Advancing Medical Resident Scheduling

Improving Human-Centricity in Planning of Medical Education

Alexander Gaal, Wolfgang Dummer, Thomas Sobottka,
Fazel Ansari, Sebastian Schlund

1. Introduction

Medical residents are physicians in training who must complete a residency program to become licensed specialists or general practitioners in Austria. The residency program consists of several modules that cover different disciplines related to the chosen specialty. The modules must be completed in specific hospitals that offer duty and training positions for the residents. The process of assigning residents to these positions over a planning horizon of up to 72 months is called residency scheduling. This process is challenging and complex, as it involves multiple stakeholders, objectives, constraints, and preferences. At present, resident scheduling in Austria is mostly conducted manually using generic digital tools such as Microsoft Excel. This leads to high planning effort and inconsistent planning quality for stakeholders, depending on the individual knowledge and skills of the planners. In case of a hospital network analyzed by the authors, 197 working hours are spent each week on training planning for around 200 residents. 52% of these hours are worked by qualified and experienced physicians who are not available for day-to-day hospital operations during this time. The remaining work is carried out by dedicated planning staff, which corresponds to two full-time and one part- time staff working only on the planning and administration of resident training. Automating or at least supporting this planning with algorithms is currently impossible due to lacking available methods. Focusing on the literature, state-of-the-art planning methods for resident training cannot be directly employed due to the specifics and complexity of the Austrian system for resident training, cf. section 2. The main challenge in the Austrian training system is a two-stage allocation of resources, i.e. i) allocating residents to duty and training positions in stations at the same time, with ii) allocations subject to further requirements (e.g. hospital changes, personal preferences , cf. Section 2.1. This considerably enlarges the solution space and requires a planning method tailored to this requirement.

DOI: https://doi.org/10.30844/wgab_2023_7

In related works, exact solution methods using mixed integer programming (MIP) or column generation are applied on much smaller problem instances. Those involve far less constraints and problem features, while still requiring significant computational resources and effort. Using such methods for solving the real-world problem studied in this paper on a large scale would not be reasonable in terms of computational effort and time. To address this problem, this paper proposes a combination of heuristic and metaheuristic methods based on a constructive heuristic and a genetic algorithm (GA) – with GAs being one of the most used method classes for similar planning problems, as demonstrated in Section 2. In a previous approach by the authors, a purely metaheuristic method was applied to the same problem formulation and problem instances (Dummer et al. 2023). However, the metaheuristic approach had limitations concerning runtime performance and solution quality. Pursuing this line of research, this paper deals with the medical resident scheduling problem (RSP), which is a tactical scheduling problem that aims to find a feasible and high-quality schedule for the residents over the entire training period. A hybrid solution method Is proposed that combines a constructive heuristic and a genetic algorithm (GA). Further, two new objectives are introduced to address fairness and equity-related aspects (i.e. variance of the total training duration and consideration of trainee preferences) when scheduling, c.f. Section 4.1. Accordingly, the proposed method is evaluated on real examples and compared with human planning results as well as the earlier method developed by the authors, c.f. Section 5.

The rest of the paper is organized as follows. In Section 2, the relevant literature for the RSP is reviewed. In Section 3, a detailed problem description and a mathematical formulation of the problem is provided. Section 4 discusses the solution methodology developed. Section 5 focuses on validation of the performance of the proposed method. Finally, Section 6 discusses the key findings and provides a research outlook.

2. Related Work

The RSP has been studied in various forms and specifications. This paper analyzes the problem in terms of (i) problem characteristics and (ii) relevant solution methodologies based on a literature review. This literature analysis will be based on Akbarzadeh's comprehensive literature analysis (Akbarzadeh/Maenhout 2021), using a similar analysis structure.

2.1. Characteristics of the problem

Medical resident scheduling is a complex problem because of the stakeholders involved, namely, legislation, hospitals, residents as well as the medical board and medical schools. These stakeholders may have different and conflicting requirements and objectives. An overview of the problem characteristics that have been reviewed in the literature is discussed below.

Legislation and planning complexity: Each country has its own legal framework for medical education. Although there are certain similarities between European medical education systems, it Is practically not affordable to formulate a universal problem. In Austria, there are currently just over 8,300 residents in training. All 271 public hospitals in Austria have a training mandate and perform some degree of resident scheduling. On average, about 30 residents are trained in each hospital. This figure varies greatly depending on the size of the hospital. The largest hospitals train nearly 400-500 residents at a time. The ordinance of the Federal Ministry of Health on "Training to become a general practitioner and a specialist" (Ärzteausbildungsordnung 2015, abbreviated to AO2015) constitutes a set of regulations for the proof of successful completion of practical training in general medical and specialist training. Public medical universities currently offer a total of around 1,540 study places per year in the human medicine program. Currently, about 8,000 graduates are undergoing training to become general practitioners or medical specialists (Bundesministerium für Gesundheit 2022; Rechnungshof Österreich 2021).

Due to almost unique characteristics of national/regional legal framework for medical education, scientific publications usually refer to specific countries or regions. Several publications consider different residency training programs and hospital resources specifically dedicated to them and solve the problem for different programs separately (Kraul et al. 2019; Diponegoro/Rukman 2017 - 2017; Guo et al. 2014). Other papers feature resources shared between different training programs and distinguish between junior and senior students are trained in different programs (Bard et al. 2016, 2017; Proano/Agarwal 2018). Requirements can be adjusted within certain limits (e.g., elective modules) to suit the interests of individual residents, allowing specializations in a particular area (ITO et al. 2018). Few authors modeled direct precedence relationships between two training modules, or disciplines or sections (e.g., is anesthesiology as the basis for general surgery) (Brech et al. 2019). In the case at hand, the training stages (basic and main training) must be planned with precedence relationships, and duty and training positions at the same time.

Hospitals: Hospitals are responsible for the implementation of training programs, i.e., attending physicians supervise residents who are involved in the day-to-day operation of the hospital. As a result, many authors have included staffing requirements in their problem definition. Almost all studies consider a maximum student capacity (Proano/Agarwal 2018; Smalley/Keskinocak 2016). In some cases minimum staffing requirements are included because hospitals count on students as part of the required staff to handle the workload (Guo et al. 2014; Ryan et al. 2013). All above mentioned publications assume that hospital resources are dedicated to a particular specialty. In most cases, minimum and maximum staffing requirements are also specified for a specialty.

DOI: https://doi.org/10.30844/wgab_2023_7

Residents: A limited number of resident requirements are considered in various publications. In most publications, resident requirements play a minor role and are often not even listed as stakeholders. Student availability is a common feature in many studies (Proano/Agarwal 2018; Smalley/Keskinocak 2016). Preferences for a particular specialty were also considered in a few other publications (Diponegoro/Rukman 2017 - 2017; ITO et al. 2018; Smalley/Keskinocak 2016). Other requirements, such as preferences regarding hospitals, departments, or training physicians, were not considered, and must be part of the approach developed in the paper at hand. In addition, to the best of the authors' knowledge, there are no references considering the extent of employment of part-time employees.

Objective function: Several objectives for the RSP can be found in the literature. These objectives differ depending on the considered stakeholders. The following objectives have been already defined: i) provide training that is as fair and equal as possible (Schleyer 1994), ii) carry out the scheduling of all required modules (Guo et al. 2014; Kraul et al. 2019), iii) support the scheduling of all required modules complying with specifications about the training sequence (ITO et al. 2018). In most cases, these training requirements are formulated as hard constraints because of their importance. In addition, Bard et al. (Bard et al. 2016, 2017) consider hospital perspective and minimize violations related to student staffing requirements. Literature shows that objectives related to residents are included in most cases. Smalley and Beliën (Smalley/Keskinocak 2016; Beliën/Demeulemeester 2006) create rotation schedules by making the best use of the availability of residents. Diponegoro (Diponegoro/Rukman 2017) considers preferences for specific modules among other goals. All studies optimize goals of a single stakeholder or at most two stakeholders. To the best of the authors' knowledge, an approach that considers the objectives of all relevant stakeholders involved is not available in the literature.

2.2. Solution Methodology

Finding an exact (optimal) solution (there are multiple optima) to RSP in a real-world setting is very difficult when multiple objectives and multiple requirements are involved. The problem has been shown to be at least NP-complete (Guo et al. 2014), even though Guo et al. dealt with a highly simplified problem with few constraints. Only an assignment to a training and station has been considered. Changes of department or hospital, duty positions, training positions, preferences, etc. have not been considered. In particular, finding an optimal solution for a real-world, large-scale problem requires not only additional computational effort, but also some technical effort, e.g., reformulating the problem or developing special solution methods. The real-world problem at hand, considering the AO2015 and stakeholder interests, is an NP-hard set of problems (Kraul et al. 2019).

The following paragraph gives an overview of the relevant literature of planning methods, with an indication of i) the problem size considered, ii) a categorization of the proposed solution approach, and iii) the computational effort required. Both exact and heuristic methods have been published in the literature are considered.

Several studies propose a greedy heuristic based on decomposition and scheduling individual residents (Diponegoro/Rukman 2017 - 2017; Schleyer 1994). Other authors use mathematical programming and first solve a simpler problem by relaxing some functional constraints or the optimization constraints. Then, they use heuristic methods to optimize further (Franz/Miller 1993; Kraul et al. 2019; Bard et al. 2013). Few publications use exact solution methodology (Bard et al. 2017; Proano/Agarwal 2018; Ryan et al. 2013; Beliën/Demeulemeester 2006, 2007). To solve a large problem instance (i.e., planning many residents simultaneously over a long planning period), some authors reformulate the original problem into a set partitioning problem using Dantzig-Wolfe decomposition (Kraul et al. 2019; Beliën/Demeulemeester 2006, 2007). Column generation is applied to solve the linear programming (LP) relaxation of the reformulated model, and the optimal LP solution is converted to an integer solution using either heuristic techniques or an exact branching scheme. Note that the approaches with an optimal procedure, which requires a significant amount of computational time for large problem instances, contain far fewer problem features and constraints, optimize only a single objective, and are evaluated on instances with few problem dimensions.

Since uncertainties and changed circumstances in scheduling for people cannot be ruled out, re-rostering often occurs after a plan has been disrupted by, e.g. absences due to illness. Maenhout and Vanhoucke (Maenhout/Vanhoucke 2011) describe the problem of reassigning nurses to shifts as a "re-rostering problem", which deals with plans that are invalidated by a change in constraints due to disruptions. The aim of re-scheduling in the context of the RSP is to optimally re-schedule the module that has not been completed while at the same time maintaining most of the existing schedule, to ensure planning security for other residents. In real resident management, rescheduling and postponements occur frequently. To the best of the authors' knowledge, there is no described solution for a resident re-scheduling algorithm in the literature. It is worth noting that Aickelin (Aickelin/Dowsland 2004) uses an approach similar to the one presented. It uses an indirect coding-based approach by creating permutations of nurses with a GA and a heuristic decoder that generates schedules from these permutations. The results are further improved by introducing a hybrid crossover operator and using simple bounds to reduce the size of the solution space. Results show that the algorithm can find high quality solutions while being faster and more flexible than a published tabu search approach. In contrast to the problem at hand, this approach works in the context of short-term deployment planning.

DOI: https://doi.org/10.30844/wgab_2023_7

Akbarzadeh and Maenhout (Akbarzadeh/Maenhout 2021) Introduce a special form of RSP applied in Belgium. They use a heuristic solution procedure consisting of a constructive heuristic and two local search heuristics to optimize the initial solution. The treated RSP corresponds to the Belgian medical training and differs significantly from the practice in Austria, Germany and in particular the USA, where training in hospitals with an exact duration of 12 months must be completed already during the studies. Re-rostering is also not treated. Building on these results, Zanazzo et al. (E. Zanazzo et al. 2022) were able to show that the problem could be solved with a metaheuristic method (i.e. Simulated Annealing) with the same quality but in a significantly shorter time. This evidences that metaheuristic methods have advantages over exact solution methods as problem size and complexity increase, and thus, this will shape a basis for the method development herein.

The authors' investigations also reveal that training quality, fairness or equal treatment are not addressed in the existing optimization literature, i.e. none of the aforementioned planning methods explicitly considers these factors. To sum up, the solutions described in the literature are not affordably and effectively applicable to the complex case of RSP in the Austrian medical education system. In contrast, an optimal solution for the Austrian problem can be transferred to the reduced complexity of other regulations. Hence, the Impact of the approach proposed in this paper can be justified beyond the special case of Austria.

3. Problem definition

This section formally describes the RSP, in which individuals are assigned to a predefined set of medical disciplines (e.g., internal medicine) at multiple hospitals for a period of up to 72 months during their medical training to ensure that they receive adequate medical education. The time unit of month is specified by the AO 2015 and the planning horizon is divided into several sections. First, general practitioners (as well as aspiring specialists) undergo nine months of basic training, followed by actual training in the hospital followed by an internship with an established physician in a certified teaching practice. After basic training, specialists begin basic specialty training and then undergo specialty training in their field. Scheduling must be carried out simultaneously for all people involved, irrespective of their educational progress (i.e., different stages and subjects) as they can possibly be assigned to the same resources. During a given training, residents are assigned to different departments and hospitals, where they work for one or more (consecutive) months under the supervision of an attending physician to ensure that they acquire the required competencies. In particular, the data set was derived from the case of an Austrian hospital organization with 9 hospitals and headquarters in Vienna, Austria, which trains 240-300 residents annually.

Various constraints and multiple objectives related to the three main actors, i.e. legislation, local hospitals and residents, characterize the problem and are described in this section.

Requirements and constraints:

Legislation

The legislation in AO2015 (Bundesministerium für Gesundheit 2022) establishes different educational requirements to ensure that each resident acquires the intended competencies. These requirements are:

- Training should be scheduled in full months.
- Training stages must be completed in the predefined sequence, e.g., basic training before general medicine training.
- Basic training must include both conservative and surgical subjects. Ideally, these should be evenly distributed (e.g. 5 months of conservative subjects and 4 months of surgical subjects).
- Training stages (e.g., general medicine training, specialist training) include a series of educational subjects supplemented by electives that residents may choose from, based on their personal preferences. Successful completion of the training stage requires that all subjects be completed in the defined length and a minimum number of electives be completed.
- Only fully completed electives contribute to overall educational progress; partially completed electives do not.
- Subjects do not have to be completed in one go.
- training can be suspended for a while and continued later.
- Subjects cannot be taken more than once and the duration of training in a subject cannot be extended.
- A person's educational progress is proportional to the extent of his or her employment (half-time employment means half the educational progress of a full-time employee, resulting in a doubling of educational time.).
- A complete training schedule for a particular resident must be compiled at the time a person enters training.

DOI: https://doi.org/10.30844/wgab_2023_7

Hospitals

The set of rules and requirements for planning often vary from hospital (group) to hospital (group) where resident's training takes place. Yet, some scheduling requirements are common in most Austrian hospitals:

- Residents complete training subjects in multiple hospitals, as not every hospital has every medical department needed for training.
- Duty positions are available at different hourly extents, and training positions are available only for 35 hours per week. Both can be shared in different ways among several trainees.
- The resource capacity of the positions may be dedicated to a single specialty or shared across several related subjects.
- A minimum and maximum resident requirement is established for each department and hospital.
- At all times, a resident must be assigned to one or more duty positions that meet their scope of employment. To achieve educational progress, a resident must also be assigned to a training position. Otherwise, this results in a month(s) without training.
- To facilitate planning, department-independent duty positions (i.e., positions that are shared among multiple departments of a certain hospital) can be used as a fallback option if there is insufficient capacity in a department or more training positions have been approved than duty positions.

Residents

In Austria, a resident is a physician in training to become a general practitioner or a physician in training to become a specialist. Their scheduling requirements include:

- Residents can start training at the beginning of any given month.
- Residents can only be trained in disciplines according to their ability (i.e., educational progress and specialty) and when available.
- certain training subjects can be assigned according to preferences.
- Residents may pause, de- or increase their extent of employment at any time.
- The training schedule for the next 12 months must always be known and available to the residents.

Objective Function

All the requirements mentioned so far are specified as hard constraints. The objective function consists of several components that aim to ensure planning quality with respect to the corresponding objectives. To evaluate a resident schedule, the following objectives of the various stakeholders must be considered:

- **Sum** of Months without training, department changes, hospital changes, single month assignments, violated preferences.
- **Variance** of months without training, violated preferences.

Consequently, the objective function considers multiple objectives summed over all residents, i.e. a global sum of negative effects is computed. These summed values are normalized to a range [0,1], with respect to the characteristic sizes of a metric (i.e., total planned months, planned persons, considered preferences) to ensure problem size independence. The variances of negative factors across all residents enforce fairness aspects and equal treatment. This follows the definition of Stolletz and Brunner (Stolletz/Brunner 2012), who define that "fairness can be seen as how violations of preferences are balanced across employees". The objective function is formulated as a weighted sum function, where all individual objectives are considered linearly with appropriate weights depending on a planner's preferences and goals.

4. Method development

The problem dimensions and features considered in this work result in multiple binary decision variables and constraints. As stated in Section 2, a hybrid method of heuristics and metaheuristics Is proposed featuring a greedy constructive heuristic and a GA to provide the scheduling orders for the heuristic. Akbarzadeh and Maenhout (Akbarzadeh/Maenhout 2021) employ a heuristic procedure to solve the RSP relying on a constructive heuristic and two local search heuristics trying to improve the constructed initial solution. In contrast, the method developed herein does not iteratively improve and optimize an initial solution, but rather performs continuous replanning by generating many solutions guided by evolutionary behavior. The constructive heuristic follows all the constraints defined in AO2015 when scheduling residents by mimicking real-world scheduling processes performed by human planners (cf. Section 4.2). The execution of the greedy heuristic for a given input-sequence yields a set of training assignments. Such a training assignment represents the assignment of a person to one or more duty positions and to zero to many training positions for a particular month. These attributes imply the assigned training subject (aka module), the department and thus the hospital. This set of training assignments is in turn used to evaluate the solution's fitness with respect to the previously described objectives.

DOI: https://doi.org/10.30844/wgab_2023_7

4.1. Development of the objective function

Table 1 shows a schematic representation of a scheduling calendar, a 2-D array of assignments that is the basic representation of a training schedule for a given number of residents and time horizon. Each row of the array represents an individual resident's schedule, and each column represents a time period, i.e. a month. Each resident's schedule must contain at least one valid duty position that defines the person's schedule and should contain a valid training position to ensure that the person's training progresses. Duty and training positions can be represented as integers. Therefore, a valid assignment (pi, qj) consists of a tuple of two integers defining the position p_i and the training position q_j occupied by a person.

R/T	$\mathbf{T_1}$	$\mathbf{T_2}$	**…**	$\mathbf{T_n}$
R_1	(pi, qj)	(pi, qj)		(pi, qj)
R_2	(pi, qj)	(pi, qj)		(pi, qj)
…				
R_n	(pi, qj)	(pi, qj)		(pi, qj)

Table 1: Assignments of residents (R) to duty (p) and training positions (q) (Dummer 2022)

Table 2 provides an overview of all components of the objective function considered. The objectives associated with these metrics were identified in collaboration with an Austrian expert and former scheduler in the field of staff scheduling for residents, considering all identified relevant stakeholder interests. Separately, an attempt is made to obtain a holistic view of all related interests as well as fairness aspects in resident scheduling. To ensure minimal training time, it is critical to avoid months without training (MWT). Unnecessary changes between departments (DC) or even hospitals (HC) negatively impact training quality by wasting more training time on administrative and familiarization processes. In addition, more frequent rotations may negatively impact patient outcomes (Denson et al. 2015). Single-month assignments (SMAs) are the worst-case scenario, when it comes to changes of location during training. SMAs are assignments to a specific department that last only one month, so an individual enters and leaves that department in the same month. As described, residents can choose between a few electives in their curriculum, depending on personal preferences. Any preferred subject that is not included in a created plan is considered a violated preference (VP). In the case of long-term absences (e.g. illness) or drop-outs, no training progress is taken into account. The monthly rescheduling examines which months of training have been completed and adjusts the schedule accordingly. The planner is advised to set an interruption period in the system for longer absences, or to remove the person from the pool if the training has been completed.

No.	Metric	Variable	Weight w_i
1	∑ Months without Training	MWT	200
2	∑ Department Changes	DC	50
3	∑ Hospital Changes	HC	500
4	∑ Single Month Assignments	SMA	500
5	∑ Violated Preferences	VP	50
6	Variance of MWT among residents	V-MWT	5
7	Variance of VP among residents	V-VP	5

Table 2: Metrics used in the objective function.

To satisfy the residents' need for fairness and equal treatment, two additional metrics have been added in this paper, each assessing the variance, i.e., the equal distribution of negative aspects of planning. This approach is discussed more often in research on workforce scheduling and is also known, for example, in the scheduling of train drivers, i.e. unpopular routes or shifts are distributed as equally as possible among all drivers (Jütte et al. 2017). When creating a shift plan with multiple objectives, the pursuit of equity is usually accompanied by a deterioration in other objectives, such as cost or efficiency. For the present planning case, the two metrics are i) the variance of months without training (V-MWT) and ii) the variance of violated (preferred subject) preferences (V-VP) among the trainees. The weighting of the individual goals is currently subject to a planner's strategy . The weights used in the objective function (c.f. Table 2) were determined experimentally in collaboration with a planning expert and iteratively refined in the evaluation process. Accordingly, the optimization problem consists of 17 stakeholder requirements, 7 objectives considered for evaluation and is formulated as a minimization problem. The objective function is shown in Equation 1.

$$O(\mathrm{x}) = \sum_{i=1}^{7} f_i\,(\mathrm{x}) \times w_i$$

Equation 1: Objective function.

The associated objective function f_i for each individual objective (c.f. Table 2) is defined on x, which is the mathematical representation of a training schedule (i.e. set of training assignments). The corresponding weight for a particular objective is denoted with w_i.

DOI: https://doi.org/10.30844/wgab_2023_7

4.2. Optimization method

In this section, the development of the planning method is described. In the optimization process, a GA modulates the order in which a greedy heuristic schedules the necessary training modules for all residents. This approach is inspired by Aickelin (Aickelin/Dowsland 2004) who presented a similar approach to solve the Nurse Scheduling Problem, which deals with short-term operational scheduling. To enforce elitism concerning the considered objectives, all six individual goals and $f_i\ (\mathrm{x})$ are used as criteria for a non-dominated sorting survival mechanism proposed by Deb (Deb et al. 2002). The selection operator (i.e. tournament selection), utilizes the value of $O(\mathrm{x})$ (c.f. Equation 1). The consideration of the objective function, comprising all previously defined weights w_i , in the selection operator, allows for user-specific customization of the algorithm by introducing bias to the selection process. Thus, it gives control over the planning process with respect to prioritization of specific objectives. To maintain diversity in the population of the GA, a population size of $\mu = 40$ is used, which was determined experimentally considering convergence speed, overall obtained solution quality and runtime. Populations of this size have been shown to provide sufficient robustness for computationally intensive evaluation functions in multimodal search spaces (Kamhuber et al. 2020).

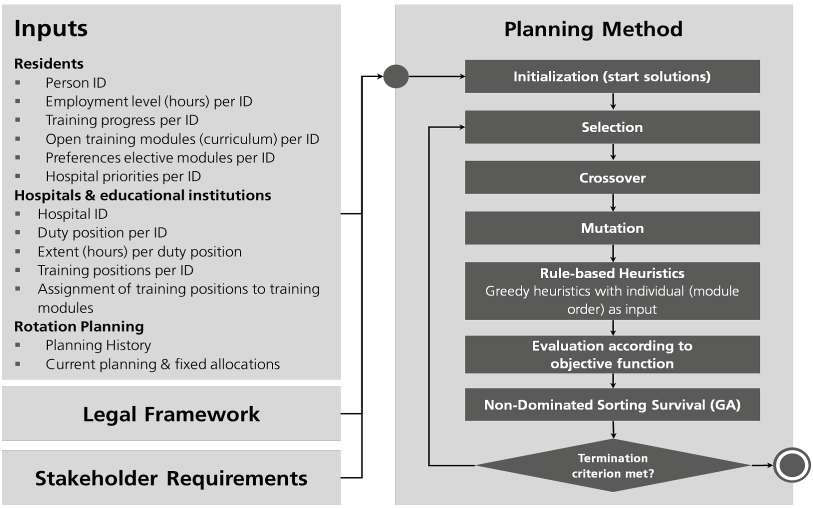

Figure 1: Flowchart of the presented planning method

A general overview of the algorithm's operating principle is shown in a flowchart in Figure 1. Initially, a set of integers, containing all necessary training subjects for all considered residents, is created. Initial solutions for the algorithm are obtained by creating randomly shuffled vectors from said set. Those vectors are the genotypes of the Individuals of the GA. The genotype of an individual determines the order in which the greedy scheduling algorithm plans each module. Thus, the optimization problem solved by the GA is a permutation problem of a 1-dimensional vector. The algorithm uses a simple swap mutation operator, and no crossover operator is used as it did not prove to be beneficial in experimental studies. The genetic algorithm used can be classified as similar to the NSGA-II (Deb et al. 2002) due to its survival mechanism. After creating new individuals by selecting parents and applying the genetic operators to them, individuals are evaluated by passing their genotype to the greedy scheduling heuristic. This rule-based heuristic initially applies a simple repair algorithm that reorders the list of integers to prevent infeasible planning solutions (i.e., sequence violations). After repairing the input list, the heuristic schedules each listed training module sequentially by greedily selecting the best of the remaining combinations of training and duty positions. In this case, "best" means the algorithm tries to prevent department or even hospital changes at all costs. Notably, a person will only switch departments or hospitals if necessary. A month without training is only assigned if there is no free combination of training and duty position available. After the planning process for each genotype is completed, the created plans are evaluated using the objective function and the obtained objective values are returned and assigned to the individuals of the GA. A new population is formed by applying the non-dominated sorting mechanism to the newly created population and the previous one. The process terminates, once the specified termination criterion is met (i.e., number of generations or convergence of fitness).

5. Computational Study

This section provides computational insights into the proposed method. Further, It carries out a benchmark comparison with the results of the previous GA-only approach by the authors (Dummer et al. 2023) (hereafter referred to as GA-method) and a reference solution created by a human planning expert. The proposed algorithm was implemented in Python and Rust, and all tests were performed on a Core i7-10510U CPU, 16 GB of RAM and Ubuntu 22.04 LTS. The test dataset is based on anonymized data that has been enriched with additional information such as preferences and absence times. It contains 78 residents and 5 hospitals. The dataset is the same as in (Dummer et al. 2023). When enriching the test dataset with preferences, the final schedule was used to retroactively assign preferences, i.e., the reference solution is optimal concerning VP. Thus, this reference solution is harder to beat.

The comparison of the three schedule, namely expert solution, GA-method and proposed method is shown in Figure 2. Note that the original objective function (i.e., without the variance metrics) from the GA-method is used to establish a fair comparison of the three methods. The achieved objective-function values of the expert solution (40.98) and the GA-method (81.53) are plotted as horizontal

DOI: https://doi.org/10.30844/wgab_2023_7

dashed lines. The vertical dotted line indicates the maximum number of objective-function evaluations performed in (Dummer et al. 2023). The proposed method outperforms the GA-solution after a few generations and the expert's solution is surpassed after around 1,000 generations. After 50,000 generations an objective-function value of 17.81 is achieved, but already after 20,000 generations a value of 20 is reached. The expert solution is outperformed by the solution method presented in all single objectives except, unsurprisingly, for the violated preferences (VP) (c.f. Figure 3). In Figure 3, all other objective values almost immediately surpass the expert's solution in terms of quality, only the improvements regarding personal preferences take more time.

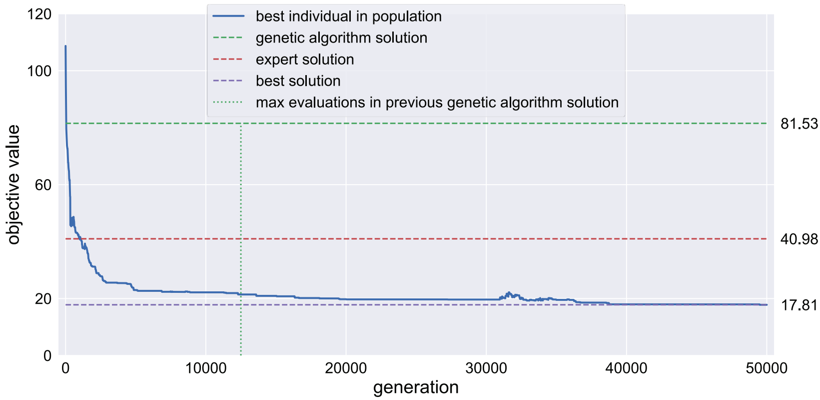

Figure 2: Benchmark: proposed method vs. previous GA method vs. expert solution

Table 3 compares the runtimes of the GA-method and the hybrid method presented and shows the influence of the population size for the hybrid method. The differences in runtime can be explained by the efficient design of the planning method and the use of the Rust programming language compared to Python. For an optimization run of 50,000 generations with a population size of $\mu = 10$, the Python-based GA-method requires 4,200 minutes, while the presented hybrid solution requires only 235 minutes. Thus, the older method is almost 18 times slower than the presented method. This value could be further reduced by using more powerful hardware. Since the new planning method generates better solutions than the expert and the GA-method after only 1,000 generations, good schedules can be generated within one minute. However, this value will increase with the size of the solution space (i.e., more residents, hospitals etc.). The exact time spent to create the manually compiled expert solution has thus far not been recorded. In particular, a few hours were reported.

Method	Pop Size	500 Gens [s]	50.000 Gens[min]	Runtime /Gen [s]	Runtime/ Individual [s]
GA	10	252	4200	0,504	0,0504
Hybrid	10	14.11	235,17	0.028	0.0024
Hybrid	40	34.63	577.23	0.069	0.0017
Hybrid	100	69,76	1162.63	0.140	0,0014

Table 3: Runtime comparison

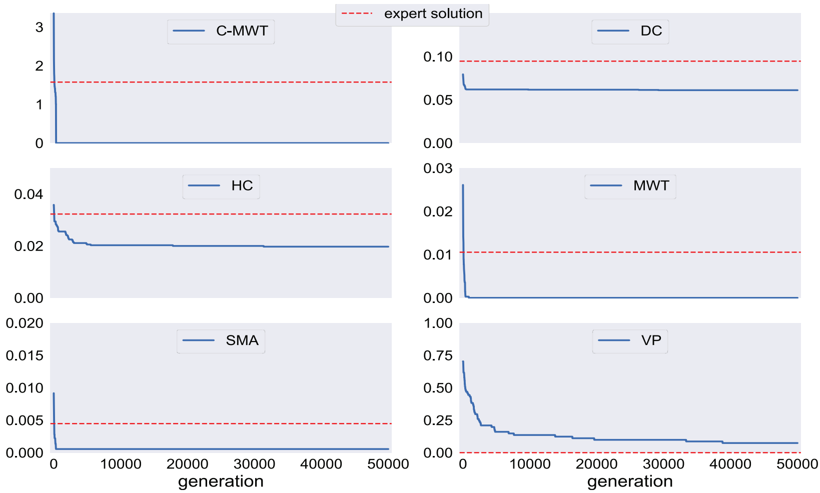

Figure 3: Benchmark normalized single objective values - proposed method vs. expert

The new objective function no longer takes consecutive months without training (C-MWT) into account because the metric had a negative impact on the optimizers efficiency and is rendered obsolete due to the introduction of V-MWT. The resulting function value reaches its minimum at 16.46. Thus, even with two additional objectives and otherwise the same weights as in the prior objective function, the optimization method can find even better solutions than before. In addition, the proposed method performs better in terms of convergence speed when using the newly introduced objective function (compare Figure 2 and Figure 3) almost reaching the final fitness value after about 13,000 generations.

DOI: https://doi.org/10.30844/wgab_2023_7

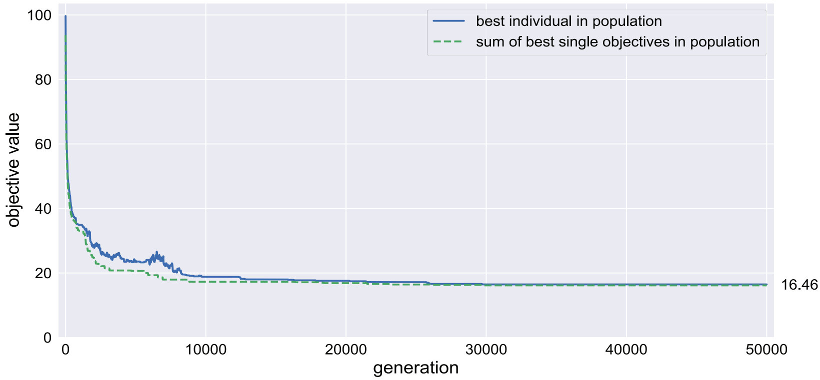

Figure 4: Objective function value of the proposed method with proposed weights

Figure 5 shows the qualitative evolution of the part goals of the new objective function over time for 50,000 generations, each normalized on a range [0,1]. The objective values MWT, V-MWT, SMA and DC and drop almost immediately to their final value. The other objective values converge more slowly towards their optimum.

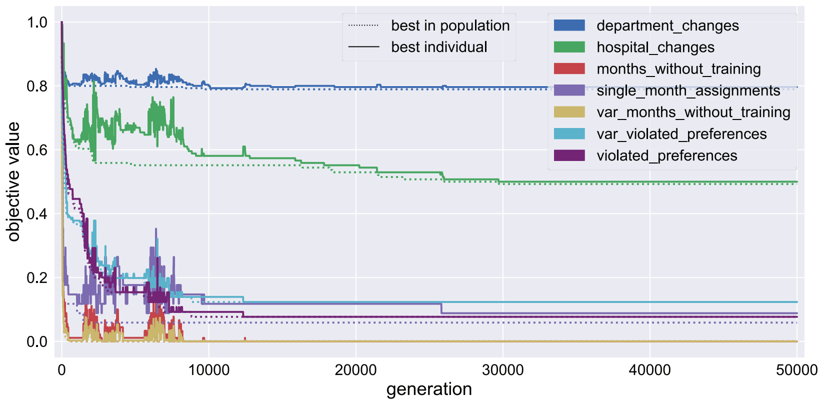

Figure 5: Single objective values for the new configuration of objectives (normalized)

HC, VP and V-VP reach their final value after around 10,000 generations but continue to gradually improve until 30,000 generations. For SMA, some improvement can be seen until just over 25,000 generations, after remaining on a plateau for most of the optimization.

6. Discussion & Outlook

This paper shows that a hybrid optimization method with an indirect GA and rule-based heuristics is suitable to solve a complex RSP and to generate solutions with high planning quality in a short time. At the same time, it was demonstrated how fairness and equity for the residents can be better pursued by extending the objective function with variance metrics. It could successfully be demonstrated that the method achieves solutions with higher quality than human schedulers. The proposed method can potentially reduce the planning effort and improve the planning quality for hospitals by (semi-)automating or supporting the scheduling process. It can also potentially increase the satisfaction and motivation of residents by considering their preferences and ensuring fairness in their training.

Despite merits, the proposed approach faces some limitations concerning the method evaluation and real-world testing, which Identify the pathways for future research. In particular, the criteria of the objective function used in this paper have not yet been finalized. In the next step, the exchange with residents (as target groups) will be intensified to define and better evaluate further aspects of fairness and thus to refine the objective function. According to an initial survey, a final objective function with 12-13 criteria is envisaged. It could make the solution finding more difficult to calculate but will provide solutions with better quality and better acceptance by all stakeholders. Integrating feedback or evaluation (from planner and residents) mechanisms into the solution methodology to track and improve the training performance of residents could be one way to implement this. Another requirement for the solution is that existing schedules should only be changed slightly, especially for the near future, in the event of a rescheduling. This would significantly increase planning reliability for residents, as no major changes would be expected during monthly rescheduling. Also, while the dataset used is already quite large, real-world planning problems from larger hospitals or hospital groups can be more extensive. Hence, in the next step the efforts are extended to a validation with real-world use cases and additionally comparing the plans with existing manual planning from hospitals. Possibilities for employing this approach in strategic workforce planning will also be investigated. Currently, the approval process for new duty and training positions is not transparent, and it is usually not known where these positions are missing to achieve better planning quality. With an extension of the proposed method, it would be possible to search for optimization potentials through new positions by including fictitious additional positions in the planning process. Investigating possible benefits of transferring the method developed for the complex Austrian regulations to other countries and possibly related planning problems will be another research trajectory.

Finally, yet importantly, a further research direction is to use machine learning or artificial intelligence (AI) methods for optimization. One option is to use data-based methods that learn from data, based on manually obtained planning

DOI: https://doi.org/10.30844/wgab_2023_7

outcomes. However, this may introduce bias and limit the solution quality to human performance. Another option is to use learning algorithms such as deep reinforcement learning, which could train on a digital model – similar to the one used for the GA and heuristic – before the actual planning, and then execute it faster with a learned strategy. This poses challenges in terms of explainability and responsibility (i.e. fairness, equity, inclusiveness), as well as dealing with the complexity of the search space.

References

Aickelin, Uwe; Dowsland, Kathryn A. (2004): An indirect Genetic Algorithm for a nurse-scheduling problem. In: Computers & Operations Research 31 (5), S. 761–778. DOI: 10.1016/S0305-0548(03)00034-0.

Akbarzadeh, Babak; Maenhout, Broos (2021): A decomposition-based heuristic procedure for the Medical Student Scheduling problem. In: European Journal of Operational Research 288 (1), S. 63–79. DOI: 10.1016/j.ejor.2020.05.042.

Bard, Jonathan F.; Shu, Zhichao; Leykum, Luci (2013): Monthly clinic assignments for internal medicine housestaff. In: IIE Transactions on Healthcare Systems Engineering 3 (4), S. 207–239. DOI: 10.1080/19488300.2013.857370.

Bard, Jonathan F.; Shu, Zhichao; Morrice, Douglas J.; Leykum, Luci K. (2016): Annual block scheduling for internal medicine residents with 4+1 templates. In: Journal of the Operational Research Society 67 (7), S. 911–927. DOI: 10.1057/jors.2015.109.

Bard, Jonathan F.; Shu, Zhichao; Morrice, Douglas J.; Leykum, Luci K. (2017): Constructing block schedules for internal medicine residents. In: IISE Transactions on Healthcare Systems Engineering 7 (1), S. 1–14. DOI: 10.1080/19488300.2016.1255284.

Beliën, Jeroen; Demeulemeester, Erik (2006): Scheduling trainees at a hospital department using a branch-and-price approach. In: European Journal of Operational Research 175 (1), S. 258–278. DOI: 10.1016/j.ejor.2005.04.028.

Beliën, Jeroen; Demeulemeester, Erik (2007): Building cyclic master surgery schedules with leveled resulting bed occupancy. In: European Journal of Operational Research 176 (2), S. 1185–1204. DOI: 10.1016/j.ejor.2005.06.063.

Brech, Claus-Henning; Ernst, Andreas; Kolisch, Rainer (2019): Scheduling medical residents' training at university hospitals. In: European Journal of Operational Research 274 (1), S. 253–266. DOI: 10.1016/j.ejor.2018.04.003.

Bundesministerium für Gesundheit (2022): Verordnung der Bundesministerin für Gesundheit über die Ausbildung zur Ärztin für Allgemeinmedizin/zum Arzt für Allgemeinmedizin/zur Fachärztin/zum Facharzt. Ärztinnen-/Ärzte-Ausbildungsordnung 2015 - ÄAO 2015, of the version dated 05.06.2022.

Deb, K.; Pratap, A.; Agarwal, S.; Meyarivan, T. (2002): A fast and elitist multiobjective genetic algorithm: NSGA-II. In: IEEE Transactions on Evolutionary Computation 6 (2), S. 182–197. DOI: 10.1109/4235.996017.

Denson, Joshua L.; McCarty, Matthew; Fang, Yixin; Uppal, Amit; Evans, Laura (2015): Increased Mortality Rates During Resident Handoff Periods and the Effect of ACGME Duty Hour Regulations. In: The American Journal of Medicine 128 (9), S. 994–1000. DOI: 10.1016/j.amjmed.2015.03.023.

Diponegoro, Ahmad; Rukman, Fajar (2017 - 2017): Modeling of a Scheduling Method for Organizing Training Assignments. In: Proceedings of AISTEEL 2017. 2nd Annual International Seminar on Transformative Education and Educational Leadership (AISTEEL 2017). Medan, Indonesia, 10/16/2017 - 10/17/2017. Paris, France: Atlantis.

Dummer, Wolfgang (2022): Ein metaheuristisches Optimierungsverfahren für Ausbildungspläne im Ärzt*innenausbildungsmanagement. Master Thesis. Technische Universität Wien. Online verfügbar unter https://repositum.tuwien.at/handle/20.500.12708/136043.

DOI: https://doi.org/10.30844/wgab_2023_7

Dummer, Wolfgang; Gaal, Alexander; Sobottka, Thomas; Ansari, Fazel (2023): A Genetic Algorithm Approach for Medical Resident Scheduling in Austria. In: Yuri Borgianni, Dominik T. Matt, Margherita Molinaro/Guido Orzes (Hg.): Towards a Smart, Resilient and Sustainable Industry. Proceedings of ISIEA 2023. 2nd ed. 2023. Cham: Springer (Lecture Notes in Networks and Systems). DOI:10.1007/978-3-031-38274-1_27

E. Zanazzo et al. (2022): Local Search Techniques for a Medical Student Scheduling Problem. University of Udine, Udine, Italy.

Franz, Lori S.; Miller, Janis L. (1993): Scheduling Medical Residents to Rotations: Solving the Large-Scale Multiperiod Staff Assignment Problem. In: Operations Research 41 (2), S. 269–279. DOI: 10.1287/opre.41.2.269.

Guo, Jiayi; Morrison, David R.; Jacobson, Sheldon H.; Jokela, Janet A. (2014): Complexity results for the basic residency scheduling problem. In: J Sched 17 (3), S. 211–223. DOI: 10.1007/s10951-013-0362-9.

ITO, Mari; ONISHI, Aino; SUZUKI, Atsuo; IMAMURA, Akira; ITO, Takuya (2018): The Resident Scheduling Problem. - A Case Study at Aichi Medical University Hospital -. In: Journal of Japan Industrial Management Association 68 (4E), S. 259–272. DOI: 10.11221/jima.68.259.

Jütte, Silke; Müller, Daniel; Thonemann, Ulrich W. (2017): Optimizing railway crew schedules with fairness preferences. In: J Sched 20 (1), S. 43–55. DOI: 10.1007/s10951-016-0499-4.

Kamhuber, Felix; Sobottka, Thomas; Heinzl, Bernhard; Henjes, Jan; Sihn, Wilfried (2020): An efficient hybrid multi-criteria optimization approach for rolling production smoothing of a European food manufacturer. In: Computers & Industrial Engineering 147, S. 106620. DOI: 10.1016/j.cie.2020.106620.

Kraul, Sebastian; Fügener, Andreas; Brunner, Jens O.; Blobner, Manfred (2019): A robust framework for task-related resident scheduling. In: European Journal of Operational Research 276 (2), S. 656–675. DOI: 10.1016/j.ejor.2019.01.034.

Maenhout, Broos; Vanhoucke, Mario (2011): An evolutionary approach for the nurse rerostering problem. In: Computers & Operations Research 38 (10), S. 1400–1411. DOI: 10.1016/j.cor.2010.12.012.

Proano, Ruben A.; Agarwal, Akshit (2018): Scheduling internal medicine resident rotations to ensure fairness and facilitate continuity of care. In: Health care management science 21 (4), S. 461–474. DOI: 10.1007/s10729-017-9403-9.

Rechnungshof Österreich (2021): Bericht des Rechnungshofes zur Ärzteausbildung. Reihe BUND 2021/42. Online verfügbar unter https://www.rechnungshof.gv.at/rh/home/home/A-rzteausbildung.pdf, last checked on 21.05.2023.

Ryan; S. Pruksaritanon; A. W. Reber; P. G. Aman; K. P. White; M. C. Smith; S. N. Hwang (2013): Preference based scheduling for medical residents and interns at the University of Virginia. In: 2013 IEEE Systems and Information Engineering Design Symposium. 2013 IEEE Systems and Information Engineering Design Symposium, S. 110–115.

Schleyer, Titus K. L. (1994): Temple Dental School Uses an Expert System to Schedule Students' Clinical Rotations. In: Interfaces 24 (5), S. 13–21. DOI: 10.1287/inte.24.5.13.

Smalley, Hannah K.; Keskinocak, Pinar (2016): Automated medical resident rotation and shift scheduling to ensure quality resident education and patient care. In: Health care management science 19 (1), S. 66–88. DOI: 10.1007/s10729-014-9289-8.

Stolletz, Raik; Brunner, Jens O. (2012): Fair optimization of fortnightly physician schedules with flexible shifts. In: European Journal of Operational Research 219 (3), S. 622–629. DOI: 10.1016/j.ejor.2011.10.038.