Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 58-19 (2024) 929–934

IFAC Papers Online
CONFERENCE PAPER ARCHIVE

A Novel Personnel Planning Method to Improve Operations Management: Transferring lessons learned from manufacturing to healthcare

Alexander Gaal* Wolfgang Dummer* Paul Lindorfer* Fazel Ansari*,**

* *Center for Sustainable Production and Logistics, Fraunhofer Austria Research GmbH, Theresianumgasse 7, Vienna, Austria (e-mail: alexander.gaal@fraunhofer.at)*

** *Research Group of Production and Maintenance Management, TU Vienna, Theresianumgasse 27, Vienna, Austria (e-mail: fazel.ansari@tuwien.ac.at)*

Abstract: There is a solid body of knowledge on personnel planning in production and logistics, showcasing potential applications across various sectors, particularly in operations management in healthcare. This paper focuses on Medical Residency Scheduling Problems (RSP) in a cross-facility context, employing a real dataset from an Austrian hospital group to assess the applicability of production planning and control (PPC) optimization techniques. The study examines approximate, expert-driven, and exact mixed-integer programming methods, underscoring the approximate method's effectiveness and rapidity in optimizing schedules against four objectives within a constrained period. The successful application of this novel method not only marks a significant advancement in scheduling systems but also demonstrates the potential for these methods to address broader scheduling challenges, significantly improving operational efficiency and quality. This approach offers insights for time-sensitive personnel planning, suggesting a versatile applicability of production-derived methods in healthcare scheduling.

Keywords: Medical Resident Scheduling, Hybrid Optimization Methods, Branch & Cut Algorithm, GAMS Modeling, Healthcare Workforce Management, Scheduling Method Comparison

1. INTRODUCTION

In healthcare management, medical resident scheduling is a multifaceted challenge, crucial for balancing educational needs, hospital operations, and legal requirements. Building on our prior research (Gaal et al. 2023), which proposed a novel approximate heuristic-metaheuristic approach, the is further refer to as hybrid approach, for addressing these challenges within the Austrian context. The research adapts optimization methods from production planning to medical residency scheduling, aiming to enhance decision-making processes across sectors. This interdisciplinary approach offers novel perspectives and underscores the broad applicability of optimization techniques in boosting operational efficiency in operations management. This paper advances our methodology by applying it to a real dataset from an Austrian hospital group. By comparing our approach with expert-generated schedules and an exact GAMS model, we aim to assess the efficiency and fairness of the developed scheduling method. This research evaluates the results based on a defined objective function and graphical comparisons, also considering the effort involved in schedule production. We structure this paper to provide a comprehensive literature review, problem description, solution methodology, performance validation, results comparison and conclude with key findings and directions for future research.

2. RELATED WORK

The Residency Scheduling Problem (RSP) garners attention due to its complexity, influenced by diverse stakeholders including legislation, hospitals, and medical residents, each with unique requirements and objectives. This complexity is heightened in Austria, where over 8,300 residents navigate training across 271 public hospitals (Ärztinnen-/Ärzte-Ausbildungsordnung 2015 of 6/5/2022; Rechnungshof Österreich 2021). The legal framework (Bundesministerium für Gesundheit 2015) and hospital-specific constraints necessitate a customized approach to RSP, differing significantly from models applicable in other regions (Kraul et al. 2019; Diponegoro and Rukman 10/16/2017 - 10/17/2017; Guo et al. 2014) Studies have explored various facets of RSP, from legislation impacts to hospital and resident requirements, highlighting the need for tailored solutions that accommodate specific program and individual preferences (Bard et al. 2016; Bard et al. 2017; Proano and Agarwal 2018; ITO et al. 2018; Brech et al. 2019). However, achieving an optimal scheduling solution that addresses all stakeholders' needs remains a challenge, with existing literature focusing on partial aspects of the problem or simplified models (Guo et al. 2014; Kraul et al. 2019).

Solution methodologies vary, ranging from heuristic approaches to exact mathematical programming, reflecting the problem's NP-hard nature (Guo et al. 2014; Bard et al. 2017). Notably, studies employing Dantzig-Wolfe decomposition and column generation as well as branch-and-price approaches address large-scale instances but often simplify constraints and objectives, limiting applicability to real-world, multifaceted RSP scenarios (Kraul et al. 2019; Beliën and Demeulemeester 2006; Beliën and Demeulemeester 2007). Recent works also acknowledge the importance of flexibility and fairness in scheduling, introducing objectives that strive for equitable

10.1016/j.ifacol.2024.09.162

training opportunities and accommodating resident preferences (ITO et al. 2018). Yet, a comprehensive approach that integrates the goals of all relevant stakeholders, including training quality and fairness, is lacking in current literature.

This paper builds upon the insights from (Akbarzadeh and Maenhout 2021a; Akbarzadeh and Maenhout 2021b), who also explored exact Branch-and-Price approaches as well as heuristic solutions in the Belgian context, and (Zanazzo et al. 2022), demonstrating the efficiency of metaheuristic methods like Simulated Annealing in solving RSP with reduced time. Later the authors presented a solution approach that combines local search techniques with a simulated annealing procedure (Zanazzo et al. 2024). This method was able to find optimal solutions for all instances of the dataset proposed by Akbarzadeh and Maenhout (2021) in significantly shorter runtimes. This research aims to fill the gap in the literature by proposing a method that not only considers the complex requirements of Austrian medical education but also sets a precedent for adaptable solutions applicable to other systems. Moreover, recent studies like (Seizinger and Brunner 2023) work on optimizing nursing curricula in German dual vocational schools, offer insights into tailored curriculum planning within healthcare education, underscoring the importance of structured approaches in addressing the sector's unique challenges. Similarly, (Kraul and Brunner 2023) focus on stabilizing medical resident scheduling through prioritized multiple training schedules, highlighting strategies to counteract operational uncertainties, a critical aspect of ensuring adaptability and continuity in medical education.

In summary, while the literature provides foundational knowledge and various approaches to RSP, this study seeks to advance this field by offering a method that prioritizes efficiency, fairness, and stakeholder satisfaction. This approach, validated against real-world data and compared with existing methodologies, aims to contribute particularly within the Austrian medical education context.

3. PROBLEM DEFINITION

Expanding on the Residency Scheduling Problem (RSP), this task entails the allocation of medical residents across multiple hospitals (cross-institutional scheduling) for up to 72 months, guided by the legislative framework AO2015 (Ärztinnen-/Ärzte-Ausbildungsordnung 2015). This scheduling is pivotal for ensuring a comprehensive medical education that spans basic and specialty stages, accommodating the diverse needs and preferences of stakeholders including legislation, hospitals, and residents. Legislation mandates that training be conducted in full months and adhere to a structured sequence, blending mandatory and elective elements. Hospitals grapple with capacity limitations, necessitating a balance between training demands and available duty positions, whereas residents are afforded flexibility regarding start times, disciplines, and the extent of employment.

Our prior work (Gaal et al. 2023) offers an in-depth examination of these intricacies, providing essential insight into the nuanced requirements of RSP within the Austrian healthcare context. For an elaborate discussion on these themes and their application to Austrian healthcare, our previous paper serves as a comprehensive resource, establishing the foundation for the enhancements proposed in this study. This research underscores the necessity for a scheduling solution that addresses the requirements of RSP while embedding principles of fairness and efficiency. This aligns with the perspective of (Stolletz and Brunner 2012), who articulate fairness in terms of equitable preference consideration across employees. Further, echoing (Abdalkareem et al. 2021) contemporary studies highlight the significance of adaptability and personalized education pathways in healthcare scheduling. Future RSP solutions are advised to favor customizable frameworks that adeptly cater to individual resident preferences, potentially elevating the overall caliber of medical training and promoting equitable educational opportunities.

4. METHOD ENHANCEMENT

In tackling the RSP, we face a scenario with complex variables and constraints reflecting the intricate nature of assigning medical residents. Drawing on the foundational principles outlined in our preceding work and leveraging the insights from (Akbarzadeh and Maenhout 2021a), our proposed hybrid approach merges heuristic and metaheuristic strategies, employing a greedy constructive heuristic in conjunction with a Genetic Algorithm (GA). This method distinctively generates multiple evolutionary-based solutions rather than iterating on an initial solution, ensuring adherence to the detailed constraints outlined in AO2015 and emulating real-world scheduling executed by human planners.

Crucially, the enhanced objective function now includes four additional metrics to thoroughly evaluate the results of planning and are listed in Table 1. These metrics were carefully developed to address the diverse interests of all stakeholders - legislators, hospitals, and residents.

Table 1. Metrics used in Objective Function

No.	Metric	Abr.
1	Month without training	MWT
2	Consecutive Month without training	CMWT
3	Department changes	DC
4	Number of departments per module	DPM
5	Departments without training	DWT
6	Hospital changes	HC
7	Months at cooperation partner	MACP
8	Single month assignments	SMA
9	Violated preferences of residents	VP
10	Variance of month without training	V-MWT
11	Variance of violated preferences	V-VP

By integrating these additional dimensions, the objective function becomes even more comprehensive, aiming to minimize negative impacts on training quality and resident satisfaction while fostering fairness and equity. For instance, CMWT and VMWT focus on reducing and evenly distributing the duration of gaps in training, whereas DPM and DWT aim to optimize the allocation of departments to training modules and minimize periods without departmental training, respectively. Metrics like MACP reflect the importance of external partnerships in resident training and highlighting the developed method's adaptability.

Our optimization technique, inspired by an indirect GA for a nurse-scheduling problem (Aickelin and Dowsland 2004) and adjusted for the RSP's requirements, utilizes a GA to determine the optimal sequence of scheduling. This approach, incorporating elements such as non-dominated sorting and tournament selection based on comprehensive objective function evaluations, prioritizes the attainment of a balanced and efficient scheduling solution. The algorithm's design, including its population size and genetic operators, has been fine-tuned to ensure an optimal balance between solution quality, computational efficiency, and diversity.

The operational framework of the method, depicted in Figure 1, commences with the generation of initial solutions through randomized sequences of training modules, guiding the scheduling order. This process, aimed at minimizing unnecessary transitions and optimizing educational continuity, iterates until meeting specified termination criteria, ensuring the derivation of optimal schedules within the expanded objective function's constraints. For an in-depth discussion on the augmented dimensions, constraints, and the specifics of the refined hybrid scheduling methodology, we refer readers to our latest paper (Gaal et al., 2023). To further validate the effectiveness of the proposed method, we intend to re-validate its performance by conducting a comparative analysis with both an expert generated schedule and an exact solution based on an identical dataset, providing a comprehensive evaluation of its capabilities.

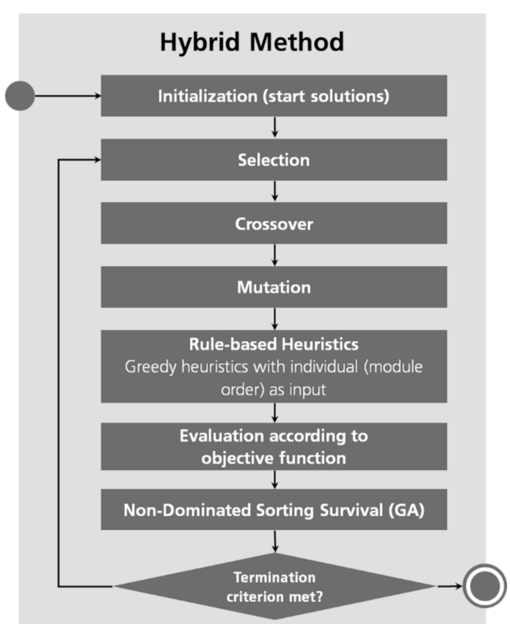

Figure 1. Hybrid Method Flowchart

5. VALIDATION APPROACH

Building upon the foundation laid in our previous paper, this validation marks a significant advancement by transitioning from a hypothetical test dataset to utilizing a real-world dataset. This shift not only enhances the realism and applicability of the findings but also allows to assess the practicality of the developed hybrid method under genuine operational conditions faced by hospitals.

Advancements in Validation Approach: Expanding from the initial comparison with expert generated schedule, this study incorporates a rigorous evaluation against both expert generated schedule and an exact GAMS-generated mathematical model. This broader framework enhances the method's demonstrated robustness and its ability to compete with advanced optimization techniques.

Operational Dataset Utilization: Employing a dataset from the day-to-day operations of an Austrian hospital group provides a tangible context for the method's effectiveness, introducing real-world complexities and challenges into the evaluation process.

Comparison with GAMS and CPLEX Optimization: The study includes a model that uses GAMS with the CPLEX solver. It is important to note that due to the premature termination of the solver, this model does not guarantee exact solutions and does not provide convergence certainty. This approach makes it possible to evaluate compliance with complex constraints and the efficiency of the method without the certainty of achieving optimal solutions.

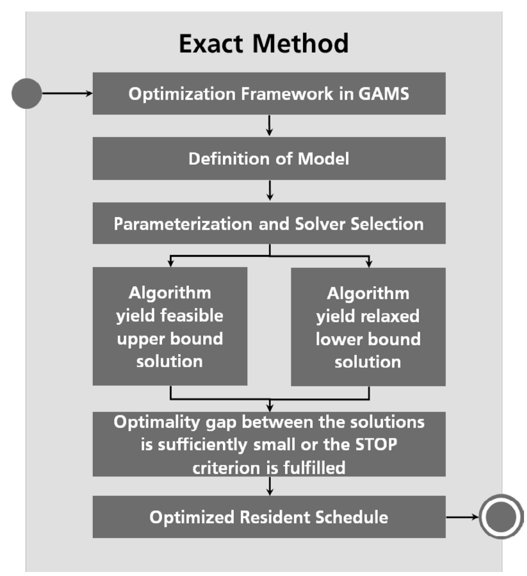

Figure 2. Exact Model Flowchart

Figure 2 showcases the exact optimization process within the GAMS framework through Mixed-Integer Programming (MIP). This method integrates both the hard constraints essential for resident training and the soft constraints reflecting objective values for the RSP. The CPLEX solver, renowned for its dual simplex algorithm, is utilized due to its capability to efficiently navigate and solve substantial MIP problems. The solver is tasked with finding solutions that provide feasible upper bounds and, through relaxation, lower bounds that serve as a comparative baseline. The primary aim is to minimize the gap between these bounds to achieve a solution that is not only feasible but optimal within the complex solution space defined by the residency scheduling problem.

Consistent Evaluation Criteria: Equally important is the fact that all schedules, irrespective of their method of generation, are evaluated against a uniformly applied objective function. This function serves as the common benchmark across all methods. The alignment of evaluation criteria across different methods guarantees that the comparative analysis accurately reflects the strengths and weaknesses of each approach in relation to the same set of goals and constraints.

Uniform Dataset and Objective Function Basis: It is crucial to underscore that all scheduling solutions in this study — whether derived from the hybrid method, the expert schedule, or the GAMS modeling — are formulated based on a singular, consistent dataset. This uniformity in data ensures that each method is subjected to the same set of operational variables and constraints, providing a level playing field for comparison.

Clarification of Evaluation Process: By utilizing the identical dataset and the objective function used for evaluation, we ensure that the comparative analysis truly measures the effectiveness of each scheduling method under identical conditions. This approach provides the foundation for drawing conclusions about the performance of compared methods.

Objective and Comprehensive Evaluation: The evaluation measures the quality of scheduling solutions, emphasizing the significance of using an operational dataset to assess the method's ability to address both the complexities of hospital scheduling and the optimization of designated constraints.

Computational Resources: The transition to a more powerful Standard_E16as_v4 (16 cores, 128GB RAM) virtual machine in the Azure Cloud was imperative for effectively solving the exact method, as the initial Standard_D4s_v3 setup (4 cores, 16GB RAM) proved inadequate. This upgrade significantly enhanced our computational capacity, aligning the efficiency of our calculations with the demands of real-time scheduling.

During the development and testing of the exact model, significant computational challenges emerged, particularly when optimizing across all 11 initially outlined objectives. The solver's limitations in processing time and memory usage, exacerbated by the model's complexity and numerous objectives, necessitated a strategic refinement of focus. Consequently, we selected four critical objectives: Months Without Training (MWT), Hospital Changes (HC), Department Changes (DC), and Single Month Assignments (SMA), which directly address the core challenges in residency scheduling, including minimizing training disruptions and optimizing logistical transitions. To ensure comparability, the evaluation of both the expert generated schedule and the hybrid approach concentrates on these four objectives. Although the hybrid method has the capability to optimize across all 11 objectives, focusing on these selected objectives allows for consistent comparisons. In real-world applications, optimizing selectively for specific objectives, rather than all simultaneously, allows for more tailored and practical scheduling solutions, reflecting the varied priorities within healthcare settings.

Mindful of the imperative for swift computational times demanded by stakeholders, we have imposed a 72-hour cap on the processing times for both the developed method and the exact model. This measure guarantees that our solutions are not only strategically optimized but also viable for timely implementation in healthcare environments.

6. RESULTS AND COMPARATIVE ANALYSIS

In this section, we delve into the empirical evaluation of the enhanced hybrid scheduling method, extending the analytical scope beyond the initial groundwork laid in our previous research. The implementation of the algorithm leverages the robust capabilities of Python and Rust, with computational tests conducted on a virtual machine in the Azure Cloud, reflecting a significant upgrade in processing power and efficiency. The provided dataset, grounded in the operational realities of an Austrian hospital group, involves scheduling 82 residents across 7 hospitals. This real-world dataset includes comprehensive details such as preferences, absence times, and specific constraints. This approach assesses the practical applicability of the developed method under conditions that closely mimic the challenges faced by hospital administrators.

Figure 3 offers a visual representation of the performance in the first 100 minutes of three scheduling methods over time. Each line depicts the objective value, calculated based on four objectives, plotted against computation time in seconds. This illustrates how each algorithm progresses toward an optimized schedule, highlighting their convergence behavior over time. Given the NP-hard complexity class of the problem, we allocated a total of 72 hours for the exact method's computation, understanding that such a duration would be impractical and unfeasible in a real-world application scenario.

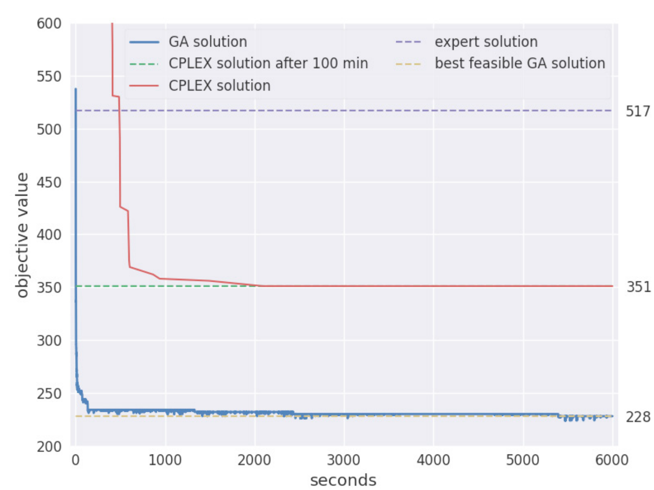

Figure 3. Results by Method after 100-min runtime

The hybrid method shows a rapid decline in the objective value, indicating quick convergence to a lower objective value which then plateaus, suggesting that the GA quickly finds a near-optimal solution and maintains it throughout the remainder of the run time. The CPLEX solution, on the other hand, demonstrates a gentler decline, reflecting its ability to find highly optimized solutions, if not constrained by computational time or resources. Notably, the CPLEX solution after 100 minutes (dashed line, 351) and underscores the computational intensity of the exact method, which may not be practical for real-time applications due to its slower convergence. This also highlights the practical limitations of this exact method in scenarios where quick decision-making is crucial. The expert solution, created over an 8-hour period by a human planner is a single, static solution. This solution serves as a benchmark for the other methods, illustrating the level of optimization that human planners can achieve without computational assistance. Table 2 provides a detailed comparison of the hybrid, exact, and expert methods against four key objectives, alongside their respective weightings. The hybrid meta-heuristic method showcases a substantial advantage in achieving the lowest objective value of 228, indicating the most optimized scheduling outcome. This demonstrates the method's capacity to reach high quality solutions with low computational effort, making it a viable option for complex scheduling tasks. The exact approach presents a middle ground, with a moderate objective value of 245 and a longer runtime compared to the hybrid method. In contrast, the expert method, relying on human expertise, results in the highest objective value of 517. This underscores the limitations of manual scheduling in achieving the level of optimization possible through computational methods. It is worth mentioning that the developed method found a solution with an objective value of 234 after less than 2 minutes. The exact method closely approaches the solution quality of the hybrid method. However, it requires a significantly longer processing time.

Table 2. Results by Method after 72-hour runtime

Objective	Hybrid	Exact	Expert	Weight
MWT	0	6	12	3
DC	107	111	162	1
HC	64	68	184	2
SMA	57	60	159	3
SUM	**228**	**245**	**517**	
Duration [min]	43	4200	480	

Overall, these results highlight the trade-offs between the time required and the quality of planning solutions. The hybrid method shows a balanced performance across all objectives, leading to the lowest overall result score, indicative of its efficiency. The exact method, while outperforming the expert approach, falls short against the hybrid method, especially in MWT metrics. However, it is conceivable that the exact method will find an even better solution if it is given enough time (>72h). The expert method, despite its higher scores, reflects the challenges manual scheduling faces in optimizing complex criteria. Based on the analysis, the key findings from the comparison can be summarized as follows:

1. **Efficiency and Optimality:** The hybrid method stands out for its ability to quickly identify near-optimal schedules, surpassing the other methods in terms of both the quality of the solution and runtime.

2. **Computational Intensity:** The exact method, while capable of delivering highly optimized solutions, exhibits significant computational demands in terms of time, even when the objective scope is narrowed. The hybrid solution demonstrates its efficiency by quickly finding near-optimal solutions even on less powerful computers, showcasing its practicality for real-world application scenarios.

3. **Methodological Synergy:** Hybrid method demonstrates a synergistic blend of a greedy constructive heuristic in conjunction with a GA, striking a balance that enhances both operational feasibility and solution quality. The hybrid solution's flexibility and efficiency suggest its applicability extends beyond the specific case studied.

These findings suggest that the hybrid method offers a viable and effective alternative for residency scheduling, particularly in settings where time efficiency is as crucial as the quality of the scheduling outcome. The comparative analysis underscores the hybrid method's efficiency in addressing RSPs, affirming its potential to significantly streamline healthcare operations. Not only does this method offer a more refined solution by effectively balancing multiple objectives, but it also adapts to the complexities and dynamic nature of hospital environments. These results, demonstrating a clear preference for the hybrid approach over traditional expert and exact methods, pave the way for future advancements in healthcare scheduling. This study's findings highlight the role of innovative scheduling solutions in optimizing healthcare delivery and medical education, promising significant contributions to operational efficiency and education quality.

7. DISCUSSION AND FURTHER DIRECTIONS

This study's evaluation was conducted with four objectives and a 72-hour runtime limit, balancing thoroughness with practicality. The developed hybrid method for medical residency scheduling integrates multiple objectives effectively, rapidly converging to near-optimal solutions. It significantly outperforms expert methods and offers competitive results against GAMS-modeled CPLEX solutions within practical time constraints. The method's flexibility to prioritize various objectives makes it suitable for real-world scenarios with diverse operational demands, and it shows promise for other planning problems in production and logistics. Despite the effectiveness of our hybrid method, several challenges remain. The high computational demands, especially for large-scale instances, and the complexity of fine-tuning the GA parameters (e.g. mutation operator) are notable limitations.

While our method shows promising results, a more detailed comparison with existing approaches in the literature is essential. Studies such as those by Akbarzadeh and Maenhout (2021), who explored branch-and-price methods and heuristic solutions, and recent works like those of Zanazzo et al. (2024) utilizing simulated annealing, provide valuable benchmarks. The data set provided by the authors could be used to compare methods further. Presented method for the Austrian education regulations would have to be adapted to the Belgian system. Integrating these comparisons will highlight the relative strengths and limitations of our approach.

Future research could explore the application of the hybrid method to other areas of scheduling, such as nursing rosters or surgical planning, where similar complexities and multi-objective optimization needs are present. Research into time-critical applications in this area appears to be promising. Investigating the method's adaptability to real-time changes, such as emergency cases or sudden staff shortages, could significantly enhance the responsiveness of healthcare scheduling systems. By continuing to refine and expand upon the hybrid method's capabilities, future research can forge paths toward more resilient, adaptive, and efficient healthcare scheduling systems, ultimately contributing to improved patient care and medical education.

8. REFERENCES

Abdalkareem, Zahraa A./Amir, Amiza/Al-Betar, Mohammed Azmi/Ekhan, Phaklen/Hammouri, Abdelaziz I. (2021). Healthcare scheduling in optimization context: a review. Health and Technology 11 (3), 445–469. https://doi.org/10.1007/s12553-021-00547-5.

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

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

Akbarzadeh, Babak/Maenhout, Broos (2021b). An exact branch-and-price approach for the medical student scheduling problem. Computers & Operations Research 129, 105209. https://doi.org/10.1016/j.cor.2021.105209.

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

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

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

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

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

Bundesministerium für Gesundheit (2015). Verordnung der Bundesministerin für Gesundheit über die Ausbildung zur Ärztin für Allgemeinmedizin/zum Arzt für Allgemeinmedizin und zur Fachärztin/zum Facharzt 2015. Available online at https://www.ris.bka.gv.at/GeltendeFassung.wxe?Abfrage=Bundesnormen&Gesetzesnummer=20009186 (accessed 4/22/2024).

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

Gaal, Alexander/Dummer, Wolfgang/Sobottka, Thomas/Ansari, Fazel/Schlund, Sebastian (2023). Advancing Medical Resident Scheduling. In: Jochen Deuse (Ed.). How can industrial management contribute to a brighter future? GITO mbH Verlag, 115–134.

Guo, Jiayi/Morrison, David R./Jacobson, Sheldon H./Jokela, Janet A. (2014). Complexity results for the basic residency scheduling problem. Journal of Scheduling 17 (3), 211–223. https://doi.org/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 -. Journal of Japan Industrial Management Association 68 (4E), 259–272. https://doi.org/10.11221/jima.68.259.

Kraul, Sebastian/Brunner, Jens O. (2023). Stable annual scheduling of medical residents using prioritized multiple training schedules to combat operational uncertainty. European Journal of Operational Research. https://doi.org/10.1016/j.

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

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

Rechnungshof Österreich (2021). Bericht des Rechnungshofes zur Ärzteausbildung. Reihe BUND 2021/42. Available online at https://www.rechnungshof.gv.at/rh/home/home/A-rzteausbildung.pdf (accessed 5/21/2023).

Seizinger, Markus/Brunner, Jens O. (2023). Optimized planning of nursing curricula in dual vocational schools focusing on the German health care system. European Journal of Operational Research 304 (3), 1223–1241. https://doi.org/10.1016/j.ejor.2022.05.007.

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

Zanazzo, Eugenia/Ceschia, Sara/Dovier, Agostino/Schaerf, Andrea (2024). Solving the medical student scheduling problem using simulated annealing. Journal of Scheduling. https://doi.org/10.1007/s10951-024-00806-z.

Zanazzo, Eugenia/Ceschia, Sara/Schaerf, Andrea (2022). Local Search Techniques for a Medical Student Scheduling Problem. In: Patrick De Causmaecker, Ender Özcan, Greet Vanden Berghe (Ed.). Proceedings of the 13th International Conference on the Practice and Theory of Automated Timetabling - PATAT 2022 - Volume III, 84–87.