AI-DRIVEN SHIFT SCHEDULING: INSIGHTS FROM A PILOT IN SAFRA CHILDREN'S HOSPITAL

David Gilad MD MhD[1],
Tzofnat Farbstein-Aljanati MD[1],
Arnon Afek MD MHA[2],
Itai M. Pessach MD PhD[1],
and Moshe Ashkenazi MD MBA[1]

[1]Safra Children's Hospital, Sheba Medical Center, Tel Hashomer, Israel
[2]Recanati School of Medicine, Reichman University, Herzliya, Israel

TO THE EDITOR:

The escalating complexity of healthcare delivery underscores the demand for digital transformation in medicine. Among impactful innovations, artificial intelligence (AI)-driven scheduling solutions are being developed and deployed [1].

Safra Children's Hospital at Sheba Medical Center is Israel's first medical institution to implement an AI-powered scheduling platform (Equina Scheduling: https://www.equinascheduling.com/). Equina was selected for its capacity to handle the complex, rule-based environment of a training program in a tertiary pediatric hospital. Comprising approximately 70 residents, our pediatric residency program is among the nation's largest. The program manages approximately 1500 monthly shifts, encompassing 49 distinct service shifts, 22 diverse training shifts, and more than 20 different block rotations.

With no constraints, the search-space for such a schedule is in the order of 10^{2768}. Even with a system of approximate constraints, the search-space remains in the order of 10^{570}, easily dwarfing the chess game possibility tree ($\sim 10^{123}$) by many orders of magnitude. Scheduling at this scale exceeds human capability and is even unfeasible to brute-force compute.

The traditional process of healthcare scheduling, particularly in a large academic hospital, is a major contributor to administrative fatigue and remains error prone. Residency scheduling is a high-dimensional optimization problem characterized by dynamic hospital needs, personal requests, regulatory constraints, evolving resident skills, and frequent last-minute changes. The resulting schedule is often a fragile construct, difficult to modify and prone to inequities in workload distribution, further impacting morale and contributing to burnout.

Our previous software solution for scheduling required extensive manual effort of approximately 50 hours from two chief residents each month. This time demand arose from the simultaneous planning of future schedules while managing requests and replanning the concurrently published active schedule. Our preliminary findings indicated that chief residents processed a monthly average of 509 requests for leave and shift off requests. In addition, they processed an average of 181 shift changes per month for the published schedule. This administrative quagmire detracted from the primary missions of medical training: delivering the best patient care and resident education.

The foundation of an effective AI-driven scheduling system requires capturing expert scheduler knowledge into a hierarchical ruleset. This ruleset comprises hard constraints like mandatory conditions that must be satisfied (e.g., work the morning after a night on-call shift is prohibited) and soft constraints, which represent preferential conditions (e.g., reducing unwanted shift patterns such as consecutive night shifts, balance rules). The system optimizes schedule permutations by assigning numerical rewards or penalties to these constraints [Figure 1]. While generating a possible schedule that adheres to all hard constraints is straightforward, achieving a truly optimized and balanced schedule depends on sophisticated trade-offs among the soft constraints. In the context of a residency program, this factor is particularly critical.

The meticulous integration of annual block rotation planning with monthly service schedules is essential to reconcile residents' educational requirements with the clinical workforce needs of the hospital. Consequently, an effective AI-scheduling solution must navigate the overlapping and often contradictory sets of constraints imposed by both the yearly block schedule and the immediate monthly roster. While software will never be a magic-bullet solution and still requires fine-tuning with a human scheduler in the loop, this reclaimed time could be used for better management of the residency program, focusing on long-term planning, advocating for and supporting residents.

Implementation of the scheduling agent involved a structured pilot, including a preliminary setup with dry run for quality assurance for 3 months. This pilot was followed by a full launch with rapid deployment and onboarding for all staff within 2 weeks. The administrative workload of the chief residents and program leadership was drastically reduced after implementation, allowing for more transparent schedules that were published with and earlier lead time.

Figure 1. Example artificial intelligence-scheduling demonstrating rule scores over time

Scores for hard (red), medium (yellow), and soft (blue) constraints are shown over elapsed solver time (x-axis). Hard-rule violations are eliminated within seconds and remain at near-zero throughout. Medium-rule scores, representing negative penalties for unassigned shifts, improved over time as required shifts were staffed, then stabilized. Soft-rule score improved gradually throughout the run from a net large starting negative penalty to more positive values reflecting progressive satisfaction of scheduling preferences.
Left y-axis: hard and medium scores; right y-axis: soft score.
Collectively, the optimizer first ensured feasibility, then shifted coverage, and finally optimized for preferences.

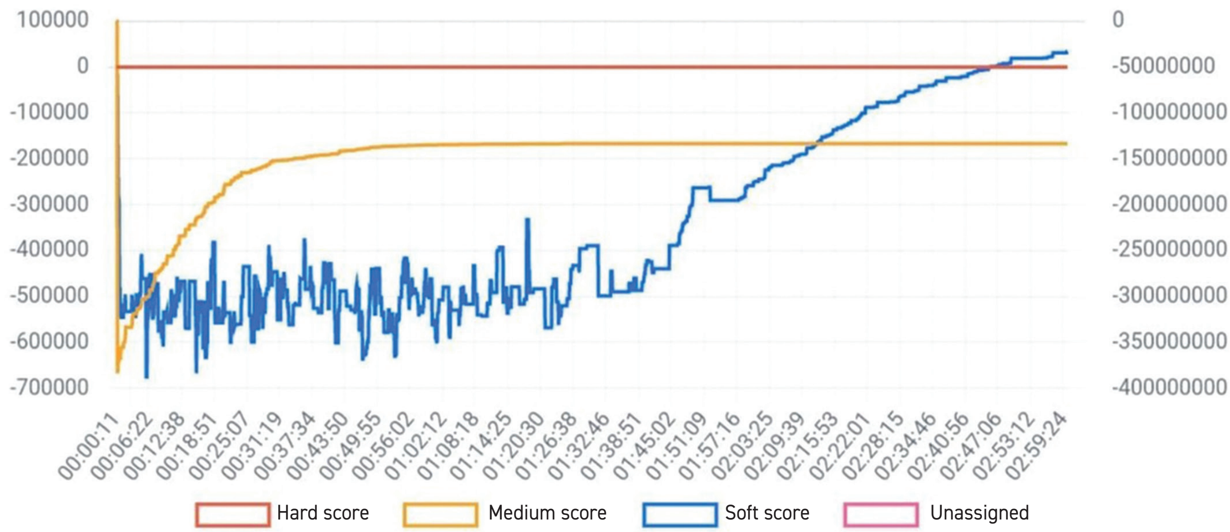

Successful implementation necessitated a scheduler's adaptability to workflow modifications, embracing new software with an aptitude for programmatic reasoning. From an administrative standpoint, the project was granted priority, financial backing, and expedited institutional IT and cybersecurity clearances, thereby mitigating potential obstacles associated with off-site cloud-based software.

Beyond operational improvements, AI-scheduling served as a powerful in-silico scheduling laboratory. By simulating schedules based on expected constraints and optimization goals, administrators could determine whether scheduling challenges arose from suboptimal rule-building or human resource limitations. Together with a system of dynamic stat tables and reports, the AI-scheduling solution provided actionable strategic insights that facilitated targeted, data-driven decision-making and intervention plans.

We continue to collaborate with Equina as a clinical innovation partner and testing site. This synergy leverages our frontline operational expertise to guide development to address real-world clinical needs. Ongoing collaboration aims to expand to other departments and provider sectors, including attending physicians and nursing teams, and to integrate language model-driven smart AI agents for no-code automation of scheduling tasks that could not be achieved by constraint optimization, leveraging clinical insights for real-world applicability.

CONCLUSIONS

The adoption of AI-scheduling is not a matter of mere convenience. It is an essential component of building a sustainable, efficient, and humane schedule that may mitigate burnout [2-4]. The successful deployment of AI-scheduling software in our pediatric residency program serves as a powerful case study. It demonstrates that by automating complex but deterministic tasks, we can significantly improve physician work–life balance, operational agility, and, most critically, preserve the finite resource of physician time for its highest and best use: the direct care of patients.

Correspondence

Dr. D. Gilad
Safra Children's Hospital, Sheba Medical Center, Tel Hashomer 52621, Israel
Email: dudugilad@gmail.com

References

1. Sumrall WD 3rd, Oury JV, Gilly GM. Enhancing physician satisfaction and patient safety through an artificial intelligence-driven scheduling system in anesthesiology. *Ochsner J* 2025; 25 (1): 44-9.
2. Underdahl L, Ditri M, Duthely LM. Physician burnout: evidence-based roadmaps to prioritizing and supporting personal wellbeing. *J Healthc Leadersh* 2024; 16: 15–27.
3. Busireddy KR, Miller JA, Ellison K, Ren V, Qayyum R, Panda M. Efficacy of interventions to reduce resident physician burnout: a systematic review. *J Grad Med Educ* 2017; 9 (3): 294–301.
4. Abbas A, Toor J, Margalik DA, et al. Machine learning to improve resident scheduling: harnessing artificial intelligence to enhance resident wellness. *J Reg Med Campuses* 2025; 8 (1). [Available from https://pubs.lib.umn.edu/index.php/jrmc/article/view/6332].

Capsule

Low-dose aspirin for PI3K-altered localized colorectal cancer

Alterations in PI3K pathway genes were detected in 1103 of 2980 patients (37.0%) with complete genomic data. Of 515 patients with group A alterations and 588 patients with group B alterations, 314 and 312, respectively, were assigned to receive aspirin or placebo. The estimated 3-year cumulative incidence of recurrence was 7.7% with aspirin and 14.1% with placebo (hazard ratio [HR] 0.49; 95% confidence interval [95%CI] 0.24–0.98; *P* = 0.04) among patients with group A alterations and 7.7% and 16.8%, respectively (HR 0.42; 95%CI 0.21–0.83), among those with group B alterations. The estimated 3-year disease-free survival was 88.5% with aspirin and 81.4% with placebo (HR 0.61; 95%CI 0.34–1.08) among patients with group A alterations and 89.1% and 78.7%, respectively (HR 0.51; 95%CI 0.29–0.88), among those with group B alterations. Severe adverse events occurred in 16.8% of aspirin recipients and 11.6% of placebo recipients. The authors **Martling** et al. concluded that aspirin led to a significantly lower incidence of colorectal cancer recurrence than placebo among patients with *PIK3CA* hotspot mutations in exon 9 or 20 and appeared to have a similar benefit among those with other somatic alterations in PI3K pathway genes.

N Engl J Med 2025; 393: 1051

Eitan Israeli

Capsule

Assembly of prefibrillar amyloids

The amyloid hypothesis is one of the leading explanations for the cause of Alzheimer's disease. It proposes that the accumulation of amyloid-β (Aβ) peptide oligomers and fibrils in the brain is the primary initiating event that triggers a cascade of neurodegenerative processes. **Liang** et al. conducted time-resolved imaging of the prefibrillar assembly of Aβ using a combination of transmission electron microscopy, cryo-electron tomography, atomic force microscopy, and single-particle cryo-electron microscopy. The collection of recorded images proposes a universal mechanism of assembly in which all observed structures, from oligomers to curvilinear protofibrils and ring-shaped annular assemblies, represent a continuum. Patch-clamp conductance measurements suggested that the observed annular structures form channels that span lipid membranes and increase their ion permeability.

Sci Adv 2025; 11 (35): eadx9030

Eitan Israeli

Capsule

Localized pain relief

The brain mechanisms underlying pain modulation remain to be fully elucidated. **Crawford** and colleagues used functional magnetic resonance imaging in humans to explore neural activity in the periaqueductal gray brain area, a key region involved in pain processing, and tested the organizational principle of placebo analgesia. The authors exposed participants to heat stimuli of varying intensities on the face, forearm, and leg with or without the application of a placebo analgesic cream and showed that placebo analgesic response was somatotopically (spatially) organized. The results provide insights into how the brain can enable precise and localized pain relief to optimize defensive responses.

Science 2025;389: 893

Eitan Israeli