

Kling, Sebastian; Kraul, Sebastian; Brunner, Jens O.

Article — Published Version

Customized GRASP for rehabilitation therapy scheduling with appointment priorities and accounting for therapist satisfaction

OR Spectrum

Provided in Cooperation with:

Springer Nature

Suggested Citation: Kling, Sebastian; Kraul, Sebastian; Brunner, Jens O. (2024) : Customized GRASP for rehabilitation therapy scheduling with appointment priorities and accounting for therapist satisfaction, OR Spectrum, ISSN 1436-6304, Springer, Berlin, Heidelberg, Vol. 46, Iss. 3, pp. 821-872, <https://doi.org/10.1007/s00291-023-00742-y>

This Version is available at:

<https://hdl.handle.net/10419/313816>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<http://creativecommons.org/licenses/by/4.0/>



Customized GRASP for rehabilitation therapy scheduling with appointment priorities and accounting for therapist satisfaction

Sebastian Kling¹ · Sebastian Kraul² · Jens O. Brunner^{1,3,4} 

Received: 29 May 2022 / Accepted: 5 December 2023 / Published online: 12 January 2024
© The Author(s) 2024

Abstract

Physical therapy in acute care hospitals plays an important role in the rehabilitation of patients. Nevertheless, the profession must deal with staff shortages caused by a lack of potential employees and absenteeism which are results of high physical and mental workloads. The therapist shortage negatively affects the total number of daily appointments the department can fulfill. For appointments that can be successfully scheduled, continuity of care with the same therapist cannot be guaranteed for individual patients. Lack of continuity of care negatively influences the therapist's satisfaction. Therapist preferences for individual appointments in general cannot always be guaranteed when designing schedules, which also hurts satisfaction. This paper develops a multi-criteria model for the daily therapy appointment-scheduling problem. The primary objective is to minimize the total sum of priority violations for unscheduled appointments. To improve therapist satisfaction, we consider therapist preferences including continuity of care as a secondary objective. Here, our integer programming formulation aims to minimize the total sum of preference violations for scheduled appointments. We are dealing with an operational planning problem with a daily planning horizon. The operational objective is to achieve therapist schedules in at most two hours. The therapists' schedules together need to include several hundred appointments for a planning day. Due to intractability, the developed integer program cannot provide schedules for such problem sizes. Therefore, we develop a customized Greedy Randomized Adaptive Search Procedure (GRASP) with six innovative local search operations to improve an initially constructed solution. We test the heuristic algorithm on realistic data instances. The metaheuristic provides high-quality schedules for various problem sizes in short runtimes, i.e., within minutes. Comparisons with the optimal solutions for small problem instances show very good results of the GRASP with a similar number of scheduled appointments and good adherence to continuity of care and therapist preference requirements.

Keywords Appointment scheduling · Therapist planning · Rehabilitation Department in Hospitals · GRASP heuristic · Integer programming

1 Introduction

Workforce shortages in several healthcare professions are a well-known problem for personnel planning in hospitals. A possible way to retain employees and become more attractive to new employees is to consider employee preferences for patients. Preferences might emerge due to specific patient characteristics, e.g., to avoid bodily harm due to heavy lifting, or due to additional training, or lack thereof, for specific treatments a patient requires. Preference-based scheduling to increase satisfaction has been shown for physicians (Erhard et al. 2018) and nurses (Cheang et al. 2003; Burke et al. 2004), where it usually manifests as preferences for specific shifts. Physical rehabilitation is one area with a serious labor shortage. For instance, physical therapy is listed as an area with a skilled labor shortage in Germany (Bundesagentur für Arbeit 2022). A forecast for the USA predicts a severe physical therapists shortage in 2030 (Zimbelman et al. 2010). The reasons for the shortages are manifold. High levels of moral distress and burnout lead to a high turnover rate (Lau et al. 2016). The extreme time pressure when dealing with treatments and mental burden contribute to a stressful work environment (Girbig et al. 2017). As a result of the serious workforce shortage, management struggles to schedule all needed appointments. Furthermore, for scheduled appointments, the department cannot guarantee preference adherence and continuity of care. In this context, continuity of care (CoC) is described as a patient seeing the same therapist over a sequence of appointments. When disregarding CoC, important treatment time is required when a (new) therapist needs to get acquainted with the patient (Olaleye et al. 2017). Often, in rehabilitation settings, appointment durations are determined by regulations. In Germany, health insurance is required by law to have framework contracts with different rehabilitation professional associations. These contracts specify appointment durations, where preparatory tasks and administrative tasks after the treatment are part of the appointment time (Verband der Ersatzkassen 2022). Consequently, time pressure increases when getting used to a new patient. This negatively influences employee satisfaction, while a lower level of CoC also leads to decreasing patient satisfaction and worse rehabilitation outcomes (Beattie et al. 2005; van Walraven et al. 2010). As a result, therapists tend to have a strong preference for CoC. In this complex planning environment, considering employee preferences for necessary appointments can help acute care hospitals improve rehabilitation objectives.

The paper aims to tackle the operational (offline) appointment scheduling problem in rehabilitation departments in hospitals. Figure 1 shows that the daily appointment scheduling is the last step after a series of earlier planning steps. On the strategic and tactical levels, the size and availability of the workforce, as well as general patient acceptance and prioritization rules, are decided. Before the scheduling takes place on the operational level, therapists must decide their preferences for treatments and patients. Patients must be given a ward bed or an appointment time in case of

outpatients. Finally, the rehabilitation department manager must set appointment priorities depending on the number of patients and their needs.

We consider hundreds of appointments for many therapists, and a schedule must be generated within at most two hours. The department in our partner hospital frequently faces situations where not all potential appointments for rehabilitation can be scheduled in a day. Depending on the importance and the expected rehabilitation outcome, an unscheduled appointment might again be available for scheduling the next day with a higher priority. Alternatively, some patients are released from the hospital without a potential appointment. While beneficial for recovery, these treatments are not expected to worsen the rehabilitation outcome if they are not fulfilled immediately. Instead, the necessary treatments are fulfilled later by specialized rehabilitation hospitals or institutions. Patients are classified into hierarchical *priority classes* while CoC and therapist preferences are represented by hierarchical (*preference*) *violation classes*. The partner hospital prioritizes appointments (i.e., patients) using strict priority guidelines. Within these guidelines, they weigh up the medical importance against the financial importance. For example, an inpatient might have an injury with a very good chance of healing without rehabilitation treatment, but the potential appointment might give the hospital high case payments. A second inpatient might need a very critical appointment for a good rehabilitation outcome, but the appointment only has small financial benefits. The second patient's appointment usually would have the higher priority class according to the priority guidelines. Examples of high-priority appointments are emergencies, critical appointments in the therapy path, or outpatient appointments (with fixed appointment slots) to increase department revenue. We define five priority classes where the most important appointments, e.g., emergencies and outpatients, constitute the highest priority class 5, while the least important appointments constitute the priority class 1. Since priority classes are not strictly motivated financially but also according to medical aspects, it is important to consider them lexicographically. In other words, it is never beneficial to fulfill any number of appointments with a lower priority class in favor of one appointment within a higher priority class. For therapist preferences, we use a similar idea. Besides the importance of CoC and qualification aspects, the therapists might have individual preferences based on specialized training, specific treatment types, or a patient due to patient characteristics. Since patient characteristics are available from patient records and specialized training is part of the employee records, preferences can be determined automatically before the scheduling takes place. We use five different (preference) violation classes based on therapist preferences. We can model CoC and preferences within the same class type, as therapists show a strong preference for CoC. Hence, an existing CoC relation always outweighs potential preferences for the patient from other therapists. For a therapist with an existing CoC relation, CoC also outweighs a potential negative attitude of the therapist against an appointment if a CoC relation would not exist. We apply the same hierarchical concept as for priority classes, i.e., the most important (preference) violation class is 5, and the least important (preference) violation class is 1. Satisfying employees and therefore having lower absenteeism and an attractive workplace for potential new employees is a means for the hospital to handle more appointments. Therefore, minimizing preference violation penalties is a secondary

objective to minimizing the total penalty for missed appointments. Here, the hospital follows a pooling approach, where first all qualified therapists are considered for appointments, to minimize unscheduled appointments according to their priority, before assigning necessary appointments to individual therapists according to preference criteria. Pooling in hospitals is a tactical decision that results in more scheduling flexibility. Kuiper and Lee (2022) show that pooling has benefits when preference criteria such as CoC are relaxed at first. We apply a lexicographic order between the two objectives to model the hospital's approach.

The contribution of the paper is manifold. We develop a new integer programming (IP) model for the operational (offline) appointment scheduling problem at rehabilitation departments focusing on workforce shortage. The IP model matches qualified therapists with patient appointments subject to time window restrictions. Our IP model deals with short periods of 10 min, adding to the combinatorial difficulty. The primary objective is to minimize the sum of unscheduled appointments for each priority class in hierarchical order. The secondary objective is to minimize the sum of scheduled appointments for each (preference) violation class to promote therapist satisfaction. All appointments for a specific day are known in the morning, and a schedule needs to be generated within at most two hours. The physical therapy department of the partner hospital must deal with more than 600 daily appointments. Our computational results show that the IP model can only solve smaller test instances. However, the IP model cannot fulfill operational time requirements of at most two hours, even for medium-sized instances, and becomes intractable for large instances. We develop a customized Greedy Randomized Adaptive Search Procedure (GRASP) as an alternative approach. A GRASP is a metaheuristic approach introduced by Feo et al. (1994). It consists of a construction phase 1 and a local search phase 2. We develop a problem-specific construction phase 1 where we consider different shift lengths and types (with and without lunch breaks) for therapists, precedence relations for appointments, and the important five priority classes and the five (preference) violation classes. In phase 2, we implement six neighborhoods based on relocation and interchange moves. Five of them are known from literature but must be fundamentally adapted to deal with problem-specific features. We introduce a new problem-specific sixth neighborhood dealing with a patient's therapy pathway. Using data representing real-world requirements, we evaluate the performance of our IP model and solution approach. Our computational results show the advantages of the GRASP procedure, in particular for large (realistic) problem instances. We achieve similar performance for small instances as solving the IP model with standard software. Finally, managerial insights describe further benefits of the customized GRASP.

In the following Sect. 2, we examine the state of the art of scheduling in a rehabilitation setting. We discuss related research in home health care scheduling, and we introduce GRASP literature with shared characteristics. Section 3 gives a problem description and introduces the formal IP model. Section 4 presents the customized GRASP algorithm with detailed descriptions of phase 1 and phase 2. Section 5 describes experimental studies and managerial insights are derived. Section 6 gives a summary and discusses future research directions.

2 Literature review

In this paper, we schedule patients' rehabilitation appointments where all required patient appointments are known in advance for the day. Our problem case is free of potential no-shows or walk-ins since we deal with inpatients and outpatient appointments confirmed on short notice. Finally, appointment durations are given as defined by regulations and the department manager, i.e., deterministic service durations as opposed to probabilistic service time distributions. Therefore, our problem is different from most appointment scheduling problems in healthcare, as described by Gupta and Denton (2008) and Ahmadi-Javid et al. (2017), or literature dealing with pooling effects in service industries, such as van Dijk and van der Sluis (2008), van Dijk and van der Sluis (2009), or Song et al. (2015). Nevertheless, there is existing rehabilitation therapy literature with comparable traits.

Table 1 summarizes important aspects of our problem and shows if other rehabilitation research includes these aspects with an X. The column *Time* informs about the planning horizon of the problems. The column *Problem Size* shows how many appointments are considered in the largest problem size. Literature dealing with the patient rehabilitation process is very diverse. For example, some of the following projects are set in an acute care setting, while others are set in specialized rehabilitation hospitals, which enables longer planning horizons. Our review is specifically interested in potential problem attributes originating in workforce shortages, prevalent in the acute care setting and specialized care. For detailed information concerning the individual problem settings and objectives, please consult the cited literature. Podgorelec and Kokol (1997) use a genetic algorithm to schedule patients for therapy. They consider differences in therapist qualifications and therapy pathways, where a patient has more than one appointment a day, and the appointment order is fixed. Compared to our research, which considers 655 potential appointments, the authors only schedule 45 appointments. In their problem, all appointments can be scheduled.

Chien et al. (2008) introduce a genetic algorithm to minimize patient waiting time and makespan. Their problem includes 250 patients, and they lack any consideration of workforce shortages. Similar is true for Zhao et al. (2018). They develop a genetic algorithm to minimize patient waiting time and makespan. The authors include therapy pathways but none of the other considerations for our problem. Huynh et al. (2018) minimize makespan and waiting time using a genetic algorithm. They deal with a larger problem size but do not consider the same real-world requirements we face. Ogulata et al. (2008) are dealing with appointment priority. They solve three mathematical models with time horizons from weekly to daily patient scheduling. Appointment priority for up to 90 potential appointments is only considered in the weekly stage. Griffith et al. (2012) develop a local-search-based approach for a weekly-scheduling problem. The paper considers patient priorities and therapist qualifications. They deal with an occupancy level of 21 patients. Other aspects required for our problem are not considered. Schimmelpfeng et al. (2012) use a three-stage approach to assign rehabilitation appointments to days before assigning daily appointments to time

Table 1 Relevant aspects in rehabilitation therapy scheduling

	Appointment priorities	Therapist preference	Continuity of care	Therapy pathway	Travel time	Different qualification	Time	Problem size
Podgorelec and Kokol (1997)				X		X	Day	45
Chien et al. (2008)				X		X	Day	250
Ogulata et al. (2008)	X						Week to day	90
Griffith et al. (2012)	X					X	Week	21
Schimmelpfeng et al. (2012)			X	X			Month to day	1287
Gartner et al. (2018)				X	X	X	Day	120
Huynh et al. (2018)				X		X	Day	740
Zhao et al. (2018)				X			Day	316
Jungwirth et al. (2021)				X	X	X	Day	120
This Paper	X	X	X	X	X	X	Day	655

slots. In the third step, they match resources to patients. The authors maximize fulfilled appointments without priority consideration. They enforce CoC, and therefore, no workforce shortage is assumed. Gartner et al. (2018) develop a cutting plane algorithm for a daily problem with up to 120 patients. They are considering therapist travel times. However, they do not consider any aspects of workforce shortages or therapist satisfaction and they assume unrealistically long 30-min periods. Jungwirth et al. (2021) use the same problem environment as Gartner et al. (2018), but they model the problem as a vehicle routing problem to minimize the total cost of all selected tours. The problem is solved with an exact branch-price-and-cut algorithm. Realistic 5-min periods are used, but they also only consider 120 patients at most.

To summarize, no other paper considers appointment priority for a daily planning horizon, and we are the first paper to consider therapists' treatment preferences. While Schimmelpfeng et al. (2012) consider CoC, we are the first paper to assume that CoC cannot always be fulfilled and only consider it as a secondary objective as an important part of therapist preferences. Gartner et al. (2018) and Jungwirth et al. (2021) realistically use travel times for therapists between patients, but they lack any of our workforce-related assumptions.

Due to travel times, there are similarities to home health care routing and scheduling. Cissé et al. (2017) and Fikar and Hirsch (2017) describe the home health care problem (HHCP) as patients scattered across a region requiring health services. Health service workers visit the patients at home to fulfill patient needs. Cissé et al. (2017) describe CoC as one of the most important aspects of HHCP. Cappanera and Scutellà (2015) consider CoC as a constraint with only a certain number of therapists working on a patient. However, their objective is to balance the therapists' workload, and they do not have a daily planning horizon. Bowers et al. (2015) examine the effects of focusing on CoC for travel times when visiting patients. Wirnitzer et al. (2016) describe several possibilities for a CoC-centered objective function for a monthly problem. CoC is in the objective function in Nickel et al. (2012). In a weekly problem, they also consider the overtime costs, travel distances, and the number of unscheduled tasks without prioritization. In many regulatory settings, the service provider is forced to fulfill all patient tasks. Therefore, only a few HHCP papers deal with the possibility of missed appointments (Fikar and Hirsch 2017). Dohn et al. (2009) consider a daily planning problem where they maximize the number of fulfilled tasks. They do not consider prioritization or employee preferences. Rasmussen et al. (2012) consider this a daily problem. They minimize prioritized uncovered visits, caretaker preferences, and travel costs. Different from this paper, the authors consider synchronization, where a task must be fulfilled by several caregivers. To summarize, to the best of our knowledge, no published research in HHCP or rehabilitation scheduling considers the required objectives and constraints while fulfilling the necessity of a fast solution approach for larger problem sizes. Therefore, we develop a new mathematical model and a GRASP.

A GRASP is a two-phase procedure that can find good solutions to complex problems quickly. First, in the construction phase, a feasible solution is constructed. In each iteration of the first construction phase, an insertion candidate is randomly chosen from a restrictive candidate list. The candidate list consists

of potential best-performing insertion candidates, according to a predefined greedy function. There is no guarantee that the solutions achieved in the first phase are locally optimal concerning predefined neighborhoods. In the second phase, local search is used to improve initial solutions concerning these neighborhoods. Within a neighborhood, a neighbor is a solution close to the current solution, which can be reached by a predefined change operation (e.g., moving an appointment from one therapist to another therapist) (Feo et al. 1994). We are aware that there are many implementations of GRASP in literature, ranging from problems like slab scheduling in the steel industry in Wichmann et al. (2014) to examples in the healthcare sector, e.g., in Ait Haddadene et al. (2016). Here, we introduce two procedures with the most shared characteristics to our problem in more detail. Ait Haddadene et al. (2016) develop a GRASP with iterated local search for an HHCP with synchronization. They consider precedence constraints and preference penalties for patient-caregiver combinations as part of their objective function. Minimizing unscheduled appointments is not an objective. Rojanasoonthon and Bard (2005) focus on job priorities. They describe a problem where spacecraft with different priorities require satellite time to communicate with ground terminals within time windows. The problem is modeled as a parallel machine-scheduling problem with nonhomogeneous machines. For priorities, their problem is comparable to ours. The objective is to maximize the number of scheduled jobs, and jobs within a higher priority class are lexicographically preferred to jobs from lower priority classes. They consider setup times on a machine, comparable to transition times between patients, and time windows for spacecraft operations. Some spacecraft with fixed time windows are comparable to outpatients in our problem. Nevertheless, several differences require adaptations for our GRASP. An appointment is planned with the same treatment duration by all different qualified therapists. Spacecraft are serviced at different speeds on different satellites. Furthermore, precedence constraints, as well as preferences, are not considered in Rojanasoonthon and Bard (2005). To summarize, while the general framework of the priority-focused GRASP by Rojanasoonthon and Bard (2005) is generally suitable to apply to our problem, substantial adaptations are necessary. We alter the existing GRASP by including (preference) violation classes, thereby changing the first phase. In our problem, each qualified therapist requires the same time for a specific appointment. A sole focus on time slack as a fitness metric, such as in Rojanasoonthon and Bard (2005), is impractical. Instead, we focus on (preference) violation classes as a fitness metric. Furthermore, therapy pathways must be respected, or infeasible solutions might result. This leads to more restricted possible insertion points in our problem. For the local search, therapy pathway constraints and different (preference) violation classes for a therapist-patient combination instead of time slacks lead to changes in the five general neighborhoods used by Rojanasoonthon and Bard (2005). To further deal with therapy pathway constraints, an additional specific neighborhood is introduced. Before we describe our GRASP in detail, the next section introduces the problem definition and the mathematical model.

3 Problem definition and mathematical model

Every morning, the planner considers all the available therapists $e \in E$. After the shift planning at a higher level (see Fig. 1 in Sect. 1), up to 50 therapists work on a typical weekday. Appointment scheduling is the chronologically last in a series of operational planning steps, as shown in Fig. 1. Therapists are qualified for one qualification Q_e . For each therapist e , Q_e holds one item $a, a' \in A$ of four possible appointment types (physical therapy, sports therapy, occupational therapy, and massage treatments). Several different shift types span from a 5-h shift without a break up to an 8-h shift, including lunch breaks of varying lengths. We introduce 10-min periods $t, t' \in T$ to model a planning day. For each therapist e , the subset $T(e)$ includes the periods the therapist can work on patient appointments.

Each required appointment for the day is known in the morning. Set P includes individual patients p, p' . A patient has at most one appointment in each of the four types: sports therapy, massage treatments, occupational therapy, and physical therapy. A patient might have appointments in more than one type, e.g., one appointment in physical therapy followed by a massage later. Here, subset $A(p)$ includes the required appointment types a for every patient p . If a patient has more than one daily appointment, we use a therapy pathway (Gartner et al. 2018) to avoid patient harm if an appointment follows the predecessor too quickly. Subset $F(p, a)$ includes all predecessor appointments for a patient's appointments $a \in A(p)$. Parameter $W_{p,p'}$ holds the transition time of a therapist between two patients p and p' . $S_{p,a}$ holds the treatment time in multiple of periods for an appointment $a \in A(p)$ of a patient p . While we only have four appointment types, treatment times within one appointment type are individual to a patient. For example, a physical therapy appointment might be a short checkup of 10 min (i.e., 1 period) or a longer treatment of 30 min (i.e., 3 periods). These treatment times are deterministic and follow rules given by framework contracts for the rehabilitation professions. $D_{p,a,a'}$ is the recovery time between appointments a and a' in a patient's therapy pathway. Individual appointments

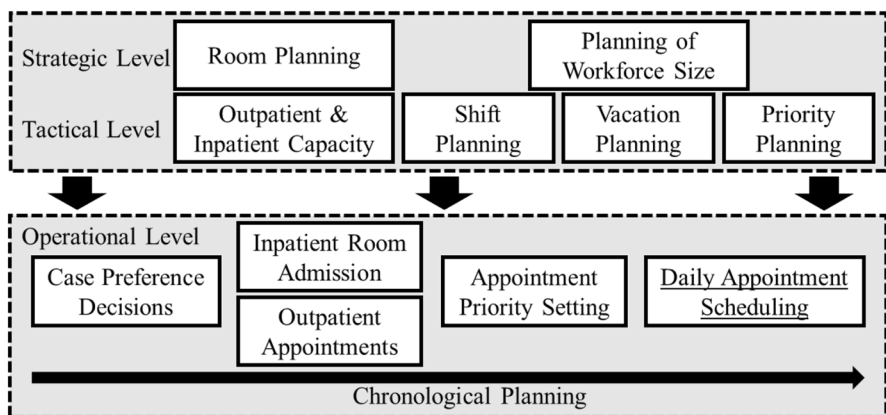


Fig. 1 Decision process environment of daily appointment scheduling

usually span more than a one-time period $t \in T$. Subset $T(p, a)$ includes the time window for each $a \in A(p)$ for a patient. Outpatient appointment time is inflexible, i.e., the exact start time is determined in an earlier step. The possible time window of an inpatient appointment can be very broad. The lower hierarchical importance of therapists restricts time windows compared to, e.g., physicians. The time window is chosen so the patient is available in their room. In the introduction, we described that the hospital fulfills appointments according to priority. To model this, we introduce priority classes $b, b' \in B = \{1, 2, 3, 4, 5\}$ and subset $A(p, b)$ of appointments $a \in A(p)$ of patient p with priority class b . The cost $C_{p,a}^{\text{prio}}$ for not fulfilling an appointment can then be defined as follows.

$$C_{p,a}^{\text{prio}} > \sum_{p' \in P} \sum_{a' \in A(p', b')} C_{p',a'}^{\text{prio}} \quad \forall p \in P, a \in A(p, b), b, b' \in B : b = b' + 1 \quad (1)$$

Inequalities (1) define an appointment's penalty cost $C_{p,a}^{\text{prio}}$ within priority class $b > 1$ as higher than the sum of penalty costs within the lower, less important priority class. Remember that priority class 5 is the most important class with the highest cost weight. When minimizing unscheduled appointments, we first schedule patients with high-priority weights. Assume a total of three appointments in class $b' = 1$. If $C_{p,a}^{\text{prio}} = 0.5$ for these three appointments, then $C_{p,a}^{\text{prio}} > 1.5$ for appointments in the next higher priority class $b = 2$, to ensure prioritization over the appointments in priority class $b' = 1$. We assume all $C_{p,a}^{\text{prio}}$ within a class b are equal due to hospital policy and costs only have positive values. This means there is no differentiation between appointments of the same patient for $C_{p,a}^{\text{prio}}$ nor are there differences regarding appointments of two or more patients if the appointments are in the same priority class.

Preference consideration is the second important objective. We introduce (preference) violation classes $r, r' \in R = \{1, 2, 3, 4, 5\}$. (Preference) violation class 5 is the most important (preference) violation class, i.e., appointments a therapist likes to work the most. It has the lowest cost weight due to our minimization objective. Subset $A(p, r)$ includes appointments $a \in A(p)$ of patient p within preference violation class r . For preference violation cost parameter $C_{e,p,a}^{\text{violation}}$, a therapist e prefers all appointments within more important preference violation classes over one appointment within a less important violation class, i.e., $C_{e,p,a}^{\text{violation}}$ for appointments within $r = 5$ are preferred to $C_{e,p,a}^{\text{violation}}$ within violation class $r = 4$. Further, we also consider CoC within cost parameter $C_{e,p,a}^{\text{violation}}$. To formalize the role of CoC within cost parameter $C_{e,p,a}^{\text{violation}}$, we introduce parameter $E_{p,a}$. The parameter holds a therapist e for an appointment $a \in A(p)$ if the therapist treated the patient last. Here we also differentiate between the four appointment qualification types. In the introduction, we described that a therapist generally prefers a continued CoC relation, even if the therapist normally would not give a patient's appointment the highest preference. Let appointments A_1, A_2, A_3 , and A_4 be appointments a therapist might fulfill within the same priority class. The therapist has a CoC relation with appointment A_4 . The therapist likes to fulfill the appointment A_1 the most, i.e., the therapist has a high preference for it. Appointments A_2 and A_3 are more bothersome for the therapist,

i.e., the therapist has a lower preference for them than for A_1 . The therapist would generally prefer appointments A_1 , A_2 , and A_3 over A_4 but as described, there is a CoC relation between the therapist and appointment A_4 . In the introduction, we described that changing patients between therapists leads to a high mental workload for the new therapist since the therapist must get acclimated with the patient within the treatment time given by the framework contract. Therefore, we want to ensure that appointment A_4 is done by the current therapist. This leads to an order of $A_4 \geq A_1 > A_2 \sim A_3$ for the therapist, with A_4 being in preference violation class $r = 5$. For appointments A_2 and A_3 , it is also not necessarily the case, that other therapists would have a high preference for them. The rationale behind this is that we try to adhere to real preferences. This enables the department management to have a long-term view of the preference adherence for individual therapists as well as the qualification types in total. To formalize the concepts, inequalities (2) and (3) define $C_{e,p,a}^{\text{violation}}$.

$$C_{e,p,a}^{\text{violation}} > \sum_{p' \in P} \sum_{a' \in A(p', r')} C_{e,p',a'}^{\text{violation}} \quad \forall e \in E, p \in P, a \in A(p, r), r, r' \in R : r = r' - 1 \quad (2)$$

$$C_{e',p,a}^{\text{violation}} < C_{e,p,a}^{\text{violation}} \quad \forall e, e' \in E : e \neq E_{p,a}, e' = E_{p,a}, p \in P, a \in A(p) \quad (3)$$

Inequalities (2) ensure for combinations of therapist e with $a \in A(p)$ of patients p that $C_{e,p,a}^{\text{violation}}$ within less important (preference) violation classes is higher than penalty costs of combinations within more important violation classes. $C_{e,p,a}^{\text{violation}}$ is the same for all patient appointments within the same class r and values for $C_{e,p,a}^{\text{violation}}$ of therapist appointment combinations in $r = 1$ are fixed to the same value. Inequalities (3) deal with a continuity of care relationship between a therapist and a patient. If a therapist treated the patient the last time, the penalty cost $C_{e,p,a}^{\text{violation}}$ of this combination must be lower than for any of the remaining possible combinations between the patient's appointment and therapists. This means, preference class $r = 5$ for a given CoC relation between an appointment and a therapist, while $r < 5$ for all remaining qualified therapists in combination with the appointment. For a therapist with a CoC relation, class $r = 5$ is not only formed with CoC appointments. Additionally, appointments without CoC relations with other therapists can be part of class $r = 5$, if the therapist has a very high preference for them. The cost parameters $C_{e,p,a}^{\text{violation}}$ and $C_{p,a}^{\text{prio}}$ are calculated in pre-processing and used as parameters for the mathematical model.

Outpatient room capacity is planned in an earlier step, as outpatients are only accepted if the designated treatment room is free. Inpatients are treated in their wardrooms. Hence, we do not require a parameter for room availability. The same holds for other resources.

We use three different decision variables. (Auxiliary) binary variable $x_{p,a} = 1$ if an appointment $a \in A(p)$ of patient p is not fulfilled, i.e., unscheduled within its time window. Ideally, variable $x_{p,a} = 0$ for important priority classes with class 5 as the most important class. Binary decision variable $z_{e,p,a,t}^{\text{start}} = 1$ if a therapist e starts appointment $a \in A(p)$ of patient p in period t . The decision variable can only be 1

if the patient has been scheduled, i.e., $x_{p,a} = 0$. Finally, the (auxiliary) binary variable $z_{e,p,a,t}$ works together with variable $z_{e,p,a,t}^{\text{start}}$. While $z_{e,p,a,t}^{\text{start}}$ indicates the exact start of an appointment if a therapist e is occupied with $a \in A(p)$ of patient p in period t , $z_{e,p,a,t} = 1$ for each consecutive period an appointment takes place.

$$\min \sum_{p \in P} \sum_{a \in A(p)} C_{p,a}^{\text{prio}} \cdot x_{p,a} \quad (4.1)$$

$$\min \sum_{e \in E} \sum_{p \in P} \sum_{a \in A(p)} \sum_{t \in T} C_{e,p,a}^{\text{violation}} \cdot z_{e,p,a,t}^{\text{start}} \quad (4.2)$$

s.t.

$$\sum_{p \in P} \sum_{a \in A(p)} z_{e,p,a,t} \leq 1 \quad \forall e \in E, t \in T \quad (5)$$

$$\sum_{e \in E} z_{e,p,a,t} \leq 1 \quad \forall p \in P, a \in A(p), t \in T(p, a) \quad (6)$$

$$x_{p,a} = 1 - \sum_{t \in T} \sum_{e \in E} z_{e,p,a,t}^{\text{start}} \quad \forall p \in P, a \in A(p) \quad (7)$$

$$z_{e,p,a,t}^{\text{start}} \geq z_{e,p,a,t} - z_{e,p,a,t-1} \quad \forall e \in E, p \in P, a \in A(p), t \in T \quad (8)$$

$$\sum_{t \in T(e)} z_{e,p,a,t} = \sum_{t \in T(e)} S_{p,a} \cdot z_{e,p,a,t}^{\text{start}} \quad \forall e \in E, p \in P, a \in A(p) \quad (9)$$

$$z_{e,p,a,t}^{\text{start}} + \sum_{t \leq t' \leq t-1+S_{p,a}+W_{p,p'}} z_{e,p',a',t'}^{\text{start}} \leq 1 \quad \forall e \in E, p \in P, a \in A(p), p' \in P \setminus \{p\}, a' \in A(p'), t \in T(e) \quad (10)$$

$$\sum_{t \in T: t < t'} \sum_{e \in E} z_{e,p,a',t}^{\text{start}} \geq \sum_{e \in E} z_{e,p,a,t'}^{\text{start}} \quad \forall p \in P, a \in A(p), a' \in F(p, a), t' \in T \quad (11)$$

$$\begin{aligned} & \sum_{t \in T(p,a')} \sum_{e \in E} \left(t \cdot z_{e,p,a',t}^{\text{start}} + (S_{p,a'} + D_{p,a',a}) \cdot z_{e,p,a',t}^{\text{start}} \right) \\ & \leq M \left(1 - \sum_{t' \in T(p,a)} \sum_{e \in E} z_{e,p,a,t'}^{\text{start}} \right) + \sum_{t' \in T(p,a)} \sum_{e \in E} t' \cdot z_{e,p,a,t'}^{\text{start}} \quad \forall p \in P, a \in A(p), a' \in F(p, a) \end{aligned} \quad (12)$$

$$z_{e,p,a,t} = 0 \quad \forall e \in E, p \in P, a \in A(p), t \in T \setminus T(e) \quad (13)$$

$$\sum_{e \in E} z_{e,p,a,t} = 0 \quad \forall p \in P, a \in A(p), t \in T \setminus T(p, a) \quad (14)$$

$$z_{e,p,a,t} = 0 \quad \forall e \in E, p \in P, a \in A(p) : a \neq Q_e, t \in T \setminus T(e) \quad (15)$$

$$x_{p,a}, z_{e,p,a,t}, z_{e,p,a,t}^{\text{start}} \in \{0, 1\} \quad \forall e \in E, p \in P, a \in A(p), t \in T \quad (16)$$

Sets

$e \in E$	Set of therapist employees
$p, p' \in P$	Set of patients
$a \in A$	Set of appointment types
$a, a' \in A(p)$	Set of appointments of patient p
$F(p, a)$	Subset of predecessor appointments for $a \in A(p)$ for a patient p
$t, t' \in T$	Set of time periods
$T(e)$	Subset of time periods a therapist is available for appointment scheduling
$T(p, a)$	Subset of time periods an appointment a of patient p can take place

Parameters

$W_{p,p'}$	Minimum number of periods between two patients p and p' , i.e., the transition time of a therapist (e.g., include walking distances as well as the varying administrative time between patients)
$S_{p,a}$	Number of treatment periods for appointment $a \in A(p)$ of patient p
$D_{p,a,a'}$	Minimum number of periods between appointments a and a' of one single patient p (e.g., needed to avoid mental or physical stress to a patient caused by follow-up appointments following too fast)
Q_e	Qualification type of therapist e
$C_{p,a}^{\text{prio}}$	Hierarchical priority penalty for missed appointments $a \in A(p)$ of patient p
$C_{e,p,a}^{\text{violation}}$	Hierarchical penalty cost for therapist preference violations for therapist e when being responsible for an appointment $a \in A(p)$ of patient p
M	Sufficiently large value

Decision variables

$x_{p,a}$	1 if an appointment $a \in A(p)$ of patient p is not fulfilled within its time window, 0 otherwise
$z_{e,p,a,t}^{\text{start}}$	1 if a therapist e starts appointment $a \in A(p)$ of patient p in period t , 0 otherwise
$z_{e,p,a,t}$	1 if a therapist e is occupied with appointment $a \in A(p)$ of patient p in period t , 0 otherwise

Fulfilling potential appointments is seen as more important than satisfying therapist preferences. Therefore, we model the two objectives lexicographically. Term (4.1) minimizes the penalty of unscheduled appointments, i.e., setting decision variables $x_{p,a}$ to 0 starting at the highest priority class 5 down to class 1. While the term (4.2) minimizes the preference and CoC violations for scheduling appointments with a therapist, i.e., setting decision variables $z_{e,p,a,t}^{\text{start}}$ to 1 starting with the most important (preference) violation class 5 down to class 1. Constraints (5) ensure that each therapist works at most one appointment in a time period. Constraints (6) enforce that at most one therapist is working on a specific appointment. Constraints (7) make sure that any appointment starts at most once if it is not unscheduled, i.e., $x_{p,a} = 1$. Constraints (8) indicate the start of an appointment. If $z_{e,p,a,t-1} = 0$ and an appointment started in period t , i.e., $z_{e,p,a,t} = 1$, then $z_{e,p,a,t}^{\text{start}} = 1$. Constraints (9) ensure the

exact length of an appointment. Constraints (10) ensure the minimum transition time between two patients. If a therapist directly treats one patient before the other, transition periods are needed. Constraints (11) deal with patients with more than one appointment a day. An appointment only takes place if all predecessors are already scheduled. Constraints (12) enforce that the recovery time between two appointments of a patient is fulfilled. A successor can be scheduled at the earliest after a preceding appointment has ended, plus the recovery time. The parameter M is used to ensure a predecessor can be scheduled, even if a potential successor cannot be scheduled. For this, it must ensure that all potential predecessors can be scheduled in the entirety of their time window. We use $M = 65$ since the latest possible end of a time window is 58 and the longest necessary recovery time is one hour, i.e., $D_{p,a,a'} = 6$. Constraints (13) ensure therapists only work within their shift assignment. Constraints (14) guarantee that appointments can only take place within their time window. The consideration of qualifications is ensured by constraints (15). Finally, constraints (16) define the binary decision variables.

Related machine scheduling problems to minimize the weighted number of tardy jobs are NP-hard even for a single machine with no setup times or release dates (Pinedo 2016). Parallel machine cases, such as our problem with several different therapists, where a therapist represents a machine and each appointment refers to a job, add more complexity due to the interplay between (therapist) schedules and the added scheduling possibilities for individual appointments. Start times, times between patient visits (i.e., setup times), and therapy pathways as well as different qualifications lead to a further increase in complexity (Pinedo 2016). First experiments (for more information see Sect. 5) show the commercial solver Gurobi (Gurobi Optimization 2022) cannot solve the problem at hand in the required time of up to two hours for smaller problem sizes. The solution becomes intractable for larger problem sizes. Accordingly, we develop a customized GRASP as an alternative solution method.

4 Greedy randomized adaptive search procedure

A GRASP is a two-phase procedure that can find high-quality solutions to complex problems in a short time. In phase 1, a feasible solution is constructed. An insertion candidate is determined in each iteration by building a candidate list out of best-performing candidates, according to a predefined greedy function. Then an item from the candidate list is chosen at random. There is no guarantee that the solutions after the construction phase are a global optimum or a local optimum regarding predefined neighborhoods. Therefore, in phase 2, a local search is used to improve initial solutions regarding six predefined neighborhoods. The efficiency of the second phase improves with initial solutions close to a local optimum (Kontoravdis and Bard 1995). In the literature review, we described that the GRASP introduced by Rojanasoonthon and Bard (2005) has a similar objective to our first objective stated in (4.1). We, therefore, use a comparable high-level design. The high-level architecture is illustrated in Fig. 2 and formally described in the following.

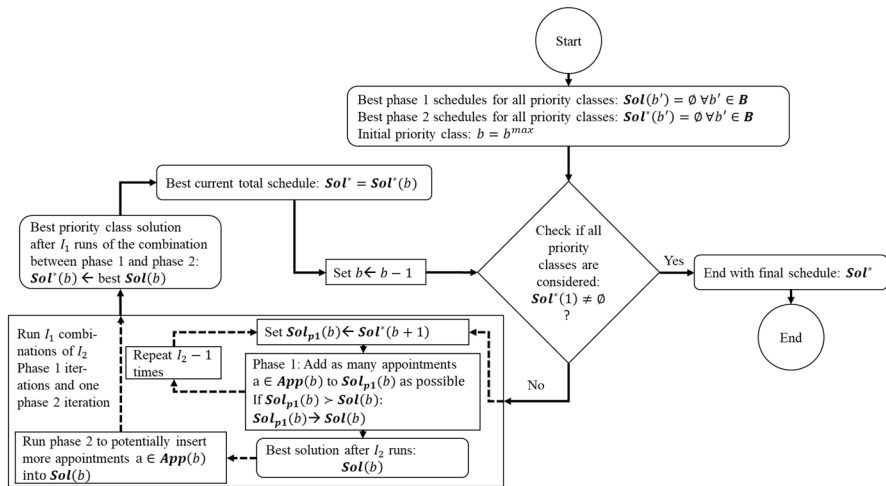


Fig. 2 High-level architecture of the customized GRASP

The GRASP is run for a subset of appointments $\mathbf{App}(b) \subset \mathbf{App}$ within a priority class $b \in B$. Each combination between appointment type $a \in \mathbf{A}(p)$ and patients $p \in P$ from the introduced IP in Sect. 3 is considered in the set of all appointments $a \in \mathbf{App}$ and therefore in one of the subsets $\mathbf{App}(b)$. We start with the most important priority b^{\max} (i.e., 5 in our case), and within the construction phase—phase 1—we add appointments to an initially empty set $\mathbf{Sol}_{p1}(b^{\max})$. We run the first phase I_2 times and the preliminary solution schedule $\mathbf{Sol}(b)$ represents the best $\mathbf{Sol}_{p1}(b)$. Solution schedules $\mathbf{Sol}(b)$ (or $\mathbf{Sol}_{p1}(b)$, respectively) contain all current schedules $\mathbf{Sol}(e)$ of any therapist e , i.e., $\mathbf{Sol}(e) \subset \mathbf{Sol}(b)$. We run the local search (phase 2) only on the best preliminary solution $\mathbf{Sol}(b)$, which results after several I_2 iterations of phase 1. Phase 2 tries to add unscheduled appointments from $\mathbf{App}(b)$ to $\mathbf{Sol}(b)$. Potential effects and drawbacks are discussed in Sect. 5.3. The process of several phase 1 iterations followed by a local search (phase 2) on the best solution is repeated for I_1 iterations within a priority class to guarantee a diverse set of good solutions. The numbers of I_1 and I_2 iterations were chosen with a tradeoff between time requirements and solution stability in mind. Details are described in Sect. 5.2.3. The best achieved final schedule after a priority class was examined is represented by $\mathbf{Sol}^*(b)$. The $\mathbf{Sol}^*(b)$ for a current priority class is then used as input for the next less important priority class $b - 1$. This means $\mathbf{Sol}_{p1}(b)$ for $b < b^{\max}$ already includes schedule $\mathbf{Sol}^*(b + 1)$ and phase 1 adds to the corresponding generated schedules of therapists. The process is repeated until the algorithm is run for the lowest priority class $b = 1$. In the end, the solution output $\mathbf{Sol}^*(1)$ is accepted as final best solution \mathbf{Sol}^* . A pseudo-code of the high-level architecture of the described GRASP can be seen in Appendix A. In the following, assumptions of the GRASP are described. These include therapist availability and rules for appointment insertion and start time updates. Afterward, we go into more detailed descriptions of the construction phase and the defined neighborhoods within the local search.

4.1 Definitions concerning insertion possibilities and time requirements

In this section, we describe concepts that are relevant for the construction phase as well as for the local search of the GRASP. First, we formalize how we include the possibility of lunch breaks into the algorithm. Further, we describe how we achieve feasible results in every step of the algorithm. Second, we describe how time window restrictions of appointments are considered when inserting the appointments in therapist schedules. Third, we describe how the start time of scheduled appointments is updated if another appointment is inserted ahead of already scheduled appointments.

Therapist and Appointment Availability. For some therapists, we must account for lunch breaks. In the GRASP, we view the shift time before and after the lunch break as two individual therapists $e \in E$. These representations of the real therapist have the same parameter values, e.g., for violation costs or qualification. Further, our algorithm only allows feasible schedules during the GRASP. Hence, appointment $a \in A(p) \cap \mathbf{App}(b)$ is only scheduled if the predecessor appointments $a' \in F(p, a)$ are part of $\mathbf{Sol}_{p_1}(b)$ in phase 1, $\mathbf{Sol}(b)$ if an appointment is inserted in phase 2, or if $F(p, a) = \emptyset$ from the start. Predecessors have a higher or equal priority to successor appointments due to priority rules. Hence, it is impossible that an appointment cannot be scheduled due to precedence restrictions from the start.

Appointment Insertion. An appointment can be scheduled with a qualified therapist if time restrictions are met. Assume T_e^{earliest} is the first period in a therapist's shift and T_a^{earliest} is the first period in the appointment time window. The earliest time an appointment can be scheduled with a therapist is $T_{e,a}^{\text{earliest}} = \max\{T_e^{\text{earliest}}, T_a^{\text{earliest}}\}$. Similar is true for the last therapist and appointment time windows, T_e^{latest} and T_a^{latest} . The last possible starting time for an appointment with a therapist is defined as $T_{e,a}^{\text{latest}} = \min\{T_e^{\text{latest}} + 1 - S_a, T_a^{\text{latest}} + 1 - S_a\}$ where S_a is an appointment's treatment time. We also consider recovery times between predecessors and an appointment. Let $t_{a'}^{\text{start}}$ be the scheduled start of a predecessor a' of appointment $a \in \mathbf{App}(b)$. The earliest start time of an appointment a is restricted by $t_a^{\text{start}} \geq \max\{t_{a'}^{\text{start}} + S_{a'} + D_{a',a} \mid a' \in F(a)\}$ where $D_{a',a}$ is the recovery time between appointment a of a patient and the predecessor appointment a' within set $F(a)$. Within the schedule of a therapist e , we must consider time relations. Assume appointment a is not the first scheduled appointment a therapist must visit. Then the relation $t_{\text{pre}(a)}^{\text{start}} + S_{\text{pre}(a)} + W_{\text{pre}(a),a} \leq t_a^{\text{start}}$ must hold, where $\text{pre}(a)$ is the appointment directly preceding a and $W_{\text{pre}(a),a}$ is the transition time a therapist needs between two appointments (patients) $\text{pre}(a)$ and a . At any time, an appointment is scheduled, it is scheduled in the first possible period to avoid unnecessary time slack before appointments. These restrictions for all appointments are summarized in Eq. (17).

$$t_a^{\text{start}} = \max \left\{ T_{e,a}^{\text{earliest}}, \max \{ t_{a'}^{\text{start}} + S_{a'} + D_{a',a} \mid a' \in F(a) \}, t_{\text{pre}(a)}^{\text{start}} + S_{\text{pre}(a)} + W_{\text{pre}(a),a} \right\} \quad (17)$$

Appointment Time Update. When a new appointment is scheduled ahead of already scheduled appointments in phase 1 or phase 2, we might have to move the succeeding appointments forward in time to insert the new appointment. Here, we

consider $T_{e,a}^{\text{latest}}$ of the candidate's appointment with a therapist and the scheduled start times of scheduled succeeding appointments. An appointment a can be inserted in a position on the schedule $\text{Sol}(e)$ if we have an existing gap of $S_a + W_{a,\text{suc}(a)}$, where $\text{suc}(a)$ is the appointment following appointment a , or if we can generate the gap by moving forward succeeding appointments $\text{suc}(a)$ without violating the succeeding appointments' time restrictions. Appointments that need to be moved forward in time for additional insertions might be part of a patient's therapy pathway. Therefore, some appointments that should be scheduled after these appointments in the therapy pathway might already be in the schedule of a different therapist. Infeasibility might be the result. Let $a' \in G(a)$ be the set of already scheduled appointments within a therapy pathway, that need to come after an appointment a that must be moved forward in time. For the scheduled start of a it must be true that $t_a^{\text{start}} \leq \min\{t_{a'}^{\text{start}} - S_a - D_{a,a'} \forall a' \in G(a)\}$ at all times. Inequality (18) ensures the discussed time restrictions.

$$t_a^{\text{start}} \leq \min\left\{T_{e,a}^{\text{latest}}, \min\{t_{a'}^{\text{start}} - S_a - D_{a,a'} \forall a' \in G(a)\}, t_{\text{suc}(a)}^{\text{start}} - S_a - W_{a,\text{suc}(a)}\right\} \quad (18)$$

When a new appointment is scheduled, it must respect equation (17) while remaining within the time limits set by inequality (18). Before a new insertion in $\text{Sol}(e)$ at a certain position, we preliminary update the start times of scheduled appointments with later positions according to equation (17). Only if the updated start times satisfy inequality (18), we accept the updated t_a^{start} of all later appointments for the considered therapist and insert the new appointment. Formulas (17) and (18) are valid for phase 1 (construction) and the local search in phase 2.

4.2 Phase 1—construction

It is important to remember that outpatient appointments with fixed timeslots are considered. Here only one appointment per patient is possible. As described in the introduction, all outpatients are part of the highest priority class 5, i.e., it is possible to first schedule all outpatients before scheduling inpatients with varying time windows. In this section, we describe how we schedule outpatients before dealing with inpatients.

Outpatients. Let $a^{\text{fixed}} \in \text{App}^{\text{fixed}}$ be an outpatient appointment where $T_{a^{\text{fixed}}}^{\text{earliest}} = T_{a^{\text{fixed}}}^{\text{latest}}$ for all $a^{\text{fixed}} \in \text{App}^{\text{fixed}}$. We first select an appointment from $\text{App}^{\text{fixed}}$ at random to achieve variety between I_2 iterations. Let $E(a)$ be the set of qualified therapists for an appointment a . For the selected appointment, the therapists $e \in E(a^{\text{fixed}}, r) \subset E(a)$ are chosen, where $r \in R$, the (preference) violation class of a therapist for treating an outpatient, is the most important remaining class, i.e., we choose therapists who like to work on an appointment the most. We define $Aq(e)$ as the appointments a therapist can be responsible for. Therapists $e \in E(a^{\text{fixed}}, r)$ are sorted in non-decreasing order according to the number of appointments for which the time window overlaps with the time window of the fixed-time appointment under consideration. We try insertion on the first therapist of the sorted set. This procedure

is chosen as fixed-time appointments are difficult to reschedule in a local search as many of the chosen neighborhoods cannot make a feasible move. We schedule the appointments within the most important (preference) violation class from class 5 to 1 while considering the opportunity cost of not scheduling other appointments in the specific time window of a therapist. By doing so, we avoid bad, unrepairable initial solutions. An appointment that cannot be scheduled after testing an insertion for all therapists down to the least important (preference) violation class $r = 1$ is unscheduled and therefore becomes part of set $\mathbf{App}^{\text{unscheduled}}(b) = \emptyset$. A pseudo-code in Appendix B summarizes the described outpatient procedure, and Appendix C shows the complete construction phase, including inpatient scheduling as described below.

Inpatients. After the outpatients are scheduled in priority class $b^{\max} = 5$, inpatient appointments with larger time windows $T_a^{\text{earliest}} < T_a^{\text{latest}}$ are considered. For inpatients, it is important to only consider appointments for scheduling when predecessors in the therapy pathway are already scheduled, or there are no pathway relations. For these appointments, we determine the possible therapists $e \in E(a, r) \forall a$, where $r = r^{\max}$ in the first iteration, i.e., we determine possible therapists where the therapist prefers to work an appointment the most. Let $C(r)$ be the set of resulting possible combinations (e, a) with $e \in E(a, r)$. We consider minimizing preference violations together with minimizing unscheduled appointments according to priority in the construction phase, i.e., phase 1. In contrast, the local search in phase 2 focuses on improving objective (4.1), i.e., scheduling as many appointments as possible concerning priority classes. Since we have a tradeoff between scheduling additional appointments according to priority (objective (4.1)) and preference adherence (objective (4.2)), this can come at the cost of additional preference violations for therapists after phase 2 compared to the schedule after phase 1. This effect can be seen for different instances in Table 14 in Appendix G. A possible candidate for insertion is randomly chosen from $C(r)$. For the chosen candidate combination, we try to schedule appointment a on schedule $\mathbf{Sol}(e)$ in positions $n = \{1, \dots, |N|\}$ starting with position 1 under consideration of equation (17). If the appointment can be scheduled, it is added to $\mathbf{Sol}(e) \subset \mathbf{Sol}_{p1}(b)$ and removed from the list of possible appointments. Potential successors are instead added to the possible appointments in the next iteration if the scheduled appointment was the last unscheduled predecessor in $F(a)$. If the appointment cannot be scheduled, the combination (e, a) is removed from further consideration and $C(r)$ is updated for a second iteration. If an appointment cannot be scheduled in any therapist's schedule $\mathbf{Sol}(e)$, the appointment is removed from the possible appointments. If all possible combinations in a (preference) violation class r are checked and there are appointments left to consider within a priority violation class b , the index r is set to the next lower, less important, class ($r \leftarrow r - 1$) and resulting combinations in $C(r)$ are checked for insertion. We repeat the process until all appointments in a priority class b are scheduled or every combination in $C(r = 1)$ is tested. Appointments $a \in \mathbf{App}(b)$ within a priority class b which are not scheduled after the first phase are included in the set of unscheduled appointments $\mathbf{App}^{\text{unscheduled}}(b)$. The local search tries to insert these appointments.

4.3 Phase 2—local search

We improve our solutions by relocating a subset of appointments to gain space for an additional appointment and by interchanging appointments with each other. Compared to Rojanasoonthon and Bard (2005), we require adaptations to the neighborhoods and an additional neighborhood (Neighborhood 3) since our problem has to deal with precedence relations due to therapy pathways. Further, we consider therapist satisfaction. In total, we utilize six different neighborhoods where the local search is run for schedules of each priority class $b \in \mathbf{B}$. Within each neighborhood the search restarts based on the updated solution if an insertion was found. Only if no possible insertion was detected, did we go to the next neighborhood. A high-level architecture of phase 2 can be seen in Appendix D.

(Relocation) Neighborhood 1: This neighborhood (see an example using a sequence of two appointments in Fig. 3) relocates a sequence of one, two, or three appointments in a therapist's schedule. The search starts with feasibility checking. Here, all sequence lengths are searched for alternative schedules. Next improvement checking for potential insertions is done. If an insertion was found, the search restarts with feasibility checking. Alternative schedules for all sequence lengths must be updated for a therapist if an insertion occurred on the therapist's schedule. If no more insertions are found, the neighborhood search ends and a new neighborhood is selected.

In the feasibility checking step, the sequence of appointments from appointment a up to a sequence of three appointments $(a, \text{suc}(a), \text{suc}(\text{suc}(a)))$ is relocated backward and forward into positions $j, k \in \{1, \dots, |N|\}$ in $\text{Sol}(e)$. Feasibility checking of a move cannot only consider appointments scheduled for an individual therapist. We must respect restrictions set by potential predecessors and successors in the therapy pathway of an appointment. Therefore, when we, e.g., relocate a sequence $(a, \text{suc}(a))$ backward, and we get the resulting schedule $(0, \dots, j, a, \text{suc}(a), j+1, \dots, |N|)$. Note, the starting times of appointments up to position j remain unchanged while we preliminary update t_a^{start} for all appointments in positions $(a, \text{suc}(a), j+1, \dots, |N|)$ according to (17). If t_a^{start} for appointments in positions $(a, \text{suc}(a), j+1, \dots, |N|)$ satisfy (18), the resulting schedule is accepted as an alternative for therapist e . If we move a sequence $(a, \text{suc}(a))$ forward between appointments in positions k and $k+1$, start times t_a^{start} of appointments in position before the initial position of appointment a remains the same




Fig. 3 Backward and forward relocating a sequence of appointments on the same therapist (neighborhood 1)

while appointments in initial positions $(a, \text{suc}(a), \dots, k, \dots, |N|)$ are updated according to (17) while respecting (18). For searching, we select a therapist e arbitrarily. We start with individual appointments for each therapist, i.e., sequence length one. We then iterate over the backward and forward positions to determine alternative schedules for the therapist. After sequence length one was checked for all therapists, we set the sequence length to two and three, respectively, i.e., each therapist's schedule is checked three times to check all three different sequence lengths.

After we determined alternative schedules $S^{\text{alternative}}$, we apply improvement checking. Here, we check if we can insert an additional appointment $a \in \{\text{App}^{\text{unscheduled}}(b) | F(a) = \emptyset \vee a' \in \text{Sol}_{p_1}(b) \forall a' \in F(a)\}$ into one of the schedules in $S^{\text{alternative}}$. The requirement for insertion is that no potential predecessors block an appointment from insertion. We arbitrarily choose a fitting unscheduled appointment and order qualified therapists in non-decreasing order according to the (preference) violation class of the combination between appointment and therapist. In other words, we try to schedule appointments with therapists first with the most important (preference) violation class available, where a therapist prefers to treat a patient. We go through the list of qualified therapists ordered by preference violation classes for the given appointment, until insertion is possible, or all qualified therapists are unsuccessfully tested. The improvement checking process is repeated for all unscheduled appointments, which can be scheduled in a given priority class. If there is no improvement after checking all unscheduled appointments, we go to the next neighborhood. If an appointment a can be scheduled with a qualified therapist, we update the following appointments in the schedule according to formulas (17) and (18). We need to update $S^{\text{alternative}}$ since there are changes in the schedule $\text{Sol}(e)$ and we need to remove a from $\text{App}^{\text{unscheduled}}(b)$. While it is unnecessary to update $S^{\text{alternative}}$ regarding all therapists, since most schedules remain unchanged, it is not enough to only test new combinations for the changed schedule $\text{Sol}(e)$. Appointments $a \in \text{Sol}(e)$ might have scheduled predecessors $F(a)$ or successors $G(a)$. If an appointment is inserted and $\text{Sol}(e)$ changes, so might the t_a^{start} for $a \in \text{Sol}(e)$ and the restrictions in formulas (17) and (18) for their predecessors $a' \in F(a) \vee G(a)$. So, it is necessary to update all alternative schedules for therapists with $a' \in F(a) \vee G(a)$ if t_a^{start} has changed.




Fig. 4 Relocating a sequence of appointments from one therapist to another therapist (neighborhood 2)

(Relocation) Neighborhood 2: Figure 4 shows an example of the second neighborhood with sequence length two. The neighborhood relocates up to three appointments from one qualified therapist to another. The feasibility checking works similarly to the first neighborhood, i.e., we only relocate an appointment if formulas (17) and (18) hold for the relocated appointments and the following appointments in positions $(a, \dots, n+1, \dots, |N|)$ for the potential new therapist e' . The start times for the remaining appointments in the schedule of the original therapist e remain unchanged.

In the next step, improvement checking considers the (preference) violation classes. Assume, e.g., $(a, \text{succ}(a))$ are successfully relocated from therapist e to e' . The relocation's violation cost is updated according to $U_{(a, \text{succ}(a)), e, e'} = (C_{e', a}^{\text{violation}} - C_{e, a}^{\text{violation}}) + (C_{e', a'}^{\text{violation}} - C_{e, a'}^{\text{violation}})$. Unscheduled appointments $a^+ \in \text{App}^{\text{unscheduled}}(b)$, if inserted in a schedule, contribute to the (preference) violation penalty. Therefore, $U_{a^+, (a, \text{succ}(a)), e, e'} = (C_{e', a}^{\text{violation}} - C_{e, a}^{\text{violation}}) + (C_{e', a'}^{\text{violation}} - C_{e, a'}^{\text{violation}}) + C_{e, a^+}^{\text{violation}}$ determines the (preference) violation cost change if $a^+ \in \text{App}^{\text{unscheduled}}(b)$ is scheduled to therapist e after a sequence $(a, \text{succ}(a))$ is moved from therapist e to e' . All alternative schedules (for all sequence lengths), in combination with the unscheduled appointments of the same appointment type, are sorted according to $U_{a^+, (a, \text{succ}(a)), e, e'}$ in non-decreasing order. We start insertion with the first combination. If unsuccessful, we iterate through the sorted list until insertion is possible. Insertion works similarly to the first neighborhood but updates to $S^{\text{alternative}}$, for successors' and predecessors' therapists are not only necessary for appointments at therapist e but also for appointments at therapist e' .

(Problem specific relocation) Neighborhood 3: The third relocation neighborhood is shown in Fig. 5. In this neighborhood, we specifically deal with potential failure to schedule patient appointments because the patient's predecessor appointments are scheduled too late. For the search procedure, an unscheduled appointment with all predecessors already scheduled is selected arbitrarily, one after another, within a priority class. Feasibility checking relocates all predecessor appointments $a' \in F(a)$ backwards for their scheduled therapist e' to allow scheduling of the remaining successor appointment a . All predecessor appointments are tried earlier in positions $\{0, \dots, j+1\}$, according to formulas (17) and (18). If successful, i.e., the updated $\max\{t_{a'}^{\text{start}} \forall a' \in F(a)\}$ is smaller than in the original schedule, the alternative

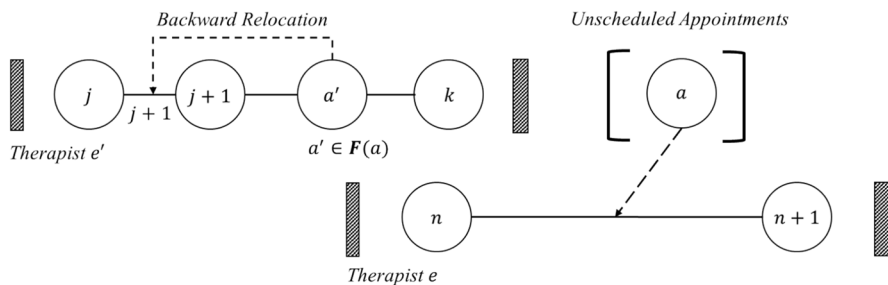


Fig. 5 Relocating predecessors backwards (neighborhood 3)




Fig. 6 Appointment interchange for the same therapist (neighborhood 4)




Fig. 7 Appointment interchange between two therapists (neighborhood 5)

schedules for all involved therapists with appointments $a' \in F(a)$ are added as combination to $\mathcal{S}^{\text{alternative}}$. Improvement checking is like the first neighborhood. We order therapists $e \in E^{\text{available}}(a)$ in non-decreasing order according to the (preference) violation class of the combination between appointment and therapist and try insertion in positions $(n = 1, \dots, |N|)$ of $\text{Sol}(e)$, until successful. If insertion is not possible, we go to the next therapist until successful or no qualified therapist remains for the unscheduled appointment.

We also use three interchange neighborhoods. We interchange the position of two appointments in a therapist's schedule, two appointments between two therapists with equal qualifications, and we exchange a scheduled appointment with two unscheduled appointments.

(Interchange) Neighborhood 4: Fig. 6 shows the first interchange neighborhood.

When checking feasibility, we try to exchange the position of two appointments a and a' . This interchange is tried for all pairs within a schedule $\text{Sol}(e)$. An interchange is feasible if all t_a^{start} of appointments in positions $(a, \dots, a', \dots, |N|)$ can be updated according to (17) without violating (18). If successful, the alternative schedule for therapist e is added to $\mathcal{S}^{\text{alternative}}$ for improvement checking. Improvement checking works similarly to the first relocation neighborhood, where an arbitrarily chosen unscheduled appointment is tried for insertion on the schedule of qualified therapists $e \in E^{\text{available}}(a)$. Therapists are picked in non-decreasing order according to $C_{e,a}^{\text{violation}}$. If insertion was successful, $\mathcal{S}^{\text{alternative}}$ is updated for the schedule where an insertion happened, and we consider potential changes for other therapists due to pathway relationships for appointments in positions $(a, \dots, a', \dots, |N|)$.

(Interchange) Neighborhood 5: Fig. 7 shows Neighborhood 5. Here, appointment a from schedule $\text{Sol}(e)$ takes the place of appointment a' in schedule $\text{Sol}(e')$ and vice

versa. In feasibility checking, appointments in positions $(a, \dots, |N|)$ and $(a', \dots, |N|)$ in the alternative schedules are updated according to (17) and (18). If an interchange is feasible, the pair is added to $S^{\text{alternative}}$. Improvement checking considers preference violations like Neighborhood 2. Interchanging a and a' affects preference violations. The change is calculated as $U_{a,a'} = (C_{e',a}^{\text{violation}} - C_{e,a}^{\text{violation}}) + (C_{e,a'}^{\text{violation}} - C_{e',a'}^{\text{violation}})$. Insertion of an unscheduled appointment also affects the sum of (preference) violation costs. Violation costs $C_{e,a^+}^{\text{violation}}$ for $a^+ \in \text{App}^{\text{unscheduled}}(b)$ must be considered for alternative combination's violation change if therapists $e, e' \in E^{\text{available}}(a)$ are qualified for the appointment. Alternative schedules are evaluated with $U_{a^+,e,a,a'} = (C_{e',a}^{\text{violation}} - C_{e,a}^{\text{violation}}) + (C_{e,a'}^{\text{violation}} - C_{e',a'}^{\text{violation}}) + C_{e,a^+}^{\text{violation}}$ if the unscheduled appointment was scheduled in the schedule of therapist e and $U_{a,a'}$ becomes $U_{a^+,e',a,a'}$ when adding a^+ to the second therapist of the interchange. We sort the combinations of interchanges together with unscheduled appointments in non-decreasing order according to $U_{a^+,e,a,a'}$. Starting with the first item, we try to insert a^+ in positions $(a', \dots, |N|)$ of therapist e , where a' was added instead of a , until successful. If we cannot add a^+ , we go to the next item which might constitute two new therapists and a different unscheduled appointment. If an insertion is successful, new $\text{Sol}(e)$ and $\text{Sol}(e')$ are accepted. Then, $S^{\text{alternative}}$ is updated with respect to therapists e and e' as well as for all therapists with predecessors or successors of appointments in $\text{Sol}(e)$ or $\text{Sol}(e')$, where t_a^{start} changed.

(Interchange) Neighborhood 6: Neighborhood 6 substitutes a scheduled appointment with two unscheduled appointments within the same priority class b . Feasibility checking is easily achieved. Every scheduled appointment a without scheduled successors, i.e., $G(a) = \emptyset$ can be replaced by unscheduled appointments since start times of appointments with later positions remain unchanged. Improvement checking considers preference violation. A therapist e is chosen arbitrarily. Appointments $a \in \text{Sol}(e) \cap \text{App}(b) | G(a) = \emptyset$ are tested for replacement, starting with an appointment with the least important $C_{e,a}^{\text{violation}}$. Unscheduled appointments are sorted in non-decreasing order according to $C_{e,a}^{\text{violation}}$ and the first two unscheduled appointments are tried for insertion instead of appointment a . The insertions can happen at any position in the schedule, i.e., the position of appointment a is not necessarily the insertion point. The two unscheduled appointments are scheduled independently from each other regarding positions. If the unscheduled appointments a' and a'' are




Fig. 8 Exchange scheduled appointment with unscheduled appointments (neighborhood 6)

inserted, they are removed from $\mathbf{App}^{\text{unscheduled}}(b)$ and added to $\mathbf{Sol}(e)$ while appointment a is added to $\mathbf{App}^{\text{unscheduled}}(b)$. If an insertion is not possible, we go through the list of unscheduled appointments where a'' and the next appointment are tested. We continue to go through $a \in \mathbf{Sol}(e) \cap \mathbf{App}(b) | G(a) = \emptyset$ sorted by $C_{e,a}^{\text{violation}}$ in non-increasing order (Fig. 8).

5 Experimental study

In this section, we show the performance of the GRASP compared with a (standard) solver, and we show the effectiveness of the neighborhoods. We use the workforce situation of a German university hospital with roughly 1700 beds. Patient and appointment data is generated to represent a typical weekday within the maximum care environment that inspired the project. The GRASP was built using Python 3.8, and experiments were run on a virtual machine with an Intel Xeon Gold 5218 CPU @ 2.30 GHz processor and 16 GB RAM. Gurobi 9.0.2 (Gurobi Optimization 2022), in combination with Python, was used as the solver for the mathematical model. In Sect. 5.1, we describe the data for the experiments. Sections 5.2 and 5.3 discuss the results and managerial insights.

5.1 Data

The therapy department works in shifts with different working times with and without lunch breaks. The department is on duty each day. On weekends only emergency staffing is available. The therapist's unavailability due to compensation days for weekend duty and vacations influences the planning on weekdays. Hence, we consider about 80% of the therapist workforce to be available. Among the 50 considered therapists, there is one sports therapist, four masseurs, eight occupational therapists, and 37 physiotherapists. 20 therapists work shifts without lunch breaks. The rest has a fixed lunch break starting at noon. Depending on the shift, work starts between 07:00 AM and 09:30 AM and ends between 11:20 AM and 04:40 PM. A day is split into 10-min time periods, with the earliest shift start $t = 1$ and the latest shift end $t = 58$. For demand data, we randomly construct a data set of at most 500 patients, representative of daily demand. The data are based on hospital information concerning required appointment types and treatment durations given to us. Demand data plausibility was discussed with the hospital department. The hospital has more than 1700 beds, but not each patient is treated by a therapist daily. Some clinics do not require therapist work, and therapy might only be required a certain number of days after surgery. According to their respective therapy pathways, some of the 500 patients require more than one therapy appointment a day. We consider 138 patients with more than one appointment in our data set. The constructed data set includes 655 appointments divided into five priority classes: 35 are in the most important priority class 5, 61 are in class 4, 117 are in class 3, 268 are in class 2, and 174 are in class 1. Note that outpatients and emergency cases have the highest priority. While outpatients boost the financial result, their treatment area is rather small. Only

16 outpatient appointments with a fixed treatment time are part of the data set. For inpatients, time windows vary from a very short window for a patient who needs an appointment before leaving the hospital in the morning to a very large time window for a patient covering the whole day since no physician visits are planned. Treatment time varies between 10 min (i.e., 1 period) for a check-up to 50 min (i.e., 5 periods) for extensive treatments. CoC relation exists for 218 out of the 655 appointments. The importance of CoC outweighs the preference violation of therapists. We consider five different (preference) violation classes where class 5 is the most important and class 1 is the least important. Therapists travel to patients' ward rooms for treatment which usually requires transition times of one period (i.e., 10 min).

The described data set is used to create several smaller instances to test the (standard) solver and the GRASP. As shown in Sect. 2, the size of our problem is larger than most other daily physical therapy scheduling problems. Since only one sports therapist and four masseurs work on a day, we ensure that all qualification types are represented in each instance. The same is true for outpatients and the different priority and (preference) violation classes.

5.2 Results

This section discusses the performance of the GRASP phase 1 and the individual neighborhoods (phase 2). Section 5.2.2 discusses the time requirements of the different neighborhoods and introduces a promising neighborhood search order. Afterward, we discuss fitting hyperparameter settings and compare the GRASP with the solver concerning solution quality and time requirements. Finally, we test the GRASP on larger instances and determine that the introduced neighborhoods are effective and frequently used.

5.2.1 Performance of the GRASP components

To compare the GRASP with (standard) solver solutions, we determine whether all six neighborhoods are effective. For determining the effectiveness, we fix the result achieved in phase 1, i.e., instead of choosing a candidate randomly, we choose the best insertion candidate in each iteration. We use a medium-sized instance with 250 patients, with 307 appointments among them and 25 therapists. Of the 307 appointments, 16 are in priority class 5, 26 in class 4, 50 in class 3, 123 in class 2, and 92 in class 1. After the first phase, we run each of the six neighborhoods individually. Table 2 shows the results. US indicates the number of unscheduled appointments. Prio5 to Prio1 break down the unscheduled appointments according to their priority class, with Prio5 being the most important class. *S* shows the total number of scheduled appointments with Pref5 to Pref1 showing the resulting number of appointments within the (preference) violation classes, with Pref1 being the least important (preference) violation class.

For priority classes, it is more beneficial to reduce the number of unscheduled appointments in more important classes by even one before considering lower-priority classes. Table 2 shows the ability of phase 1 (row labeled *Phase1*) to generate

Table 2 Comparison of the effectiveness of individual neighborhoods

Number of unscheduled appointments		Number of scheduled appointments									
<i>U/S</i>	Prio5	Prio4	Prio3	Prio2	Prio1	<i>S</i>	Pref5	Pref4	Pref3	Pref2	Pref1
Phase1	0	0	1	13	52	241	189	27	11	11	3
N1	0	0	1	12	51	243	190	27	11	12	3
N2	0	0	1	12	49	245	190	28	10	12	5
N3	0	0	1	12	51	243	189	28	11	11	4
N4	0	0	1	12	50	244	190	28	11	12	3
N5	0	0	1	11	46	249	192	31	10	13	3
N6	0	0	1	13	51	242	188	27	13	11	3
Legend:	Phase I—result without local search					N3—neighborhood 3					
						N2—neighborhood 2					
						N5—neighborhood 5					
						N4—neighborhood 4					
						N1—neighborhood 1					
						N6—neighborhood 6					

good initial solutions since no differences in the three most important priority classes Prio5 to Prio3 can be observed compared to the neighborhoods. The effectiveness of the six neighborhoods (rows labeled N1 to N6) is shown as each neighborhood schedules additional appointments compared to Phase1. In most cases, additional appointments can be scheduled for Prio2 and Prio1. In N6, only one additional appointment in Prio1 is scheduled. N5 can schedule the highest number of appointments by exchanging appointments between two therapists, with the most improvement in Prio2 and Prio1, when used standalone for the given instance. However, while the results show, that all neighborhoods are generally effective, Table 2 does not suggest, that N5 is the most effective neighborhood for all potential instances, or when used in combination with other neighborhoods. Individual results of the neighborhoods depend on the schedules provided by phase 1 on a given instance.

5.2.2 Neighborhood time requirements and neighborhood search order

Table 3 shows the processing time of the neighborhoods for the same data set. The neighborhoods require more time for less important priority classes, where more appointments are already scheduled. Interchanging and relocating within one therapist's schedule, i.e., N1 and N4, requires less time than operations considering two therapists, i.e., N2 and N5. N1 takes longer than N4 since N1 searches sequence lengths from one appointment to three appointments for improvements, while N4 only includes one type of exchange operation. N6 and N3 require the least amount of time since they combine feasibility and improvement checking and consider therapists one at a time for individual unscheduled appointments. The effectiveness of N5 comes at the cost of a long processing time compared to the other neighborhoods. This can be explained by the necessity to update starting times as well as affected predecessor and successor relations in a therapy pathway for two therapists.

In the next step, we carried out experiments for the order in which we use the neighborhoods. The detailed experiments can be seen in Appendix E. In general, we could observe that every tested combination of the six neighborhoods leads to better results than using individual neighborhoods for the medium-sized instance. Testing the most promising orders on other medium-sized instances, we concluded that the order N4–N1–N5–N2–N6–N3 seems to be working best, considering solution quality and processing time.

Table 3 Processing time of the individual neighborhoods in seconds

	N1	N2	N3	N4	N5	N6
Prio5	0.004	0.004	0.007	0.003	0.020	0.007
Prio4	0.003	0.003	0.007	0.003	0.020	0.004
Prio3	0.016	0.030	0.004	0.005	0.020	0.005
Prio2	0.230	3.160	0.009	0.080	4.170	0.012
Prio1	3.240	6.320	0.047	0.980	73.990	0.051
Total	3.493	9.527	0.075	1.071	78.220	0.079

5.2.3 Parameter setting

To determine the best time tradeoff between phase 1 (I_2) and phase 2 (I_1) iterations, we started the GRASP five times with $I_1 = 1$ and $I_2 = 1$ for each of five different instance sizes, ranging from 100 patients and 10 therapists up to the largest, most practically relevant instance of 500 patients and 50 therapists. We determined that phase 2 requires far more time for all instance sizes across all five GRASP starts per instance size. It was also observed that across the five runs, the coefficient of variation (CoV) of processing times for phase 1 lies between eleven percent and 22 percent depending on the instance. For phase 2, we experienced higher CoV between 13 and 49 percent, which can be explained by the different neighborhoods leading to additional insertions, depending on the solution candidate achieved in phase 1. Due to these results, we run phase 2 only on the best solution of several phase 1 iterations. Nevertheless, I_2 should not be set too high to guarantee a diverse set of phase 1 solutions. Too many iterations might result in the same local optimum for each GRASP repetition. For I_1 , we focus on the time limit of two hours for our problem. Preliminary tests showed that 10 phase 1 iterations ($I_2 = 10$) combined with 15 phase 2 iterations ($I_1 = 15$) delivers consistently good results across all problem sizes while satisfying the two-hour time limit for the realistic problem size of 500 patients and 50 therapists. Hence, we use this setting for comparing the GRASP with the standard solver.

5.2.4 Comparison between GRASP and commercial solver

We evaluate five different final GRASP results for each problem size.

Table 4 shows the best result of the GRASP (GRBE) as well as the worst result (GRWO). SOF shows the (standard) solver, i.e., Gurobi's, results considering both objectives (4.1) and (4.2). The first column shows the number of patients $|P|$, appointments $|A|$ and therapists $|E|$ in the different instances. We compare results between the GRASP and the solver for up to 100 patients and 10 therapists. Larger problem sizes become intractable with the commercial solver due to memory restrictions. As described in Sect. 3, we solve the problem in hierarchical order, where minimizing the total priority of missed appointments dominates the total preference violation cost of scheduled appointments, which also includes CoC consideration.

The GRASP solutions manage to schedule all appointments within the three most important priority classes, like the standard software. Minimizing unscheduled appointments according to priority, GRBE does not achieve the optimal result in any instance size, but high-quality solutions are achieved compared to SOF. For the 61 appointment and the 122 appointment instances, for example, GRBE schedules only one appointment less than SOF in the least important priority class. Table 4 also shows that GRWO (the worst GRASP run) performs slightly worse than GRBE. However, high-quality solutions are still achieved concerning the main objective of fulfilling appointments according to their priority class with *Prio5* being the most important class. Concerning the second objective of preference violations, the GRASP solutions are comparable to SOF for all instances. GRBE and GRWO

Table 4 Comparison of results between standard solver and GRASP

$ P / A / E $	Method	Number of unscheduled appointments					Number of scheduled appointments						
		US	Prio5	Prio4	Prio3	Prio2	Prio1	S	Pref5	Pref4	Pref3	Pref2	Pref1
50/61/5	SOF	15	0	0	0	2	13	46	21	8	10	6	1
	GRWO	17	0	0	0	3	14	44	17	10	8	8	1
	GRBE	16	0	0	0	2	14	45	18	9	9	7	2
75/93/5	SOF	38	0	0	0	10	28	55	23	13	10	6	3
	GRWO	41	0	0	0	13	28	52	21	12	8	8	3
	GRBE	39	0	0	0	12	27	54	25	11	8	7	3
75/93/7	SOF	27	0	0	0	0	27	66	35	15	8	6	2
	GRWO	29	0	0	0	2	2	64	35	13	8	6	2
	GRBE	29	0	0	0	1	28	64	34	10	8	7	5
80/98/10	SOF	3	0	0	0	0	3	95	46	34	8	5	2
	GRWO	6	0	0	0	3	3	92	63	14	9	4	2
	GRBE	5	0	0	0	1	4	93	65	14	8	4	2
85/103/10	SOF	5	0	0	0	1	4	98	50	27	12	6	3
	GRWO	8	0	0	0	4	4	95	61	16	10	5	3
	GRBE	7	0	0	0	2	5	96	62	15	7	9	3
100/122/10	SOF	20	0	0	0	4	16	102	63	23	8	5	3
	GRWO	23	0	0	0	7	16	99	65	16	9	6	3
	GRBE	21	0	0	0	4	17	101	69	15	8	5	4

Table 5 Comparison of processing times between commercial solver and GRASP in seconds

$ P / E $	50/5	75/5	75/7	80/10	85/10	100/10	150/15	200/20	300/30	400/40	500/50
Solver	130.36	212.91	696.45	3037.51	920,671.14	63,677.39	—	—	—	—	—
GRASP mean	5.93	12.49	26.36	10.10	23.35	121.13	298.80	746.48	2866.21	4274.87	5750.96
GRASP min	5.62	11.56	24.06	6.09	20.27	102.15	250.76	447.43	1968.54	3699.96	4352.54
GRASP max	6.58	13.63	28.54	13.57	25.63	149.98	376.13	1124.69	5636.9	5312.01	7051.78




Fig. 9 Comparison of schedules between GRASP (schedule G) and Standard Solver (schedule S) for 100 Patients and 10 Therapists

provide good solutions, especially considering the differences in processing time (see Table 5).

The GRASP also provides an equally tight schedule compared to the (standard) solver results. Figure 9 compares the best GRASP solution *G* to the solver solution *S* for each of the 10 therapists when considering the 100-patient problem size. The workforce consists of one sports therapist *S1*, one occupational therapist *E1*, two masseurs *M1* and *M2*, and six physiotherapists *P1* to *P6*. Therapists are sorted according to their ID within the implementations.

The figure shows different shifts with fluctuating start and end times and some therapists without a lunch break. The schedules of the masseurs and the sports therapist are not tightly packed. Unlike the (standard) solver solution *S*, the appointments are scheduled as early as possible according to the GRASP logic, leading to no appointments for *M2* after lunch. There is little difference between preference violations. *P5* and *P6* mainly work on appointments outside of the most preferred violation class 1 for both solution methods. For occupational therapists and physiotherapists, the GRASP manages to generate plans with equally little slack compared to the (standard) solver in only a fraction of the time.

Table 5 shows processing times for the same instances shown in Table 4. Due to intractability caused by memory restrictions, processing times for the standard solver (Gurobi) are shown for up to 100 patients $|P|$ and 10 therapists $|E|$. For larger instances, no feasible solution could be found until intractability issues occurred. For GRASP, the minimum and maximum processing times are shown for problem sizes up to the largest problem of 500 patients and 50 therapists. For the small instance of 85 patients and 10 therapists, the standard solver achieves the optimal result only after several days. For 100 patients and 10 therapists, the standard solver needed more than 17 h to optimality, i.e., to achieve the results of Table 4. After two hours, the instance with 85 patients did not find a practical solution, and for 100 patients and 10 therapists, surprisingly the optimality gap was slightly above 4 percent after two hours. The GRASP provides solutions much more quickly. The 100-patient size problem is solved after 121 s on average, with 150 s required for the most time-consuming GRASP run. The GRASP requires more time for larger problem sizes

Table 6 GRASP results for larger problem sizes

$ P / A / E $	Method	Number of unscheduled appointments					Number of scheduled appointments						
		U/S	Prio5	Prio4	Prio3	Prio2	Prio1	S	Pref5	Pref4	Pref3	Pref2	Pref1
150/181/15	GRWO	24	0	0	1	5	18	157	81	52	11	10	3
	GRBE	23	0	0	1	1	21	158	79	55	12	7	5
200/242/20	GRWO	38	0	0	2	7	29	204	128	58	9	7	2
	GRBE	37	0	0	1	8	28	205	120	64	13	6	2
300/372/30	GRWO	66	0	0	1	21	44	306	194	83	20	9	0
	GRBE	63	0	0	1	19	43	309	196	85	18	7	3
400/515/40	GRWO	98	0	1	1	40	56	417	296	87	26	6	2
	GRBE	97	0	0	1	39	57	418	295	94	18	10	1
500/655/50	GRWO	128	0	1	3	50	74	527	396	107	15	9	0
	GRBE	124	0	0	2	51	71	531	405	101	15	7	3

but stays under the 2-h limit for the largest problem size with 500 patients and 655 appointments.

5.2.5 GRASP results for realistically large instances

Table 6 shows the best and worst GRASP results for five runs of the GRASP.

The GRASP avoids unscheduled appointments in *Prio4* or *Prio5* for problem sizes of up to 300 patients. For 400 and 500 patients, one appointment in *Prio4* remains unscheduled in GRWO, while all *Prio4* appointments were scheduled in *GRBE*. The vast majority of *Prio4* appointments are scheduled. The total number of scheduled appointments does not vary much between the best and the worst achieved result per problem size, meaning even the worst achieved result can provide good quality solutions in the required timeframe. On the right, we show the number of scheduled appointments within the five preference (violation) classes. It can be seen for all instances that most scheduled appointments are in the class *Pref5* which therapists prefer the most. From a patient perspective, not every patient might have a therapist with a high preference for the patient's appointment available, i.e., for some appointments, e.g., *Pref3* might be the best scheduling possibility. Therefore, we additionally examined the relative preference adherence which determines the percentage of appointments assigned to the best (second best, etc.) possible preference (violation) class. Assuming two appointments can be at best scheduled in *Pref3* and the appointments are indeed scheduled in *Pref3*, the percentage is 100 percent. If one of the two cannot be scheduled in the highest possible class, then the percentage is 50 percent. Depending on the instance, between 71.18 percent and 76.39 percent of appointments are scheduled within their best possible preference (violation) class. The percentage is between 93.68 percent and 95.41, if we consider the best or the second-best preference (violation) class. We also witnessed hardly any standard deviation (SD) across the five GRASP runs for each instance. Detailed results are shown in Table 15 in Appendix H.

Table 7 shows the performance of the GRASP on larger instance sizes with five GRASP runs for each instance, i.e., for the instance with 500 patients and 50 therapists for example, we included 5×50 schedules = 250 schedules in the calculations. Therapists spend about 75 percent of their shifts directly treating patients. We see similar means and medians across all five instances. The slightly higher medians compared to the means are due to similar effects as can be seen in Fig. 9, where a qualification type has fewer potential daily appointments than the other qualification types. SD varies between 2.4 percent and 8.0 percent. Including the necessary travel time between patients, on average, therapists spend 95.3 percent of their shifts with productive tasks for all instances. The percentage is slightly higher when looking at the medians. For the instance with 300 patients and 30 therapists, the median therapist does not have any idle time. The SD varies between 3.1 percent and 11.8 percent between the different instances. Finally, we track the percentage of schedules without idle time, i.e., time not spent walking between patients or treating patients. For the four smaller instances, 45.0% to 51.0% of the schedules do not have idle time. For the largest instance with 500 patients and 50 therapists, 31.6% of the 250

Table 7 Percentages of shifts spent with patients (including travel times) and shifts without idle time

$ P / A / E $		150/18/15	200/24/20	300/37/30	400/51/40	500/65/50
Percentage of shift spent with patient	Median	0.750	0.750	0.750	0.750	0.739
	Mean	0.723	0.743	0.749	0.737	0.727
	SD	0.080	0.036	0.024	0.062	0.064
Percentage of shift spent with patient and traveling between patients	Median	0.978	0.978	1.000	0.978	0.968
	Mean	0.939	0.965	0.976	0.965	0.953
	SD	0.118	0.053	0.031	0.082	0.083
Percentage of schedules without any idle time (across 5 GRASP Runs)	Mean	0.480	0.450	0.510	0.465	0.316
	SD	0.049	0.045	0.107	0.056	0.040

generated schedules, i.e., 79 schedules do not have idle time. This is another indication of the ability of the GRASP to provide tightly packed schedules. Please note, the therapists do not work 100% of their available working time, e.g., travel time is less than assumed and (lunch) breaks are considered.

Using the largest problem size with 500 patients, we varied problem parameters to test the sensitivity of the problem and the GRASP results. First, we removed differences in appointment priority by assigning all appointments to the same priority class. Second, we removed recovery times in the individual therapy pathways of patients for more scheduling flexibility. In neither of the two variations, we saw large differences in the number of scheduled appointments or different scheduling patterns. Results for these tests are described in more detail in Appendix I.

5.2.6 Frequency of neighborhood use

Table 8 indicates that the neighborhoods remain effective in the chosen neighborhood order. Within each GRASP procedure, we chose 15 I_1 iterations. The whole GRASP procedure was run five independent times with $I_1 = 15$ (and $I_2 = 5$). This multiplies to 75 potential times a neighborhood can be used to insert additional appointments across all priority classes. For all instances, N_2 , N_3 , and N_4 were used in most cases. N_5 seldom led to improvements because of the relatively restrictive environment. We need additional space in the therapist's schedule for a scheduled appointment to be replaced by two unscheduled appointments. Predecessor appointment times of unscheduled appointments restrict time slots for insertion, and time restrictions for scheduled appointments must be respected. N_1 varies the most, i.e., N_1 is not used within one GRASP start for some problem sizes while frequently being used in another GRASP start. This behavior is due to running the GRASP on the individual priority classes. While one interim solution restricts relocations within therapist schedules for lower priority classes, another best solution allows successful relocations using N_1 . N_6 is used 25 times over all priority classes for the smallest problem size with the smallest number of therapy pathway constraints. With more therapy pathways to consider for larger problem sizes, it is used most or all 75 possible times, mostly in *Prio2* and *Prio1*. Most

Table 8 Number of phase 2 runs where neighborhoods led to additional appointment insertions

$ P / E $	50/5	75/5	75/7	80/10	85/10	100/10	150/15	200/20	300/30	400/40	500/50
N_1	14	5	32	4	14	15	14	47	36	21	47
N_2	47	51	65	50	52	42	42	70	68	74	50
N_3	57	62	57	34	43	49	71	73	74	70	73
N_4	48	40	70	52	63	50	75	75	75	75	75
N_5	4	7	5	0	0	0	2	0	0	0	1
N_6	25	75	75	57	58	75	37	63	75	75	75

additional insertions are in these lowest-priority classes when the schedules are already tight. Overall, $N2$, $N3$, $N4$, and $N6$ lead to most insertions across all problem sizes, with most insertions in *Prio2* and *Prio1*. While construction phase 1 provides good initial results, the number of scheduled appointments increases by five to ten percent for all problem sizes using local search in phase 2.

5.3 Managerial insights

Without automated scheduling, therapists in our partner hospital are assigned to wards. Within these wards, they walk to rooms without a given schedule and often find patients occupied. If a patient is occupied, therapists return to the department offices to look up the treatment requirements of other untreated patients. This causes a lot of idle time. Therapists are generally unable to fulfill the department's daily appointment objectives, which would only be obtainable with uninterrupted tightly packed visitation schedules. This leads to a vicious circle of low revenue and the inability to hire more therapists as well as threatens rehabilitation progress. GRASP-generated schedules can help to decrease this uncertainty for therapists and management. Figure 9 showed that our modeling approach and the GRASP can provide tightly packed schedules for physical therapists. We obtained similar results for larger problem sizes using the GRASP, as was shown in Table 7. Different generated schedules, with varying time windows (i.e., patient availability), enable the department to quantify possible adaptations in physician visitation patterns, which can further decrease uncertainty. According to department management, strong preferences for treatments exist among therapists. These preferences can hardly be considered in today's manual scheduling approach. For the smaller instance sizes, where optimal results can be compared, Table 9 compares the model without preferences (*NoPref*), i.e., the

Table 9 Preference differences considering the second objective function

$ P / A / E $	Result	Number of scheduled appointments within preference classes					
		S	Pref5	Pref4	Pref3	Pref2	Pref1
50/61/5	NoPref	46	17	6	9	10	4
	WithPref		21	8	10	6	1
75/93/5	NoPref	55	19	8	8	14	6
	WithPref		23	13	10	6	3
75/93/7	NoPref	66	27	7	9	13	10
	WithPref		35	15	8	6	2
80/98/10	NoPref	95	34	13	14	17	17
	WithPref		46	34	8	5	2
85/103/10	NoPref	98	23	17	26	13	19
	WithPref		50	27	12	6	3
100/122/10	NoPref	102	34	14	16	21	17
	WithPref		63	23	8	5	3

objective function (4.2) is removed, with the model including preferences (*With-Pref*). No differences in the number of scheduled appointments (S) for small problem instances exist. The reason is the dominance of objective (4.1) over objective (4.2). For all problem sizes, we see that objective (4.2) substantially improves the schedules with respect to preference considerations. Remember, *Pref5* is the most desired preference (violation) class and *Pref1* is the least desired.

Another factor for therapist satisfaction is CoC. We examine the effects of pooling, where we consider CoC only in the objective function (4.2), by including a set of constraints into the IP problem to compare optimal results. The constraints enforce that a patient can only be treated by the existing CoC therapist. If an appointment cannot be scheduled with the CoC therapist, it remains unscheduled. For the experiments, we use six smaller instances. We guarantee that all potential appointments have a continuity of care relationship, i.e., there is no appointment without a CoC therapist. To have more CoC variety, we also adapt the staff composition slightly by replacing a physiotherapist with an additional occupational therapist for staff sizes of seven and ten therapists.

Table 10 shows the difference in scheduled appointments when enforcing CoC using constraints (first row) versus using a pooling approach, where CoC is only considered as part of our objective function (4.2) (second row). The numbers in brackets show how many appointments are scheduled with their CoC therapist. For the smaller instances with 50 or 75 patients and five therapists, all 4 qualification types are considered, i.e., only physical therapy appointments can be scheduled with another therapist than the CoC therapist. This leads to only small differences between a pooling approach and using CoC constraints. For the instances with seven or ten therapists, relaxing CoC constraint and using a pooling approach leads to a bigger effect. There are more therapist alternatives available for the different qualification types. Here, for the instance with 100 patients, i.e., 122 appointments, and 10 therapists, 77 appointments are scheduled in the optimal solution when enforcing CoC using the additional set of constraints. 97 appointments can be scheduled using a pooling approach. We conclude that pooling leads to more scheduled appointments, i.e., it is better for satisfying our objective (4.1). However, it comes at the cost of potentially not being able to satisfy all CoC relationships. Note, only 56 of the 97 appointments were scheduled with the CoC therapist, while all 77 scheduled appointments are scheduled with the respective CoC therapist when enforcing CoC using constraints. We also see that not all 122 appointments with CoC relationships are scheduled due to capacity restrictions.

Table 10 Difference in scheduled appointments—enforcing CoC versus considering it in objective function 4.2

P / A / E	Number of scheduled appointments (with CoC therapist)					
	50/61/5	75/93/5	75/93/7	80/98/10	85/103/10	100/122/10
Scheduled with CoC constraints	45 (45)	55 (55)	55 (55)	54 (54)	59 (59)	77 (77)
Scheduled without CoC constraints	46 (41)	55 (55)	57 (45)	91 (48)	93 (53)	97 (56)

Table 11 Continuity of care adherence

	Continuity of care adherence for the real-sized instance with 500 patients and 50 therapists				
GRASP run	1	2	3	4	5
Scheduled with CoC relation	179	182	181	181	182
Continuity of care adherence	48.04%	46.15%	38.87%	42.54%	43.41%

Similar observations can be made for the most realistic-sized instance. Table 7 in Sect. 5.2.5 showed tightly packed schedules for the real-world-sized problem with 500 patients and 50 therapists. This instance has 218 appointments with CoC relationships.

Table 11 shows the results of five GRASP runs with the real-world-sized problem. In the five runs, scheduled appointments with an existing CoC relation vary between 179 and 182 out of 218. This shows that not all appointments with an existing CoC relationship can even be scheduled, and some remain unscheduled when primarily focusing on scheduling appointments accounting for their priority. Within the scheduled CoC appointments, between 38.87% and 48.04% of appointments with an existing CoC relationship were scheduled with their respective therapist. Following our logic for preference violation classes, an appointment is scheduled with the corresponding CoC therapist, if possible, in tight schedules. Scheduling the appointment with another therapist leads to a higher total penalty. However, preference penalties are much smaller than penalties if fewer appointments would be scheduled to increase CoC.

The developed GRASP provides tightly packed schedules for therapists and includes therapist preference to a large extend. It can do so under the tight time restrictions given by the problem characteristics. We assumed the GRASP is run on the department manager's personal computer with limited computing power. Nevertheless, GRASP procedures can be implemented in parallel (Feo et al. 1994). With more computing power, managers can increase I_1 iterations and reduce the processing time if desired. On a single computer, the manager can still choose input parameters I_1 and I_2 depending on the requirements. With several GRASP starts, we showed variance in results between the best and worst achieved solutions. The manager can balance the trade-off between solution quality and runtime by choosing the parameters appropriately. More iterations lead to more stable solutions while requiring more runtime.

6 Summary and outlook

We describe the challenging workforce situation of physical therapy departments worldwide and the resulting tensions between fulfilling appointments according to their importance and maximizing therapist satisfaction to fight the workforce shortages. We derive a complex operational planning problem that has not yet been considered in the literature. We propose a lexicographical IP model to minimize unscheduled patient appointments, according to their importance for a hospital, while minimizing therapist preference violations. Finding optimal solutions for small problem sizes in an acceptable time is impossible, and the problem is intractable for large instances. Hence, we develop an innovative GRASP which can deliver high-quality solutions while satisfying the time requirements. In experiments, we show the effect of considering therapist preferences on solutions and the effectiveness of the chosen neighborhoods for the problem. We show that the GRASP delivers tight schedules compared to the IP model while requiring a fraction of the processing time.

Based on this contribution, future research possibilities can be derived. We are the first paper to focus on workforce shortages and therapist satisfaction in rehabilitation therapy, but we only consider acute care. The shortage is also prevalent in rehabilitation hospitals, where usually all patients can be treated. Here are opportunities for different objective functions, focusing on therapist preferences and fairness aspects between therapists. We deal with hierarchical appointment priority. The situation might be more fluid in other hospitals requiring adaptations to the solution procedure or different procedures. Comparable to Gartner et al. (2018) and Jungwirth et al. (2021), other hospitals might require designated rooms or instruments for treatment. Stochasticity might be introduced for parameters. Examples are time windows, treatment durations, or no-shows. One might assume therapist-specific treatment times, e.g., depending on the experience of a therapist. Group therapies, where a therapist treats several patients simultaneously or one patient requires several therapists (synchronization), might be considered together with preferences. Finally, practical time and computational capacity restrictions motivated our research and solution procedure. Future research might develop theoretically motivated solution procedures or examine effects of daily scheduling for long-term therapist satisfaction or patient rehabilitation.

Appendix A: High-level architecture of the GRASP

```

1  for  $b = b^{max}$  to  $b = 1$  begin
2      for  $i = 1$  to  $l_1$  begin
3          for  $j = 1$  to  $l_2$  begin
4              Phase 1:
5               $Sol_{p1}(b) \leftarrow Sol^*(b + 1)$ 
6              Insert appointments from  $App(b)$  into  $Sol_{p1}(b)$ 
7              if  $Sol_{p1}(b)$  better than  $Sol(b)$  then
8                   $Sol(b) \leftarrow Sol_{p1}(b)$ 
9              end if
10             end for
11             Phase 2:
12             Local search on  $Sol(b)$  to try additional insertions
13             Evaluate and store improvements in  $Sol(b)$ 
14             if  $Sol(b)$  better than  $Sol^*(b)$  then
15                  $Sol^*(b) \leftarrow Sol(b)$ 
16             end if
17             end for
18             if  $b = 1$  then
19                  $Sol^* \leftarrow Sol^*(b)$ 
20             end if
21         end for

```

This pseudo-code shows the high-level architecture of the GRASP as described in Sect. 4.

Appendix B: Initial scheduling of outpatient appointments in the construction phase of the GRASP

```

1  while  $App^{fixed} \neq \emptyset$  begin
2      select  $a^{fixed}$  from  $App^{fixed}$  at random
3      for  $r = r^{max}$  to  $r = 1$  begin
4          sort  $E(a^{fixed}, r)$  in non-decreasing order according to the number of appointments  $App \setminus a^{fixed}$  with time window overlap with  $a^{fixed}$ 
5          for  $e$  in  $E(a^{fixed}, r)$  begin
6              try insertion of  $a^{fixed}$  in  $Sol(e) \subset Sol_{p1}(b^{max})$ 
7              if  $a^{fixed} \in Sol_{p1}(b^{max})$  then
8                   $App^{fixed} \leftarrow App^{fixed} \setminus a^{fixed}$ 
9                  break
10             end if
11             elif  $r = 1$  and  $e = |E(a^{fixed}, r = 1)|$  then
12                  $App^{fixed} \leftarrow App^{fixed} \setminus a^{fixed}$ 
13                  $a^{fixed} \in App^{unscheduled}(b^{max})$ 
14                 break
15             end elif
16             end for
17             if  $a^{fixed} \notin App^{fixed}$  then
18                 break
19             end if
20         end for
21     end while

```

This pseudo-code shows the procedure of constructing an initial solution for outpatient appointments with fixed treatment slots as described in Sect. 4.2.

Appendix C: Construction phase of the GRASP

```

0  Initialization:
    $A^{possible} = \{App(b) \setminus App^{fixed} | F(a) = \emptyset \vee a' \in Sol_{p1}(b) \forall a' \in F(a)\}$ 
    $E^{available}(a) = E(a)$ 
1  if  $b = b^{max}$  begin
2      process Algorithm 2
3  end if
4  while  $A^{possible} \neq \emptyset$  begin
5      for  $r$  in  $R$  begin
6           $Y \leftarrow 0$ 
7          while  $Y = 0$  begin
8               $C(r) = \{(e, a) | \forall a \in A^{possible}, e \in E(a, r) \cap E^{available}(a)\}$ 
9              Pick  $(e, a)$  from  $C(r)$  at random
10             for  $n$  in  $\{1, \dots, |N|\}$  begin
11                 attempt insertion of  $a$  in  $Sol(e) \subset Sol_{p1}(b)$ 
12                 if  $a \in Sol_{p1}(b)$  begin
13                      $A^{possible} \leftarrow \{App(b) \setminus App^{fixed} | F(a') = \emptyset \vee a'' \in Sol_{p1}(b) \forall a'' \in F(a'), a' \in App\}$ 
14                      $Y \leftarrow 1$ 
15                     break
16                 end if
17                 elif  $n = |N|$  begin
18                      $E^{available}(a) \leftarrow E^{available}(a) \setminus e$ 
19                     if  $E^{available}(a) = \emptyset$  begin
20                          $A^{possible} \leftarrow A^{possible} \setminus a$ 
21                     end if
22                 end elif
23             end for
24             if  $C(r) = \emptyset$  begin
25                 break
26             end if
27         end while
28         if  $Y = 1$  begin
29             break
30         end if
31     end for
32 end while
33  $App^{unscheduled}(b) \leftarrow \{a \in App(b) | a \notin Sol_{p1}(b)\}$ 

```

This pseudo-code illustrates the complete construction procedure starting with outpatients and followed by inpatients for one priority class b , as described in Sect. 4.2.

Appendix D: High-level architecture of the local search

```

1  for neighborhood in neighborhoods begin
2      Inserted  $\leftarrow$  True
3      while Inserted
4          update  $S^{alternative}$ 
5          Improvement checking
6          if no improvement
7              Inserted  $\leftarrow$  False
8          end if
9      end while
10 end for

```

This pseudo-code shows the high-level architecture of the local search as described in Sect. 4.3.

Appendix E: Determination of the neighborhood order

For neighborhood order tests, we used the same solutions after the construction phase 1 for the test across all neighborhood orders to be able to only compare differences in neighborhood orders. Table 12 tests different orders of neighborhoods for the medium-sized instance with 250 patients, comprising of 307 appointments, and 25 therapists.

N_6 is the least effective of the neighborhoods, therefore, we do not test it in early spots. N_3 is tested in positions six and as a fallback in position one. N_3 moves predecessors of an appointment which requires many already scheduled predecessors. Additionally, moving several appointments on different therapists might prevent insertions by other neighborhoods, therefore it likely is most successful in position six. For the other four neighborhoods processing time influences the decision which orders are tested. The combination $N_3-N_1-N_2-N_4-N_5-N_6$ performs the worst among all orders, we therefore do not consider it as a possibility anymore. Combinations $N_2-N_5-N_1-N_4-N_6-N_3$ and $N_5-N_2-N_4-N_1-N_6-N_3$ require more processing time since schedules for two therapists must be reevaluated when a new insertion is made. Therefore, we do not consider combinations with interchanges or relocation between two therapists in first positions anymore. For the data instance, $N_4-N_1-N_5-N_2-N_6-N_3$ provides the best result.

Table 13 shows that the combination $N_4-N_1-N_5-N_2-N_6-N_3$ also outperforms the other neighborhood orders for other medium sized data instances, with a smaller objective function value, while it performs worse once. For the medium sized instances, we chose six problem sizes from 175 patients and 15 therapists up to 300 patients and 30 therapists. Due to the similar solution quality, combined with the most competitive runtimes across all tested instances, the combination $N_4-N_1-N_5-N_2-N_6-N_3$ is chosen for the experiments.

Table 12 Results for different neighborhood orders for 250 patients and 25 therapists

Order	Number of unscheduled appointments						Number of scheduled appointments						Phase 2 time in s
	U/S	Prio5	Prio4	Prio3	Prio2	Prio1	S	Pref5	Pref4	Pref3	Pref2	Pref1	
N1-N2-N4-N5-N6-N3	57	0	0	1	10	46	250	190	32	10	13	5	62.44
N1-N4-N2-N5-N6-N3	57	0	0	1	10	46	250	190	32	10	13	5	61.22
N4-N1-N2-N5-N6-N3	56	0	0	1	10	45	251	188	32	12	13	6	69.98
N4-N1-N5-N2-N6-N3	58	0	0	1	9	48	249	193	30	11	12	3	59.19
N2-N5-N1-N4-N6-N3	56	0	0	1	10	45	251	187	34	12	13	5	85.08
N5-N2-N4-N1-N6-N3	57	0	0	1	10	46	250	189	33	12	13	3	85.39
N3-N1-N2-N4-N5-N6	58	0	0	1	10	47	249	190	31	12	12	4	59.05

Table 13 Objective function values for different neighborhood orders across different instances

Order	175/15	200/20	200/25	250/25	300/25	300/30
N1–N2–N4–N5–N6–N3	274,968	666,801	608,063	3,407,529	14,510,851	14,233,649
N1–N4–N2–N5–N6–N3	274,968	666,801	608,063	3,407,529	14,510,851	14,233,649
N4–N1–N2–N5–N6–N3	274,968	666,801	608,147	3,407,254	14,510,851	14,233,647
N4–N1–N5–N2–N6–N3	274,968	666,801	608,063	3,382,743	14,511,569	14,232,929

Appendix F: Used notation for the IP model and the GRASP

<i>General sets, indices and parameter</i>		$z_{e,p,a,t}$	1 if a therapist e is occupied with appointment $a \in A(p)$ of patient p in period t , 0 otherwise
$e \in E$	Set of therapist employees	Input Sets and Indices Specific to the GRASP	
$a \in A$	Set of appointment types		
$p, p' \in P$	Set of patients	$\mathbf{App}(b)$	Subset of all daily appointments within a priority violation class b
$a, a' \in A(p)$	Set of appointments of patient p	$a' \in F(a)$	Subset of predecessor appointments for $a \in \mathbf{App}$ within a patient's therapy pathway
		$a' \in G(a)$	Subset of successor appointments for $a \in \mathbf{App}$ within a patient's therapy pathway
$b, b' \in B$	Set of hierarchical priority violation classes	$a^{\text{fixed}} \in \mathbf{App}^{\text{fixed}}$	Subset of fixed time outpatient appointments, with $\mathbf{App}^{\text{fixed}} \subset \mathbf{App}$
b^{max}	Highest priority violation class	$E(a)$	Subset of therapists $e \in E$ qualified for $a \in \mathbf{App}$
$A(p, b)$	Set of a patient's appointments $a \in A(p)$ within priority violation class b	$\mathbf{Aq}(e)$	Subset of appointments $a \in \mathbf{App}$ a therapist e is qualified to manage
$r, r' \in R$	Set of preference violation classes	$E(a, r)$	Subset of qualified therapists for an appointment where the therapist has a preference r for appointment a ; $E(a, r) \subset E(a)$
$A(p, r)$	Set of appointments $a \in A(p)$ within preference violation class r	Parameters and Variables Specific to the GRASP	
r^{max}	Highest preference violation class		
$E_{p,a}$	Therapist e with an existing continuity of care relation with appointment $a \in A(p)$ of patient p	I_1	Number of GRASP runs within a priority violation class
		I_2	Number of runs of the construction phase within a GRASP run

<i>Sets specific to the IP formulation</i>		$\mathbf{Sol}_{p_1}(b)$	Solution after one construction run: including appointments down to priority violation class b
$t, t' \in T$	Set of time periods	$\mathbf{Sol}(b)$	Best solution out of I_2 solutions $\mathbf{Sol}_{p_1}(b)$
$T(e)$	Subset of time periods a therapist is available for appointment scheduling	$\mathbf{Sol}(e)$	Solution schedule of an individual therapist $e \in E$, where $\mathbf{Sol}(e) \subset \mathbf{Sol}(b)$
$T(p, a)$	Subset of time periods for an appointment a of patient p	$\mathbf{App}^{\text{unscheduled}}(b)$	Appointments remaining unscheduled after the construction
$F(p, a)$	Subset of predecessor appointments for $a \in A(p)$ for a patient p	$\mathbf{preApp}^{\text{unscheduled}}(b) = \mathbf{App}(b) \setminus \mathbf{Sol}(b)$	
<i>Parameters specific to the IP formulation</i>		$S^{\text{alternative}}$	Alternative feasible schedules after neighborhood operations are applied to $\mathbf{Sol}(b)$ within a given neighborhood
$W_{p,p'}$	Minimum number of periods between two patients p and p' , i.e., transition time of a therapist (might e.g., include walking distances as well as varying administrative time between patients)	$\mathbf{Sol}^*(b)$	Best achieved solution after I_1 GRASP runs including priority violation class b
$S_{p,a}$	Number of treatment periods for appointment $a \in A(p)$ of patient p	\mathbf{Sol}^*	Final solution
$D_{p,a,a'}$	Minimum number of periods between appointments a and a' of one single patient p (e.g., needed to avoid mental or physical stress to a patient caused by follow-up appointments following too fast)	T_e^{earliest}	First time period within a shift of therapist e
		T_e^{latest}	Last time period within a shift of therapist e
Q_e	Qualification type of therapist e	T_a^{earliest}	First time period in a time window of $a \in \mathbf{App}$
		T_a^{latest}	Last time period in a time window of $a \in \mathbf{App}$
$C_{p,a}^{\text{prio}}$	Hierarchical priority penalty for missed appointments $a \in A(p)$ of patient p	$T_{e,a}^{\text{earliest}}$	Earliest time period an appointment a can be treated by therapist e
$C_{e,p,a}^{\text{violation}}$	Hierarchical penalty cost for preference violations for therapist e when being responsible for an appointment $a \in A(p)$ of patient p	$T_{e,a}^{\text{latest}}$	Latest time period an appointment a can be treated by therapist e
M	Sufficiently large value	S_a	Treatment time in periods for appointment a
		$D_{a,a'}$	Minimum number of periods between appointments a and a' with $a' \in F(a)$
<i>Variables specific to the IP formulation</i>		$\text{pre}(a)$	Preceding appointment, directly scheduled before insertion point of appointment a within $\mathbf{Sol}(e)$
$x_{p,a}$	1 if an appointment $a \in A(p)$ of patient p is not fulfilled within its time window, 0 otherwise		

$z_{e,p,a,t}^{\text{start}}$	1 if a therapist e starts appointment $a \in A(p)$ of patient p in period t , 0 otherwise	$\text{suc}(a)$	Succeeding appointment, directly scheduled after insertion point of appointment a within $\text{Sol}(e)$
		$W_{\text{pre}(a),a}$	Minimum number of periods between two appointments $\text{pre}(a)$ and a , i.e., transition time of a therapist
		$C(r)$	Insertion candidates, i.e., combination between therapist and appointment for preference violation class r
		$C_{e,a}^{\text{violation}}$	Preference violation cost if matching therapist e with appointment $a \in \mathbf{App}$
		$U_{(a,\text{suc}(a)),e,e'}$	Change in preference violation cost if moving appointments a and $\text{suc}(a)$ from therapist e to therapist e' within neighborhood 2
		$U_{a^+,(a,\text{suc}(a)),e,e'}$	Cost $U_{(a,\text{suc}(a)),e,e'}$ plus the preference cost of the combination between unscheduled appointment a^+ and therapist e if inserting unscheduled appointment a^+ in schedule $\text{Sol}(e)$ in neighborhood 2
		$U_{a,a'}$	Change in preference violation cost if interchanging appointments a and a' in neighborhood 5
		$U_{a^+,e,a,a'}$	Change in preference violation cost if interchanging appointments a and a' in neighborhood 5, i.e., getting cost $U_{a,a'}$, and then adding unscheduled appointment a^+ to schedule $\text{Sol}(e)$

Appendix G: Comparison of phase 1 solutions and final GRASP solutions

Table 14 compares the best achieved GRASP solution (*GRBE*) using five runs for each of the eleven instances, with one run only using phase 1 (*GRPI*). The column $|P|/|A|/|E|$ includes the number of patients $|P|$, the number of appointments $|A|$, and the number of therapists $|E|$. US is the number of unscheduled appointments and S is the number of scheduled appointments. *Prio* informs about unscheduled appointments within the five priority violation classes. *Pref* are the scheduled appointments within the preference violation classes. Compared to *GRPI*, *GRBE* schedules additional appointments for all instances, improving our primary objective (4.1). For

most instances, this leads to a slightly worse preference adherence, i.e., our secondary objective (4.2), where fewer appointments are scheduled in *Pref5* and *Pref4*. However, instance 300/372/30 illustrates that sometimes scheduling more appointments in phase 2 can also lead to higher numbers of scheduled appointments in the most desired preference classes. The reason is that newly scheduled appointments have a high preference class for the therapist they are scheduled on. In this instance, *GRBE* schedules 15 appointments more than *GRP1*, and more scheduled appointments can be seen in *Pref5*, *Pref4*, and *Pref3*, the most desired preference classes. The same behavior can be observed for the instance 80/98/10.

Appendix H: Relative preference results

Table 15 shows the mean percentage of scheduled appointments within their relative preference classes. For example, if an appointment has an available therapist within preference class $r = 5$ and the appointment is scheduled with a therapist $r = 5$, we have *0 Classes Difference* between the achieved and the optimal preference class for an appointment. If the optimal available class for an appointment would be $r = 4$ and an appointment is scheduled with a therapist with $r = 1$, we have *3 Classes Difference*. Mean and standard deviation are given across five GRASP runs for our five larger-sized instances. The results show small differences between the instances and the small standard deviations show hardly any differences between the results of five GRASP runs. The mean percentage in the best class slightly increases for larger instances, where more therapists per qualification are available. For all instances, at least 71.18 percent of the appointments can be scheduled in their optimal available preference (violation) class, i.e., *0 Classes Difference*. For all instances, more than 93 percent of appointments are scheduled at the worst with one class difference, i.e., appointments scheduled in *0 Classes Difference* or *1 Class Difference*.

Appendix I: Performance of the GRASP for other problem settings

Two additional, vastly different, instances using the problem size of 500 patients and 50 therapists are used to get results with varying parameters. In the first instance (*I#1*), the recovery time within a therapy pathway is set to zero for all patients' appointments, i.e., appointments of the same patient can be performed, unless another appointment of the patient is processed. In this instance the problem complexity is reduced. 120 appointments remain unscheduled (i.e., slightly less than for the original instance with 124 unscheduled appointments in the best case and 128 in the worst case) and schedules are similarly tight as shown in Table 7 for the original instance. We could not recognize any interesting patterns for this new instance concerning scheduling results. In the second instance (*I#2*), all appointments are set to the same priority, i.e., instead of five priority classes we only use one. The added flexibility in this instance does not lead to more scheduled appointments. 127

Table 14 Comparison of Phase 1 Solutions and Best Achieved Total GRASP Solution

$ P / A / E $	Method	Number of unscheduled appointments					Number of scheduled appointments						
		US	Prio5	Prio4	Prio3	Prio2	Prio1	S	Pref5	Pref4	Pref3	Pref2	Pref1
50/61/5	GRP1	20	0	1	1	4	14	41	17	11	8	4	1
	GRBE	16	0	0	0	2	14	45	18	9	9	7	2
75/93/5	GRP1	45	0	1	1	15	28	48	23	8	8	6	3
	GRBE	41	0	0	0	13	28	52	21	12	8	8	3
75/93/7	GRP1	34	0	1	1	5	27	59	36	9	8	4	2
	GRBE	29	0	0	0	1	28	64	34	10	8	7	5
80/98/10	GRP1	7	0	1	0	1	5	91	63	13	7	5	3
	GRBE	5	0	0	0	1	4	93	65	14	8	4	2
85/103/10	GRP1	7	0	1	0	1	5	96	64	13	8	7	4
	GRBE	7	0	0	0	2	5	96	62	15	7	9	3
100/122/10	GRP1	26	0	1	0	6	19	96	69	12	8	4	3
	GRBE	21	0	0	0	4	17	101	69	15	8	5	4
150/181/15	GRP1	30	0	0	1	5	24	151	81	55	6	7	2
	GRBE	23	0	0	1	1	21	158	79	55	12	7	5
200/242/20	GRP1	49	0	0	1	11	37	193	122	55	9	6	1
	GRBE	37	0	0	1	8	28	205	120	64	13	6	2
300/372/30	GRP1	78	0	0	1	24	53	294	190	87	10	6	1
	GRBE	63	0	0	1	19	43	309	196	85	18	7	3
400/515/40	GRP1	112	0	0	1	46	65	403	300	87	10	5	1
	GRBE	97	0	0	1	39	57	418	296	87	26	6	2
500/655/50	GRP1	135	0	0	2	58	75	520	406	98	14	2	0
	GRBE	124	0	0	2	51	71	531	405	101	15	7	3

Table 15 Percentage of appointments scheduled in relative difference from best possible preference class

$ P / A / E $	Difference between scheduled class and optimal possible class				
	0 Classes difference	1 Class difference	2 Classes difference	3 Classes difference	4 Classes difference
150/181/15	Mean	0.7118	0.2250	0.0405	0.0177
	SD	0.0118	0.0216	0.0118	0.0047
200/242/20	Mean	0.7345	0.2173	0.0265	0.0157
	SD	0.0163	0.0179	0.0040	0.0058
300/372/30	Mean	0.7305	0.2084	0.0448	0.0117
	SD	0.0187	0.0133	0.0088	0.0044
400/515/40	Mean	0.7594	0.1849	0.0378	0.0134
	SD	0.0118	0.0109	0.0055	0.0038
500/655/50	Mean	0.7639	0.1902	0.0291	0.0117
	SD	0.0208	0.0123	0.0059	0.0032
					0.0030

Table 16 Results for additional instances

$ P / A / E $		Number of unscheduled appointments						Number of scheduled appointments					
		<i>US</i>	Prio5	Prio4	Prio3	Prio2	Prio1	<i>S</i>	Pref5	Pref4	Pref3	Pref2	Pref1
500/655/50	I#1	120	0	0	2	43	75	535	405	103	14	7	6
	I#2	127	–	–	–	–	127	528	416	96	12	3	1

appointments remain unscheduled. As can be expected, when preference becomes the main differencing factor, more appointments are scheduled in the most important preference class compared to the best performing GRASP run in the original instance (416 appointments vs. 405 in the original instance). The increase is restricted to the most important preference class. All remaining classes lose appointments to the most important class compared to the original instance. Schedules are similarly tight as results in Table 7 and no other pattern could be detected. The described results for the two instances are shown in Table 16 in more detail, where *US* shows the total number of unscheduled appointments within the different *Prio* classes, and *S* shows the total number of scheduled appointments within the different *Pref* classes.

Funding Open access funding provided by Technical University of Denmark.

Availability of data and materials The datasets generated during and/or analyzed during the current study are available from the corresponding author on reasonable request.

Declarations

Conflict of interest There is no conflict of interest in the study.

Ethics approval, consent to participate and consent to publication The study does not require an ethics approval, consent to participate and consent to publication.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ahmadi-Javid A, Jalali Z, Klassen KJ (2017) Outpatient appointment systems in healthcare: a review of optimization studies. *Eur J Oper Res* 258:3–34. <https://doi.org/10.1016/j.ejor.2016.06.064>
- Ait Haddadene SR, Labadie N, Prod'homme C (2016) A GRASP × ILS for the vehicle routing problem with time windows, synchronization and precedence constraints. *Expert Syst Appl* 66:274–294. <https://doi.org/10.1016/j.eswa.2016.09.002>

- Beattie P, Dowda M, Turner C, Michener L, Nelson R (2005) Longitudinal continuity of care is associated with high patient satisfaction with physical therapy. *Phys Ther* 85:1046–1052. <https://doi.org/10.1093/ptj/85.10.1046>
- Bowers J, Cheyne H, Mould G, Page M (2015) Continuity of care in community midwifery. *Health Care Manag Sci* 18:195–204. <https://doi.org/10.1007/s10729-014-9285-z>
- Bundesagentur für Arbeit (2022) Fachkräfteengpassanalyse 2021: Blickpunkt Arbeitsmarkt | Mai 2022. <https://statistik.arbeitsagentur.de/DE/Navigation/Footer/Top-Produkte/Fachkraefteengpassanalyse-Nav.html>. Accessed 14 Nov 2022
- Burke EK, de Causmaecker P, Vanden Berghe G, van Landeghem H (2004) The state of the art of nurse rostering. *J Sched* 7:441–499. <https://doi.org/10.1023/B:JOSH.0000046076.75950.0b>
- Cappanera P, Scutellà MG (2015) Joint assignment, scheduling, and routing models to home care optimization: a pattern-based approach. *Transp Sci* 49:830–852. <https://doi.org/10.1287/trsc.2014.0548>
- Cheang B, Li H, Lim A, Rodrigues B (2003) Nurse rostering problems—a bibliographic survey. *Eur J Oper Res* 151:447–460. [https://doi.org/10.1016/S0377-2217\(03\)00021-3](https://doi.org/10.1016/S0377-2217(03)00021-3)
- Chien C-F, Tseng F-P, Chen C-H (2008) An evolutionary approach to rehabilitation patient scheduling: a case study. *Eur J Oper Res* 189:1234–1253. <https://doi.org/10.1016/j.ejor.2007.01.062>
- Cissé M, Yalçındağ S, Kergosien Y, Şahin E, Lenté C, Matta A (2017) OR problems related to home health care: a review of relevant routing and scheduling problems. *Oper Res Health Care* 13–14:1–22. <https://doi.org/10.1016/j.orhc.2017.06.001>
- Dohn A, Kolind E, Clausen J (2009) The manpower allocation problem with time windows and job-teaming constraints: a branch-and-price approach. *Comput Oper Res* 36:1145–1157. <https://doi.org/10.1016/j.cor.2007.12.011>
- Erhard M, Schoenfelder J, Fügner A, Brunner JO (2018) State of the art in physician scheduling. *Eur J Oper Res* 265:1–18. <https://doi.org/10.1016/j.ejor.2017.06.037>
- Feo TA, Resende MGC, Smith SH (1994) A greedy randomized adaptive search procedure for maximum independent set. *Oper Res* 42:860–878
- Fikar C, Hirsch P (2017) Home health care routing and scheduling: a review. *Comput Oper Res* 77:86–95. <https://doi.org/10.1016/j.cor.2016.07.019>
- Gartner D, Frey M, Kolisch R (2018) Hospital-wide therapist scheduling and routing: Exact and heuristic methods. *IIE Trans Healthc Syst Eng* 8:268–279. <https://doi.org/10.1080/24725579.2018.1530314>
- Girbig M, Freiberg A, Deckert S, Druschke D, Kopkow C, Nienhaus A, Seidler A (2017) Work-related exposures and disorders among physical therapists: experiences and beliefs of professional representatives assessed using a qualitative approach. *J Occup Med Toxicol* 12:1–9. <https://doi.org/10.1186/s12995-016-0147-0>
- Griffith JD, Williams JE, Wood RM (2012) Scheduling physiotherapy treatment in an inpatient setting. *Oper Res Health Care* 1:65–72. <https://doi.org/10.1016/j.orhc.2012.08.001>
- Gupta D, Denton B (2008) Appointment scheduling in health care: challenges and opportunities. *IIE Trans* 40:800–819. <https://doi.org/10.1080/07408170802165880>
- Gurobi Optimization LLC (2022) Gurobi optimizer reference manual. <https://www.gurobi.com>
- Huynh N-T, Huang Y-C, Chien C-F (2018) A hybrid genetic algorithm with 2D encoding for the scheduling of rehabilitation patients. *Comput Ind Eng* 125:221–231. <https://doi.org/10.1016/j.cie.2018.08.030>
- Jungwirth A, Desaulniers G, Frey M, Kolisch R (2021) Exact branch-price-and-cut for a hospital therapist scheduling problem with flexible service locations and time-dependent location capacity. *INFORMS J Comput*. <https://doi.org/10.1287/ijoc.2021.1119>
- Kontoravdis G, Bard J (1995) A GRASP for the vehicle routing problem with time windows. *ORSA J Comput* 7:10–23. <https://doi.org/10.1287/ijoc.7.1.10>
- Kuiper A, Lee RH (2022) Appointment scheduling for multiple servers. *Manag Sci* 68:7422–7440. <https://doi.org/10.1287/mnsc.2021.4221>
- Lau B, Skinner EH, Lo K, Bearman M (2016) Experiences of physical therapists working in the acute hospital setting: systematic review. *Phys Ther* 96:1317–1332. <https://doi.org/10.2522/ptj.20150261>
- Nickel S, Schröder M, Steeg J (2012) Mid-term and short-term planning support for home health care services. *Eur J Oper Res* 219:574–587. <https://doi.org/10.1016/j.ejor.2011.10.042>
- Olaleye OA, Hamzat TK, Akinrinmade MA (2017) Satisfaction of Nigerian stroke survivors with outpatient physiotherapy care. *Physiother Theory Pract* 33(1):41–51. <https://doi.org/10.1080/09593985.2016.1247931>
- Ogulata SN, Koyuncu M, Karaskas E (2008) Personnel and patient scheduling in the high demanded hospital services: a case study in the physiotherapy service. *J Med Syst* 32:221–228. <https://doi.org/10.1007/s10916-007-9126-4>

- Pinedo M (2016) Scheduling: theory, algorithms, and systems. Springer, Cham, Heidelberg, New York, Dordrecht, London
- Podgorelec V, Kokol P (1997) Genetic algorithm based system for patient scheduling in highly constrained situations. *J Med Syst* 21:417–427. <https://doi.org/10.1023/A:1022828414460>
- Rasmussen MS, Justesen T, Dohn A, Larsen J (2012) The home care crew scheduling problem: preference-based visit clustering and temporal dependencies. *Eur J Oper Res* 219:598–610. <https://doi.org/10.1016/j.ejor.2011.10.048>
- Rojanasoonthon S, Bard J (2005) A GRASP for parallel machine scheduling with time windows. *INFORMS J Comput* 17:32–51. <https://doi.org/10.1287/ijoc.1030.0048>
- Schimmelpfeng K, Helber S, Kasper S (2012) Decision support for rehabilitation hospital scheduling. *Or Spectrum* 34:461–489. <https://doi.org/10.1007/s00291-011-0273-0>
- Song H, Tucker AL, Murrell KL (2015) The diseconomies of queue pooling: an empirical investigation of emergency department length of stay. *Manage Sci* 61:3032–3053. <https://doi.org/10.1287/mnsc.2014.2118>
- van Dijk N, van der Sluis E (2008) To pool or not to pool in call centers. *Prod Oper Manag* 17:296–305. <https://doi.org/10.3401/poms.1080.0029>
- van Dijk N, van der Sluis E (2009) Pooling is not the answer. *Eur J Oper Res* 197:415–421. <https://doi.org/10.1016/j.ejor.2008.06.014>
- Verband der Ersatzkassen (2022) Rahmenverträge. <https://www.vdek.com/vertragspartner/heilmittel/rahmenvertrag.html>. Accessed 12 Apr 2022
- van Walraven C, Oake N, Jennings A, Forster AJ (2010) The association between continuity of care and outcomes: a systematic and critical review. *J Eval Clin Pract* 16:947–956. <https://doi.org/10.1111/j.1365-2753.2009.01235.x>
- Wichmann MG, Volling T, Spengler TS (2014) A GRASP heuristic for slab scheduling at continuous casters. *Or Spectr* 36:693–722. <https://doi.org/10.1007/s00291-013-0330-y>
- Wirnitzer J, Heckmann I, Meyer A, Nickel S (2016) Patient-based nurse rostering in home care. *Oper Res Health Care* 8:91–102. <https://doi.org/10.1016/j.orhc.2015.08.005>
- Zhao L, Chien C-F, Gen M (2018) A bi-objective genetic algorithm for intelligent rehabilitation scheduling considering therapy precedence constraints. *J Intell Manuf* 29:973–988. <https://doi.org/10.1007/s10845-015-1149-y>
- Zimelman JL, Juraschek SP, Zhang X, Lin VW-H (2010) Physical therapy workforce in the United States: forecasting nationwide shortages. *PM R* 2:1021–1029. <https://doi.org/10.1016/j.pmrj.2010.06.015>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Sebastian Kling¹ · Sebastian Kraul² · Jens O. Brunner^{1,3,4} 

✉ Jens O. Brunner
jotbr@dtu.dk; jens.brunner@uni-a.de

Sebastian Kling
sebastian.kling@uni-a.de

Sebastian Kraul
s.kraul@vu.nl

¹ Faculty of Business and Economics, University of Augsburg, Universitätsstraße 16, 86159 Augsburg, Germany

² Department of Operations Analytics, School of Business and Economics, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

³ Department of Technology, Management, and Economics, Technical University of Denmark, Kgs. Lyngby, Denmark

⁴ Data and Development Support, Region Zealand, Denmark