Week, Employee, r, d, s ∈ ℤ {Indices for week, employee, room, day, slot}
           SlotFound ← False {Tracks slot allocation}
for Week ∈ (Weeks) do
  for Employee ∈ (Week.Employees) do
    while Employee.PlannedWorkload < Employee.RequiredWorkload do
      SlotFound ← False
      for r ← 0 to NumRooms − 1 do
        for d ← 0 to NumDaysPerWeek − 1 do
          for s ← 0 to NumSlotsPerDay − 1 do
            if RoomIsEmpty and VerifyAddition(Week.WeekNr, Employee.ID, r, d, s)
            then
              AddToWorkAndRoomSchedule(Week.WeekNr, Employee.ID, r, d, s)
              SlotFound ← True
              break
            end if
          end for
          if SlotFound then
            break
          end if
        end for
        if SlotFound then
          break
        end if
      end for
      if not SlotFound then
        LogMessage('No suitable schedule found for ' + Employee.Name)
        break
      end if
    end while
  end for
  VerifySchedule
  CalcObjectiveFunction
  StoreBestSolution
  UpdateForm
end for
```

Algorithm 9 Simulated Annealing Algorithm

```
procedure SimulatedAnnealing(w)
var Iteration: Integer // Iteration number
var T_start, T_final, CurrentTemp, Alpha: Double // SA parameters
var MarkovChainLength, MaxMarkovChainLength: Integer // SA parameters
// Stopping variable
StopLocalSearch ← False
// Parameters for simulated annealing
T_start ← StrToFloat(Form1.Edit1.Text) // Initial temperature
CurrentTemp ← T_start // Temperature
T_final ← StrToFloat(Form1.Edit2.Text) // Final temperature
Alpha ← StrToFloat(Form1.Edit3.Text) // Cooling rate
MaxMarkovChainLength ← StrToInt(Form1.Edit4.Text) // Length of Markov Chain at low temperatures (length increases with lower T)
Iteration ← 0
CalcObjectiveFunction(w)
StoreBestSolution(w)
StoreCurrentSolution(w)
// Clear chart data
ClearCharts()
while (StopLocalSearch = False) and (CurrentTemp > T_final) do
   // Process messages
   Application.ProcessMessages
   // Init data for acceptance ratio charts (updated for every temperature)
   NumAcceptedCurrentBetter ← 0 // Number of times the current solution was updated by a better solution
   NumAcceptedCurrentEqual ← 0 // Number of times the current solution was updated by an equivalent solution
   NumAcceptedCurrentWorse ← 0 // Number of times the current solution was updated by a worse solution
   NumAcceptedBest ← 0 // Number of times the best solution was updated
   NumAttempted ← 0 // Number of times a neighbour was found
   // Determine Markov Chain length for this temperature
   MarkovChainLength ← DetermineMarkovLength(MaxMarkovChainLength, CurrentTemp)
   // Perform local search with number of iterations == MarkovChainLength
   LocalSearchSA(CurrentTemp, w, MarkovChainLength)
   // Update temperature
   CurrentTemp ← CurrentTemp * Alpha
   Form1.Edit8.Text ← Format('
   // Reset Data back to best week after each Markov Chain
   RestoreBestSolution(w)
   // Update Charts
   UpdateForm(w, Iteration)
   // Increment the number of iterations of Markov Chains
   Inc(Iteration)
end while
MainFinalization()
```

C Cooling Scheme Tuning

Having created the SA program, we have to decide on our cooling scheme parameters. These parameters are the starting temperature, the stopping temperature, the Markov chain length, and the cooling rate. The starting temperature determines the acceptance ratio that we start with, and the ending temperature determines the acceptance ratio that we end with. The Markov chain length determines the number of iterations per temperature, and the cooling ratio (α) determines how quickly we decrease the temperature. The new temperature at each iteration is calculated by multiplying the previous temperature by the cooling ratio.

To determine the cooling parameters that obtain the best solution, we test a number of different configurations to create a complete schedule based on the workload in a typical non-reduction week (Week 7 of 2023). We use a non-reduction week because the workload in these weeks is higher than in reduction weeks, which provides a more robust test of the model.

We experiment with parameters where the initial acceptance ratio is approximately 100%, and the acceptance ratio at the end is approximately 0%.[33] When the acceptance ratio is 0%, the Simulated Annealing algorithm is essentially reduced to a gradient descent [33], and a longer time at 0% allows the algorithm more time to find its local optimum.

Our acceptance ratio is based on the Boltzmann distribution [33], and can be found in Equation 4. Here NewOV is the objective value (OV) of the solution after a swap has taken place, and CurrentOV is the objective value of the solution before the swap.

$$\begin{aligned} NewOV - CurrentOV > 0: \quad & P_{accept} = e^{-\frac{(NewOV - CurrentOV)}{T}} \\ NewOV - CurrentOV \leq 0: \quad & P_{accept} = 1 \end{aligned} \tag{4}$$

We can make an estimate on the typical acceptance ratio at a given temperature based on the typical difference between NewOV and CurrentOV. Through analysis of the change in the objective value, we see that it almost never increases by a value of more than four from a swap. We want the acceptance ratio for this largest typical swap to be at least 90% at the starting temperature, and the acceptance ratio for a swap that increases the OV by one to be at most at 0.005% at the ending temperature. We can calculate the starting and ending temperatures required to achieve this using 4. For the starting temperature we test several other values corresponding with 95%, 99%, and 99.5% at the start. For the ending temperatures we test several smaller values to determine whether a longer gradient descent at the end is beneficial. We attempt three different Markov Chain Lengths, and calculate α from our chosen values according to Equation 5.

$$\alpha = \left(\frac{T_{end}}{T_{start}} \right)^{\frac{MarkovLength}{NumIterations}} \tag{5}$$

Our test structure uses 27 rooms with ten slots, this means that there are a total of $27^2 = 72900$

possible swaps that we can do, disregarding any constraints and redundancy. Taking this into account, we start with 10^5 swaps using 48 different configurations. All used values can be found in Table C-VIII, where each combination of these options is tested as a possible configuration. The scheme occasionally finds solutions that are relatively worse or better, and we therefore need a procedure to avoid these outliers. To accomplish this, we run each combination three times, and use the second best result. The results can be found in Table C-X of Appendix C. When we look

Possible Markov Length	**Possible** T_{start}	**Possible** T_{end}
10^3	38	10^{-1}
$5 \cdot 10^3$	75	$5 \cdot 10^{-3}$
10^4	400	10^{-4}
	800	10^{-5}

Table C-VIII: Parameters options used for the 48 configurations for T_{start}, T_{end}, and Markov length

at the configuration parameters individually, we see that configurations with a larger Markov chain length have a lower objective value. We also see that lower starting- and ending temperatures result in lower objective values. We hypothesize that this is all due to the model benefiting from more exploitation time rather than exploration time, and the equivalent number of iterations means that a lower starting and ending temperature allow for more exploitation at lower temperatures. The longer Markov chain in turn means that the heuristic can attempt more swaps in order to find the local optimum at each temperature level, before moving on to the next level.

Based on this, we narrow our parameters down to a Markov length of 10^4, T_{start} of 400, 75, and 38, and T_{end} of $5 \cdot 10^{-3}$, 10^{-4}, 10^{-5}, and use 10^6 swaps. This time, we do five runs and take the average of the 2-4 best runs. With this, we can see in Table C that the differences between all configurations are fairly insignificant. We therefore choose T_{start} to be 400, so as to enable more exploration in other instances that might need this. T_{end} we choose to be 10^{-5}, since both this and the previous analysis showed the best average results using this temperature. Based on 10^6 iterations and a Markov length of 10^5, we calculate using Equation 5 that $\alpha \approx 0.86$. This means the cooling rate will be relatively slow, and explore many temperatures.

T_{start} \ T_{end}	10^{-5}	10^{-4}	$5 \cdot 10^{-3}$	**Average**
38	53.33	54.33	53.33	53.66
75	53.00	52.66	54.33	53.33
400	53.00	53.33	54.66	53.66
Average	53.11	53.44	54.11	53.55

Table C-IX: Results for various cooling scheme configurations using 10^6 swaps

Results for different cooling scheme configurations using 10^5 swaps						
Markov Length	T_{end} \ T_{start}	38	75	400	800	**Average**
10^3	10^{-5}	82	91	84	101	90
	10^{-4}	86	94	95	97	93
	$5 \cdot 10^{-3}$	104	95	105	108	103
	10^{-1}	127	125	147	138	134
$5 \cdot 10^3$	10^{-5}	86	88	80	96	88
	10^{-4}	94	86	87	94	90
	$5 \cdot 10^{-3}$	96	100	108	115	105
	10^{-1}	111	126	132	129	125
10^4	10^{-5}	87	78	97	83	86
	10^{-4}	91	86	83	89	87
	$5 \cdot 10^{-3}$	95	96	102	108	100
	10^{-1}	100	121	120	120	115
Average		97	99	103	107	101

Table C-X: Performance metrics for varying cooling parameters using 10^5 iterations

D Supporting tables for location capacity allocation

6.2 Supporting tables for Section 6.2.

Job Role	Aandachtsgebied	Zwolle	Meppel	Kampen	Heerde	Steenwijk
Co-assistant	0. Algemeen Interne	3.3	3.1	1.1	0.8	1.5
	4. Nefrologie	0.8	0.1	0.1	0.2	0.0
Doc. assistant	0. Algemeen Interne	1.5	0.5	3.3	0.7	0.3
	1. Acute Geneeskunde	2.2				
	4. Nefrologie	4.1	0.7	1.0	1.1	0.3
	7. Hematologie	2.7	1.7	1.3	1.0	0.8
	8. Oncologie	1.6	1.4	0.5	0.6	0.9
Phys. assistant	0. Algemeen Interne	0.3	0.0	0.2	0.3	0.0
	7. Hematologie	0.1	0.0	0.0	0.0	0.0

Table D-XI: Job Roles Excluded from Satellite Location Planning

Location	Doctors' assistants	Available demand (hrs/week)
Meppel	*Anonymized 1*, *Anonymized 2*, *Other (total)*	1.25, 1.59, *1.52*
Kampen	*Anonymized 3*, *Anonymized 2*, *Other (total)*	2.36, 1.20, *2.49*
Heerde	*Other (total)*	*3.42*
Steenwijk	*Other (total)*	*2.22*

Table D-XII: Demand in each location for specific doctors' assistants

Supporting dashboard current capacity allocation

Now that we have explained all the different resource dimensions, we can look at the distribution of workload across the resources. We express this workload in terms of hours, summed over the entirety of 2023. The results are shown in Figure D-2. We show how the *total* workload is distributed over the locations, and over the employees, and show the percentage of work that is physical/telephonic and new/returning. The graphs on the right then show how these numbers vary across the specialties. The bottom graph shows how the ratio of physical to telephonic work varies across the locations, and the table shows the number of employees working within each specialty. It may seem confusing that there are no doctors working in GIM, whereas we do see a significant workload for them in the second graph on the right. This is, again, because we are looking at the workload for each area of specialization, meaning that this graph includes the general work done by doctors from other specialties.

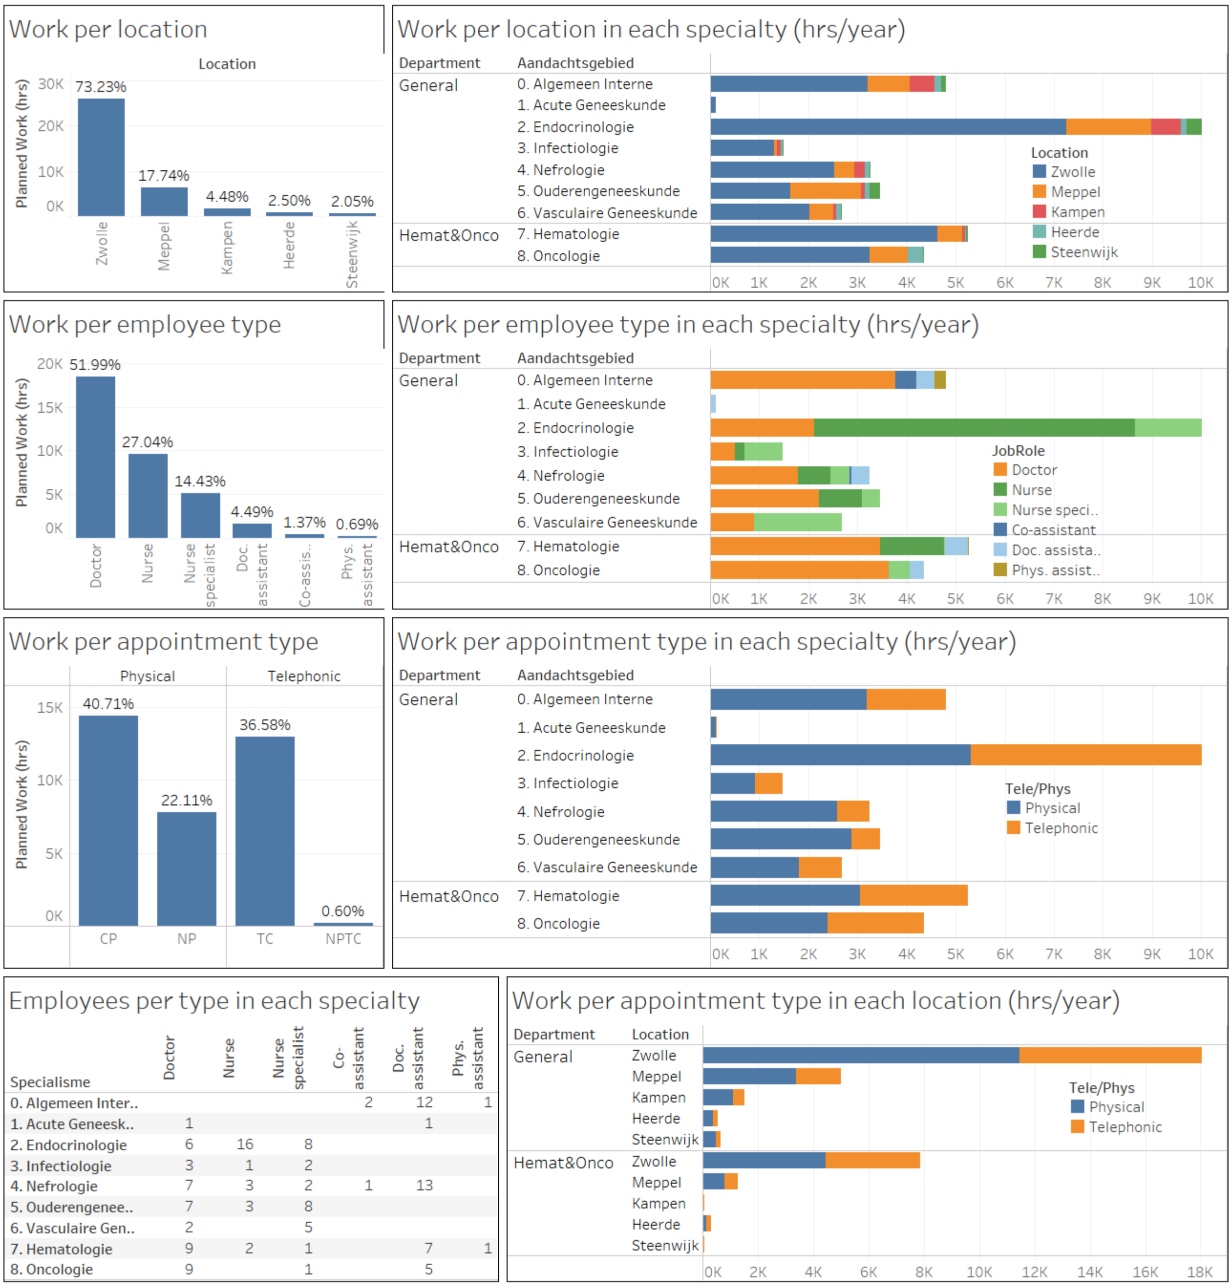

Figure D-2: Dashboard displaying hours worked across multiple resource dimensions

E Contributions of Applied Physics background

Although this thesis does not engage directly with fundamental physics, my Applied Physics (AP) background has significantly shaped its development. Even though I have at this point taken plenty of Industrial Engineering and Management (IEM) courses, my way of thinking is still largely centered in Applied Physics practices.

The typical Applied Physicist questions everything, and wants to know everything for certain before following through with analyses. This has been one of the major challenges for me to overcome throughout IEM in general. Although the approach helps me with curiosity and thoroughness, it also often means I can get overly focused on minor issues, slowing progress. Learning to balance thoroughness with the need to advance the work has been a critical growth experience during this project.

One aspect that was wholly beneficial is the programming skills I acquired in my Applied Physics studies. Writing small, task-specific programs was essential to perform the detailed analyses required for this work. My ability to develop these solutions efficiently and effectively was grounded in the coding proficiency and systematic thinking cultivated during my physics studies.

All in all, I contribute most of the analytical skills I have used throughout my thesis to AP, whereas many of the topic-specific and practical skills have been been derived from my IEM background.