INFORMS Journal on Applied Analytics

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Using Optimization Techniques and Multidisciplinary Collaboration to Solve a Challenging Real-World Residency Scheduling Problem

Young-Chae Hong, Amy Cohn, Stephen Gorga, Edmond O'Brien, William Pozehl, Jennifer Zank

informs
http://pubsonline.informs.org/journal/inte/

INFORMS JOURNAL ON APPLIED ANALYTICS
Vol. 49, No. 3, May–June 2019, pp. 201–212
ISSN 0092-2102 (print), ISSN 1526-551X (online)

Using Optimization Techniques and Multidisciplinary Collaboration to Solve a Challenging Real-World Residency Scheduling Problem

Young-Chae Hong,[a] Amy Cohn,[a] Stephen Gorga,[b] Edmond O'Brien,[c] William Pozehl,[a] Jennifer Zank[d]

[a] Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, Michigan 48109; [b] Department of Pediatrics, Division of Pediatric Critical Care Medicine, University of Michigan, Ann Arbor, Michigan 48109; [c] Department of Anesthesiology, University of Michigan, Ann Arbor, Michigan 48109; [d] Division of Newborn Medicine, University of Pittsburgh Medical School, Pittsburgh, Pennsylvania 15213

Contact: hongyc@umich.edu, http://orcid.org/0000-0001-6330-6101 (YCH); amycohn@med.umich.edu, http://orcid.org/0000-0002-8612-6098 (AC); smgorga@med.umich.edu, http://orcid.org/0000-0001-8193-7679 (SG); obriene@med.umich.edu, http://orcid.org/0000-0002-3242-0478 (EO); pozewil@umich.edu, http://orcid.org/0000-0003-4010-6195 (WP); zankjs@upmc.edu, http://orcid.org/0000-0001-5712-4627 (JZ)

Received: September 6, 2017
Revised: June 10, 2018; January 9, 2019
Accepted: March 11, 2019
Published Online in Articles in Advance: May 15, 2019

https://doi.org/10.1287/inte.2019.0995

Abstract. In this paper, we consider the problem of assigning medical residents to shifts within a pediatric emergency department. This problem is challenging to solve for a number of reasons. First, as with many other healthcare personnel scheduling problems, it has a nonhomogeneous workforce comprising residents of different characteristics, requirements, and capabilities. Second, residency scheduling problems not only must ensure adequate resources for patient care but also must meet educational training needs, adding further complexity and constraints. Finally, rather than being evaluated under a single cost metric, resident scheduling problems have multiple objective criteria, which are often in conflict with each other. Some of these challenges can be overcome through the use of operations research techniques; others depend on the process by which we apply these techniques and, in particular, the way that the operations researchers collaborate with the clinicians. We present our experience at the University of Michigan C.S. Mott Children's Hospital in building monthly schedules, focusing on both our integer programming formulation and the iterative, interactive approach in which we use this integer program as a tool within the broader process of schedule development.

History: This paper was refereed.

Keywords: residency scheduling • shift scheduling • integer programming • emergency department

Introduction

In this paper, we consider the problem of assigning medical residents to shifts within a pediatric emergency department (ED). We present our collaborative work with the University of Michigan C.S. Mott Children's Hospital (Mott) in building its monthly schedules, focusing on both our integer programming (IP) formulation and the iterative, interactive approach in which we use this integer program as a tool within the broader process of schedule development.

Healthcare personnel scheduling problems are often challenging to solve for many reasons, including heterogeneity of the workforce, personal preferences, fatigue issues, and concerns over continuity of care (Topaloglu 2006, 2009; Schwenk et al. 2010; Goldman et al. 2015). When scheduling residents, we are faced with all of these challenges and more, including significant diversity in the residents' skill sets and levels of training as well as the need not only to provide adequate patient care, but also to ensure that the residents meet their educational requirements. In addition, residents are bound by the regulatory requirements of the Accreditation Council for Graduate Medical Education, which governs rest rules and many other aspects of residents' training.

Complex residency scheduling problems are frequently solved by chief residents, either manually or with the limited help of some basic spreadsheet tools (as was the case at Mott prior to our collaboration). This is not only a tedious and time-consuming process, but it is often difficult to find even a feasible schedule, fully satisfying all rules and requirements, let alone one that is of high quality with respect to patient care, educational requirements, and personal preferences.

A natural approach to solving complex combinatorial problems is through *integer programming*. This is often a successful way to find *feasible solutions* to residency shift scheduling problems, such as the one that we faced at Mott, as is demonstrated by an interactive tool that we have developed and posted at https://cheps.engin.umich.edu/tools/shift-scheduling-game.

On the other hand, not all feasible solutions to this problem are equally acceptable. What does it mean to find an optimal solution? Unlike many industrial scheduling applications, the goal of a residency scheduling problem is not to minimize cost. Rather, there are many different criteria associated with quality of patient care, resident educational needs, and personal preferences under which a schedule is evaluated. Furthermore, these criteria may be nonlinear or qualitative and may vary from month to month.

The goal of the chief residents is not to find an "optimal" schedule solution. It quickly became clear to us in the early stages of our collaboration that such a thing does not exist; given two high-quality schedules, the chief residents often could not definitely tell us which one was better with each having characteristics that they liked and possibly a few flaws that they were willing to accept. The ultimate choice between the two was often quite arbitrary. Rather than an optimal solution, their desire is to find a feasible solution that can be found quickly, satisfies a substantial portion of the resident preferences, and is perceived by the residents as reasonably equitable.

Coming to accept this fact—that finding an optimal solution, as we had been trained as operations researchers to do, was not actually the right goal for the end user of our research—substantially changed the way we proceeded with solving the problem and ultimately led to a successful outcome: our collaborative approach has now been used for several years to build monthly schedules that are implemented operationally at Mott. The goal of this paper is to share our experience in incorporating integer-programming techniques within a collaborative, interactive approach to find solutions that are valid and useful to the end user.

Problem Statement

Residency is the phase of graduate medical education after completing medical school while continuing training to become a fully independent practicing physician. Residents are trained in progressively more specialized areas under the supervision of more experienced attending physicians, becoming increasingly more independent as they advance in seniority. In the United States, medical students enter residency through the National Resident Matching Program (NRMP). The "match" happens annually in the spring with medical students matched to residency programs across the country. In the United States, more than 30,000 residency positions are filled by NRMP annually. Pediatric residents make up almost 10% of this population.

Residency programs range in duration from three to seven years, depending on the specialty chosen. For example, pediatric residency and internal medicine residency both take three years. General or orthopedic surgery can take five years to complete. After completion of a residency program, one to three years of additional fellowship in a subspecialty may be chosen by some physicians for further training. For example, in pediatrics, a pediatrician may train to become a neonatologist. In internal medicine, a physician may subspecialize in cardiology.

Because they provide patient care while being trained, residents are both providers and learners, and thus, *residency scheduling* problems must satisfy both patient care needs and educational requirements. In addition, the quality of a schedule can have significant personal impact on residents. Poor-quality schedules can lead to fatigue, lack of sleep, professional burnout, and even depression. In addition, fatigue can lead to medical errors and negatively impact patient care (Shanafelt et al. 2002, Shanafelt and Habermann 2002, Lockley et al. 2007, Rogers 2008, Schwenk et al. 2010, Sen et al. 2010, Goldman et al. 2015).

Residency scheduling problems are largely divided into three types: *Block schedules* separate an academic year into blocks of time, typically on the order of two weeks, four weeks, or a calendar month, in which residents are assigned to specific services (for example, the ICU, an outpatient clinic, or the emergency department). The purpose of a block schedule is to rotate residents through many different clinical experiences in diverse fields (Franz and Miller 1993). *Call schedules* define periods of time over which residents are responsible for being "on call" outside of normal working hours (i.e., evenings, weekends) to meet patient needs. This on-call duty is both critical for patient care and required for residents as additional training (Ozkarahan 1994, Cohn et al. 2009). *Shift schedules* assign residents to specific tasks at specific times within a given block as defined by the residents' block rotation (Beaulieu et al. 2000, Carter and Lapierre 2001, Sherali et al. 2002, Guler 2013). In this paper, we focus on a version of the shift scheduling problem for a pediatric emergency department.

The Pediatric Emergency Department Scheduling (PEDS) Problem

The University of Michigan Pediatric Emergency Department provides services 24 hours a day, seven days a week to care for children with medical problems that cannot wait or are too severe to be seen by their primary care providers. The ED offers unscheduled care ranging from common minor pediatric problems to major medical and traumatic emergencies. The University of Michigan Pediatric Emergency Department sees approximately 20,000 children per year. It is designated as a Level I Pediatric Trauma Center, the highest designation, which means it is certified by the

American College of Surgeons Committee on Trauma to provide care to the most severely ill or injured children. In addition to attending physicians and other clinical personnel, the ED is staffed by medical residents of varying levels of seniority and from a variety of training programs (e.g., pediatrics, internal medicine, psychiatry) with roughly 22 *interns* (i.e., first-year residents) joining the program each year while more senior residents remain for additional years of training.

There are seven overlapping nine-hour shifts that are scheduled every day to staff the Pediatric Emergency Department: shift 1 (7 a.m. to 4 p.m.), shift 2 (9 a.m. to 6 p.m.), shift 3 (12 p.m. to 9 p.m.), shift 4 (4 p.m. to 1 a.m.), shift 5 (5 p.m. to 2 a.m.), shift 6 (8 p.m. to 5 a.m.), and shift 7 (11 p.m. to 8 a.m.). Shifts 1 and 2 are considered to be *morning shifts*, shifts 4 and 5 are *day shifts*, shifts 6 and 7 are *overnight shifts*, and shift 3 is considered to be a "flex" shift. The flex shift should ideally be staffed by a resident to provide additional support to the attending physicians during the peak hours of the day, but staffing is not required. All other shifts must be staffed by exactly one resident.

Residents are typically assigned to work in the ED for either half- or full-month rotations. The senior residents start and end on the first and last days of the calendar month, and the interns transition from one rotation to another on the 27th of the preceding month to ensure a smooth transition. From a scheduling standpoint, the implication of this is that certain shifts at the end of the month (known as "optional shifts") can be left unfilled to be staffed with incoming interns at the scheduling of the following month.

Note, however, that not all shifts can be staffed by interns. Specifically, senior residents must staff the ED for the shifts that have the fewest overlaps (i.e., shifts 1 and 7). In addition, there are a number of residents who are rotating through the ED from other residency programs besides pediatrics. In a given month, there are four to six residents from family medicine or emergency medicine who rotate through the pediatric emergency department for additional training. These residents also have limitations on when they can work, typically as a function of other regularly scheduled training activities.

Finally, in addition to their monthly rotations on different services (including the emergency department), residents also maintain a panel of patients for whom they care throughout the year in the form of *continuity clinics*. Ideally, residents' ED shifts should not conflict with attending their weekly clinics.

Feasibility Requirements

The primary PEDS rules can be summarized as follows:

- Coverage: All shifts must be staffed by exactly one resident except the flex shifts (which can be left unstaffed if necessary) and the optional shifts at the end of the month (which can wait to be staffed by interns coming in on the following month).
- Invalid assignments: Reasons why a specific resident cannot be assigned to a specific shift include seniority (i.e., interns cannot staff certain shifts), program conflicts (e.g., certain training programs have educational commitments on certain days of the week), and continuity clinics. In addition, we treat certain requests for time off as hard constraints. Personal requests are typically only for critical issues (interviews, certification boards, weddings and other major family events), and as such, the chiefs put the highest priority on satisfying these requests. There is typically enough flexibility in the assignment of shifts that such requests can be granted.
- Preassignment: In some cases, there are specific shifts that *must* be assigned to specific residents. For example, these might be to ensure the completion of educational responsibilities for residents who have been on leave and only have a few remaining shifts to cover.
- Pediatric coverage: For certain pairs of overlapping shifts, it is necessary to ensure that at least one of the two shifts is staffed by a resident from the pediatrics program.
- Duty-hours restrictions: Regulatory guidelines require that, after completion of a shift, residents must have at least 10 hours off duty before beginning a subsequent shift.
- Limits on consecutive days on shift: There are limits on the number of consecutive days in a row that residents can be assigned to shifts and tighter limits on the number of consecutive days in a row that residents can be assigned to overnight shifts.

This is just a representative sample of the requirements that must be enforced to ensure a valid schedule; additionally, each month brings unique "one-off" requirements that must be satisfied. The mathematical model in the appendix defines the feasible set of resident shift schedules relative to the rules outlined here.

Literature Review

Scheduling, the "allocation of resources to tasks over given time periods," is a decision-making process that occurs throughout almost all manufacturing and services industries (Pinedo 2016, p. 1). Examples include manufacturing processes (Graves 1981); transportation systems, such as airlines (Vance et al. 1997), railways (Ernst et al. 2001), and public transportation (Raff 1983); staff scheduling in call centers (Gans et al. 2003); emergency services, such as police (Taylor and Huxley 1989), ambulance (Li and Kozan 2009), and fire departments (Fry et al. 2006); and even toll booths (Edie 1954).

Scheduling problems are solved by using a variety of solution approaches, from heuristic algorithms to exact methods, based on problem-specific characteristics (e.g., cyclic versus noncyclic, deterministic versus stochastic, single machine versus parallel machine, and so on). The review papers (Brucker et al. 2011, Van Den Bergh et al. 2013, Pinedo 2016) address many of these methods and characteristics.

In the last few decades, *personnel scheduling* or *human resource allocation* problems have been studied widely because labor cost is a major direct cost in many environments (Ernst et al. 2004a, b; Van Den Bergh et al. 2013). These problems provide many unique challenges given the need to satisfy personal preferences and variability across workers' skill sets.

Within healthcare, the personnel scheduling literature can be primarily divided into three categories. The first, and most abundant, is in *nurse scheduling* (Cheang et al. 2003, Burke et al. 2004). In the second category, *physician scheduling* (Erhard et al. 2017), most of the literature focuses either on shift scheduling, primarily for emergency department physicians (Beaulieu et al. 2000, Carter and Lapierre 2001, Gendreau et al. 2006, Topaloglu 2006), or on scheduling operating room time for surgeons (Cardoen et al. 2010). Finally, the third category—residency scheduling—focuses specifically on the balance between patient care and training (Ozkarahan 1994, Turner et al. 2013, Guo et al. 2014). Our research falls into this third category. In particular, we focus on the challenges associated with finding and implementing solutions in a real-world environment with multiple and sometimes qualitative objective criteria.

Within this literature, Bard et al. (2016b) proposed several integer programming-based heuristics for constructing annual block schedules for family medicine residents with continuity clinic considerations. Bard et al. (2016a, 2017) focused on block schedules for internal medicine residents. Proano and Agarwal (2018) suggested a multistage, multiobjective optimization approach for weekly resident rotation schedules. Cohn et al. (2009) combined heuristic and mixed-integer programming approaches to generate one-year call schedules across three hospitals staffed by the psychiatry residency program at the Boston University School of Medicine. Within shift scheduling, Bard et al. (2013) proposed a three-phase solution approach for creating monthly schedules that assign residents to outpatient clinics in which preemptive goal programming minimizes violations of a prioritized set of goals. Guo et al. (2014) presented a generic version of the resident scheduling problem that produces a one-year schedule and showed a proof of its NP-completeness.

Much of the residency scheduling literature has had to address the issue of multiple objective criteria. One approach to addressing this problem is to put weights on each of the metrics to establish a single objective function (Caramia and Dell'Olmo 2008). To identify the weights, other papers (Ozkarahan 1994; Topaloglu 2006, 2009; Proano and Agarwal 2018) have utilized the analytic hierarchy process (AHP) (Saaty 1988), in which users are surveyed to provide relative preferences between pairs of metrics. In the AHP, the decision maker first decomposes the criteria into a hierarchical decision model and selects elements at each level of the hierarchy. Once the decision hierarchy is constructed, the decision maker systematically evaluates various elements at each level by making pairwise comparisons of the elements. This approach assumes a prior knowledge about the hierarchy of the criteria and linearity between priorities among the elements of the hierarchy. We note that, in our collaboration, these methods were not successful in addressing the problem in part because the chief residents' preferences varied from month to month. The time required to make sufficient pairwise comparisons (which would only be relevant for the current schedule) would be greater than the time needed for the process that we employ instead.

Solution Approach

Many of us were taught in the classroom a simplified version of how to solve problems such as PEDS using the following set of steps:

1. A problem is defined by the end user.
2. A model is formulated by the operations research (OR) practitioner.
3. An instance of the model is implemented and solved using commercial software/solvers.
4. The results are reported by the OR practitioner and implemented by the end user.

In practice, it is often substantially more challenging. The OR literature has an abundance of research addressing step 2, in which there are challenges in formulating the model, and step 3, in which specialized algorithms or heuristics must be developed. Far less common is discussion of challenges within the process as a whole.

In our experience with both PEDS and other work that we have done in applying optimization techniques to solve real-world residency scheduling problems, we have found that the formulation of the model and solving of an instance are often not particularly problematic. On the other hand, the following challenges *are* of significant consequence.

Challenge 1: Heterogeneous Resident Pool and Large Number of Complex Rules

In some residency scheduling problems, the resident pool is homogeneous in terms of skill sets and requirements, and the coverage needs are straightforward as well. For example, in a surgical residency program, there may be a simple "Q4" call schedule, in which four residents rotate daily with each resident doing call every fourth day, possibly with a few minor modifications to the schedule that can be accommodated with simple swaps. Such a resident schedule may be built manually or using simple heuristics and some basic spreadsheet support.

Other problems, however, are more complex with a heterogeneous pool of residents (ranging in level of experience and seniority, coming from multiple different training programs and even institutions with varying educational requirements to be filled) and complex rules defining appropriate coverage, rest rules, educational requirements, and more. Such is the case with the PEDS problem that we faced.

In these cases, it is very difficult for the chief residents to even build feasible schedules by hand, let alone schedules that satisfy individual personal requests, ensure equity across residents, and focus on opportunities to improve patient outcomes through continuity of care and related characteristics. The problems are simply too complex combinatorially to solve without the help of more sophisticated techniques, such as mathematical programming.

Challenge 2: Lack of a Well-Defined Objective Function

Although mathematical programming techniques can often help us to quickly and accurately find feasible solutions to residency scheduling problems, a challenge remains in identifying what objective function to optimize. Unlike many industrial applications but similar to many residency scheduling problems, the goal of PEDS is not to maximize profit or minimize cost; there is no cost function involved. On the other hand, there are often many criteria that impact the quality of a schedule. In the case of the PEDS problem, these include the following:

- Total shift equity: Because not all residents are on service for the same amount of time (e.g., some have a full-month and some have a half-month assignment), not all residents should perform equal numbers of shifts. The chiefs do, however, pay careful attention to how many shifts each resident is working, seeking to be equitable.
- Night shift equity: Similarly, the chiefs pay careful attention to how many night shifts are assigned to each resident to ensure fairness.
- Bad sleep patterns: A *bad sleep pattern* is a sequence of shifts that are legal from a duty-hours perspective but undesirable relative to circadian rhythm. For example, if a resident works a morning shift on Monday, an overnight shift on Tuesday, and then another morning shift on Thursday, it is difficult for that resident to match the resident's sleep and work schedules. For more details, see table 2 in Perelstein et al. (2016). Whenever possible, the chiefs prefer to avoid assigning bad sleep patterns.
- Postcontinuity clinic shifts: Although it is legal to work an ED shift immediately after a continuity clinic, the chiefs again prefer to avoid such assignments because it requires the resident to either leave the clinic early or arrive at the ED late in addition to creating a work day of nearly 14 hours.
- Intern-undesirable shifts: Although certain shifts are fully prohibited from interns being assigned to them, there are also shifts that are allowed but undesirable.
- Covered optional shifts: As noted earlier, interns begin their ED rotations on the 27th of the preceding month. Thus, when building a monthly schedule, it is permissible to leave shifts from the 27th through the end of the month unfilled to be filled when building next month's schedule. This can be difficult, however, if the following month is tightly constrained. Thus, the chiefs may not want to leave too many optional shifts uncovered if they anticipate difficulties in building the following month's schedule.
- Uncovered flex shifts: Similarly, it is not required that all flex shifts be covered, but these correspond to busy times during the daily cycle of the ED's patient volume, and therefore, when possible, it is desirable to staff them.

Not only are there multiple objective criteria of importance to the chiefs, but these objective functions are often nonlinear (e.g., the difference between one bad sleep pattern and two bad sleep patterns is not the same as the difference between four and five). Furthermore, the relative importance of these metrics may vary from month to month; for example, at the beginning of the academic year, greater focus may be placed on having more senior residents staffing the most critical shifts, and during a later month, greater emphasis may be made on accommodating residents' interviews for fellowship positions.

Challenge 3: Frequently Changing and Nuanced Problem Requirements

Finally, there are often frequently changing and nuanced one-off rules and preferences that crop up each month. For example, a resident in the late stages of pregnancy might require a variation in rest rules or a set of shifts that are less physically taxing. Another example occurred when a new electronic health record (EHR) was being implemented, during which period the staffing levels increased substantially to account for both resident training and the anticipated disruptions to patient care as all providers became familiar with the new system. Each of these one-offs can be critical to ensuring that the resulting schedule is viable in practice, and there are similar month-to-month changes in terms of new criteria being used to evaluate schedule quality.

How, then, do we effectively and efficiently use mathematical programming techniques to solve a problem in which the rules change every month and the objective function is not defined? We do so by leveraging some key opportunities:

First, the size and nature of PEDS is such that it is fairly straightforward to formulate most of the feasibility requirements within an integer programming model, and real-world instances of this model can be solved quickly using a standard implementation of a commercial solver, such as CPLEX. A typical feasibility instance for PEDS solves in well under a minute.

Second, we have developed a C++-based framework for implementing instances of PEDS, which we designed specifically to allow great flexibility in changing the underlying rules of the model with generalized data structures and associated Excel-based tools to facilitate the creation of supporting data files. For example, when the EHR was launched and the number of daily shifts, the timing of these daily shifts, the number and type of residents required for each shift, and the rest rules associated with these shifts all changed, we did not need to change the model or write any new code.

Third, we have the good fortune to work within a team that has both chief residents who are engaged, enthusiastic, and committed to building schedules and a team of students with strong technical as well as communication skills. Thus, we have been able to develop a truly collaborative relationship in which we first worked together to define the problem, then built the underlying software tool, and then learned together the most effective way to jointly build the monthly schedule. The overall time spent working together as a team is far less than if the chief residents built the schedule by hand, as in the past, or even if we divided the process with clear boundaries between the chiefs' role and the students' role. Even more importantly, by working together on a regular basis, the sharing of knowledge across disciplines greatly facilitates future changes and addressing future obstacles. For example, as the chiefs begin to understand the basic ideas behind IP modeling, they too start to understand patterns and structures, and when a new change needs to be implemented, they can not only communicate that need effectively, but can often draw parallels to existing rules that, although very different in clinical function, are very similar in mathematical structure.

Fourth, and perhaps both most critical to our success and most distinct from how many of us learned to solve OR problems, *the chief residents do not want an optimal solution!* It's not that they are willing to settle for something that is suboptimal; they don't know what optimal is. Arguably, there is no one optimal solution or set of solutions. We draw the analogy to buying a house. Some houses are clearly "infeasible" in failing to satisfy your basic requirements (one is too expensive, another only has one bathroom). Given two "feasible" houses, one may clearly dominate the other (less expensive, nicer kitchen, better school district). But, often, we are faced with a set of houses that are all feasible and all have lots of good characteristics and for which there is no clean way to compare them against each other and formally rank them. At the end of the day, we just pick one.

The same is true for PEDS; the notion of an optimal solution is meaningless, and trying to define an objective function to achieve optimality does not meet the needs of the chiefs. Rather, they need a mechanism for finding a feasible solution that is of high quality (is perceived as equitable by the residents and makes a reasonable trade-off between competing metrics, with which "reasonable" is largely in the eyes of the chiefs) and can be found quickly.

As such, it does not make sense to solve PEDS hierarchically; no one metric fully dominates. Nor does it make sense to use a weighted objective function. We have found in many tests that, when users apply weights, the resulting solution is rarely the solution that they actually prefer even when such a preference function is well defined, which is not the case for PEDS. Furthermore, because the priorities changes from month to month (e.g., one month prioritizing adjustments to a new EHR, another month prioritizing fellowship interviews) and shift each year as the new chief residents bring their own values and opinions, using machine learning or other techniques to infer weights is not an option because they are constantly changing.

Instead, we define metric constraints (shown in the appendix) for each metric respectively. With the feasibility formulation, which defines hard constraints that are not allowed to be violated, we solve for the feasibility problem plus the metric constraints with arbitrary lower and upper bounds. Then, based on the judgment of the chief resident for metrics, we adjust the boundary of the resulting feasible region and use it for the next new multiobjective optimization problem that optimizes the metrics of evaluation. We illustrate this process in Figure 1.

What we have found to be highly successful, enabling us to build monthly PEDS schedules for almost a decade, is a blending of very precise and technical integer programming techniques with very imprecise and nontechnical collaboration. Specifically, the process works as follows:

- Each month, the chief residents provide the engineering students on the team with a set of Excel-based

Figure 1. (Color online) The Iterative, Interactive Approach in PEDS

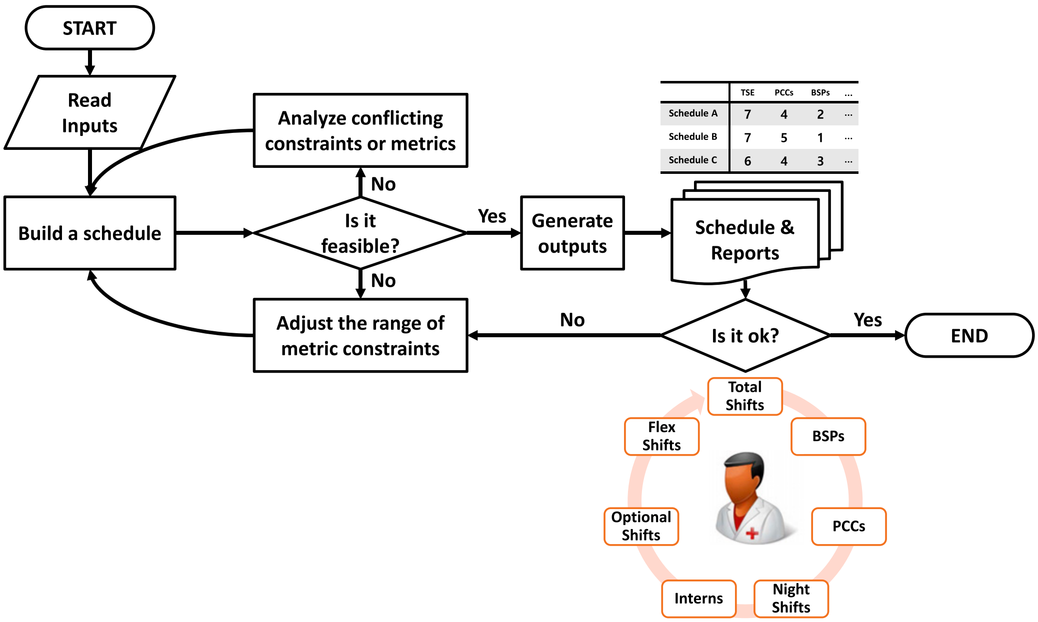

input files, containing the upcoming month's set of residents, their characteristics (e.g., their programs and levels), and their requests as well as more informal information about any expected changes or challenges for the month (for example, "We know that a lot of the seniors will be gone on Fridays this month, so we may have to allow interns to cover a few of the shifts that typically only allow seniors to cover").

• The student team then builds a preliminary schedule. The students begin by first ensuring pure feasibility; in many cases, this initially cannot be done, and email exchanges with the chiefs allow for joint decision making about when "hard" rules need to be violated given the month's unique challenges.

• Given the rules needed to enable a feasible schedule, the students then make a second pass at optimizing the schedule. By setting upper and lower bounds on each of the defined metrics, they solve multiple feasibility problems, seeking a high-quality schedule based on their knowledge of the chiefs' preferences.

• We then all meet as a team (typically for an hour or two) with the chiefs reviewing the proposed draft schedule. They are able to identify nuanced opportunities for improvement (e.g., pointing out when a resident has some special needs and a change to the schedule might address those needs). Because the integer programming feasibility problem can be solved so quickly, it is not uncommon to generate a dozen or more draft versions of the schedule, adding and subtracting bounds or imposing new rules, until the chiefs declare the schedule satisfactory.

This interactive process has the added benefits of (1) helping the students to learn about the different factors the chiefs consider and, therefore, giving them knowledge that will improve their building of higher-quality first drafts for future months; (2) identifying major changes that need more complex solutions and providing the chance to work through these together as a team (this is where most enhancements to the software have arisen as well as other collaborative projects); and (3) providing the engineering students with opportunities for developing softer skills (e.g., in communications and professionalism).

Finally, we make a brief observation about the sustainability of this approach. In our case, we have been able to do so because we have an abundance of undergraduate students eager to gain hands-on experience in applications of OR, particularly within healthcare. These students can work at low cost to the residency program with fairly limited supervision from engineering faculty and staff, and thus, it is a win–win situation (especially because this also provides a relationship on which additional research projects can be built). This is not the case for every residency program (although most such programs do naturally have affiliations with universities that may have industrial engineering or operations research programs). We also recognize that this approach is complex enough that a pure "shrink-wrapped," off-the-shelf software package would be limited in its effectiveness. Our involvement in both this project and several other similar projects stemmed from the fact that the residency programs were finding the available commercial solver to be unable to handle their unique nuances and one-off constraints. On the other hand, this does not mean that a commercialization of our approach would be unsuccessful; we suggest that the key would be in offering a monthly service rather than a stand-alone independent software tool as is seen with many other healthcare scheduling tools.

Computational Experiences and Results

We have been collaborating with the University of Michigan Pediatrics chief residents for nearly 10 years in building their monthly ED shift schedules, using the process described and working together to not only build schedules, but to continuously improve our models, tools, and processes.

From a timing standpoint, there has been marked improvement. An individual instance of the integer programming feasibility problem rarely takes more than a few seconds and never more than a minute or two. In terms of the process as a whole, it typically takes the chiefs a few hours each month to collect all of the necessary data; the primary time is in gathering the individual residents' personal requests and continuity clinic information. It then takes the OR students a few hours to build a draft schedule. Finally, the team usually spends one to two hours together each month to fine-tune and finalize the schedule. This time often includes discussion of special circumstances (e.g., preparing for the holiday season schedule, which has unique requirements) and other opportunities for collaboration, such as working together to build the annual block schedule.

In contrast, prior to this collaboration, the chiefs would typically spend between 20 and 30 hours per month building the schedule by hand. In addition, they would invest substantial time at the end of each year in training the new incoming chiefs to generate schedules.

Not only is substantial time saved through our collaborative process, but in addition, the quality of the schedule itself is substantially improved. This is in part because the chiefs can spend more time focusing on special situations and troubleshooting and in part because more complex metrics (e.g., bad sleep patterns) can be taken into consideration that were not possible to incorporate when scheduling by hand. Table 1 highlights the changes to two key metrics from the year before we began our collaboration to the years after the collaboration was fully established (we exclude the year that the process was only partially in place). Although, in some cases, metrics may worsen (this can be a function of the specific instance or of preferences placed on other metrics by the chiefs), in almost all cases there was improvement. In particular, many months under the new approach saw a complete elimination of bad sleep patterns.

Conclusion

The scheduling of medical residents, whether for shifts, call, or block schedules, presents a valuable opportunity for OR practitioners to help improve the training of residents and the quality of care that they provide to their patients. In some cases, these schedules may be amenable to straightforward heuristics that leverage repeating patterns and a limited set of rules. In other cases, as in our experience, integer programming techniques, in combination with collaborative discussion with the chief residents, can greatly facilitate the process.

An additional side benefit of our experience has been the opportunity to train undergraduate engineering students in the softer skills of interprofessional communication, project and time management, and meeting deadlines. The students have also benefited from the opportunity to meet regularly with the chiefs and, in the process, learn more about the contextual nuances of the problem domain.

Table 1. The Effects of Our Collaborative Process on Bad Sleep Patterns and Postcontinuity Clinics (the 2010–2011 Schedules Were Made by Hand)

	Bad sleep patterns				
Month	2010–2011	2012–2013	2013–2014	2014–2015	2015–2016
July	10	0	0	0	0
August	3	0	0	0	0
September	9	14	0	0	0
October	5	0	0		1
November	13	0	0	3	0
December	4	0	0	1	0
January	8	0	0	0	0
February	8	0	0	0	0
March	6	0	1	0	0
April	4	0	0	0	0
May	6	0	0	0	6
June	7	0	0	0	3
Average	6.92	1.17	0.08	0.33	0.83
Standard deviation	2.75	3.87	0.28	0.85	1.77
Median	6.50	0.00	0.00	0.00	0.00
Mode	4.00	0.00	0.00	0.00	0.00

	Postcontinuity clinic shifts				
Month	2010–2011	2012–2013	2013–2014	2014–2015	2015–2016
July	6	4	5	0	3
August	1	7	0	6	2
September	6	5	3	3	0
October	9	0	2	7	3
November	11	8	2	3	3
December	1	0	1	0	2
January	7	0	2	5	0
February	6	0	3	0	0
March	8	8	3	5	1
April	4	0	0	1	1
May	6	0	4	5	6
June	7	0	11	0	6
Average	6.00	2.67	3.00	2.92	2.25
Standard deviation	2.80	3.32	2.80	2.53	2.01
Median	6.00	0.00	2.50	3.00	2.00
Mode	6.00	0.00	3.00	0.00	3.00

Finally, this interactive process has led to opportunities for improving the existing tool/process and adjusting for new program needs as well as new collaborative opportunities.

Future opportunities in this area include improving the usability of the software so that chief residents might, in some cases, be able to build and modify schedules independently; developing mechanisms for capturing and interpreting the chief residents' preferences to provide a feedback loop for schedule modification; creating heuristics that identify multiple schedules with varying characteristics to enhance the iterative process; and exploring commercialization for the benefit of residency programs at institutions where partnering with engineering students is not an option.

Acknowledgments
We gratefully acknowledge the support of all students and collaborators from the Center for Healthcare Engineering and Patient Safety.

Appendix
Feasibility Formulation
The following demonstrates an integer programming formulation of the basic PEDS feasibility problem.

Notation

R Set of residents, $r \in \{1, 2, \ldots, |R|\}$.

D Set of days, $d \in \{1, 2, \ldots, |D|\}$; in our data instances, D is either 35 or 36.

S Set of daily shifts, $s \in \{1, 2, \ldots, |S|\}$; in our data instances, there are seven shifts.

T Set of training programs, $t \in \{\text{PED, EM, FM}, \ldots\}$.

t_r Training program of resident r, $t_r \in T\ \forall r \in R$.

C_r Set of days on which resident r has continuity clinic, $C_r \subseteq D\ \forall r \in R$.

W_r Set of days on which resident r can work, $W_r \subseteq D\ \forall r \in R$.

F Set of flex shifts, $F \subset S$; in our data instances, shift 3 (i.e., 12 p.m. to 9 p.m.) is the only element of F.

N Set of night shifts, $N \subset S$; in our data instances, shifts 6 and 7 (i.e., 8 p.m. to 5 a.m. and 11 p.m. to 8 a.m.) are in N.

K Set of shift pairs such that at least one shift in the pair must be covered by a resident r of type *PED*, $k \subset S$; in our data instances, (shifts 1 and 2), (shifts 4 and 5), and (shifts 6 and 7) are in K.

H_r Set of day–shift pairs that cannot be assigned to resident r, $H_r \subset D \times S\ \forall r \in R$.

A_r Set of day–shift pairs that must be assigned to resident r, $A_r \subset D \times S\ \forall r \in R$.

$J_{(d,s)}$ Set of day–shift pairs that would cause a duty-hours violation if assigned in addition to day–shift (d, s), $J_{ds} \subset D \times S\ \forall (d, s) \in D \times S$. Note that (d, s) is included in the set $J_{(d,s)}$.

$\overline{D}$ Maximum allowable number of days in a row that a resident can work a shift.

$\overline{N}$ Maximum allowable number of days in a row that a resident can work a night shift.

Decision Variables

x_{rds} Binary variable, equals one if resident r is assigned on day d to shift s; otherwise $0\ \forall r \in R, \forall d \in D, \forall s \in S$.

Formulation

$$\min\ 0 \tag{A.1}$$

$$\text{subject to} \sum_{r\in R} x_{rds} = 1 \qquad \forall s \in S\backslash F, \forall d \in D \tag{A.2}$$

$$\sum_{r\in R} x_{rds} \le 1 \qquad \forall s \in F, \forall d \in D \tag{A.3}$$

$$\sum_{(d,s)\in H_r} x_{rds} = 0 \qquad \forall r \in R \tag{A.4}$$

$$\sum_{(d,s)\in A_r} x_{rds} = |A_r| \qquad \forall r \in R \tag{A.5}$$

$$\sum_{\{r\in R: t_r = \text{PED}\}} \sum_{s\in k} x_{rds} \ge 1 \qquad \forall d \in D, \forall k \in K \tag{A.6}$$

$$\sum_{(d,s)\in J_{(d,s)}} x_{rds} \le 1 \qquad \forall r \in R, \forall d \in D, \forall s \in S \tag{A.7}$$

$$\sum_{i=d}^{d+\overline{D}} \sum_{s\in S} x_{ris} \le \overline{D} \qquad \forall r \in R, \forall d \in 1, \ldots, |D| - \overline{D} \tag{A.8}$$

$$\sum_{i=d}^{d+\overline{N}} \sum_{s\in N} x_{ris} \le \overline{N} \qquad \forall r \in R, \forall d \in 1, \ldots, |D| - \overline{N} \tag{A.9}$$

$$x_{rds} \in \{0, 1\} \qquad \forall r \in R, \forall d \in D, \forall s \in S \tag{A.10}$$

Constraints (A.2) ensure that exactly one resident is assigned to every nonflex shift. Constraints (A.3) ensure that flex shifts are covered by at most one resident. Constraints (A.4) ensure that residents are not assigned to prohibited shifts. Constraints (A.5) ensure that residents cover preassigned shifts. Constraints (A.6) ensure that at least one of two shifts in a required pair is assigned to a pediatrics resident. Constraints (A.7) ensure that residents are not assigned to shifts that begin within 10 hours of a previously worked shift. Constraints (A.8) ensure that no resident works more than a maximum allowable number of days in a row. Constraints (A.9) ensure that no resident works more than a maximum allowable number of nights in a row.

Metric Formulation
The following demonstrates an integer programming formulation of metric constraints for the PEDS problem.

Notation

G Set of resident *groups* in which $g \subset G$ is the set of residents who have the same number of working days, $g = \{r \in R : |W_r| \text{ is the same}\} \subset R$.

I Set of intern-undesirable shifts, $I \in S$; in our data instances, shifts 1 and 7 are in I.

P Set of shifts that are defined as the postcontinuity clinic shifts, $P \subset S$; in our data instances, shifts 6 and 7 (i.e., 8 p.m. to 5 a.m. and 11 p.m. to 8 a.m.) are in P.

E Set of day–shift pairs that are defined as the optional shifts around the end of the planning horizon, $E \subset D \times S$.

U Set of bad (undesirable) sleep patterns in which $u \in U$ is a combination of shift offsets on multiple days. Note that $|\, u \,|$ represents the number of shift offsets in u.

$U_{(d,u)}$ Set of date–shift pairs associated with bad sleep pattern u on day d, $U_{(d,u)} \subset D \times S\ \forall (d, u) \in D \times U$.

Input Parameters

$\underline{S}_g$, $\overline{S}_g$ Lower and upper bounds on the number of total shifts for a resident in group g, $\forall g \in G$.

$\underline{N}_g$, $\overline{N}_g$ Lower and upper bounds on the number of night shifts for a resident in group g, $\forall g \in G$.

$\underline{U}_g$, $\overline{U}_g$ Lower and upper bounds on the number of bad sleep patterns for a resident in a group g, $\forall g \in G$.

$\underline{P}_g$, $\overline{P}_g$ Lower and upper bounds on the number of postcontinuity clinic shifts for a resident in a group g, $\forall g \in G$.

$\underline{I}$, $\overline{I}$ Lower and upper bounds on the number of intern residents assigned to intern-prohibited shifts in the planning horizon.

$\underline{E}$, $\overline{E}$ Lower and upper bounds on the number of covered optional shifts in the planning horizon.

$\underline{F}$, $\overline{F}$ Lower and upper bounds on the number of uncovered flex shifts in the planning horizon.

Metric Variables

y_{rdu} Binary variable, equals one if resident r is assigned to bad sleep pattern u on day d; otherwise, $0\ \forall r \in R, \forall d \in D, \forall u \in U$.

z_{rd} Binary variable, equals one if resident $r \in R$ is assigned to work a postcontinuity clinic shift on day $d \in C_r$; otherwise, $0\ \forall r \in R, \forall d \in D$.

s_r Number of total shifts for resident r, $s_r = \sum_{d\in D}\sum_{s\in S} x_{rds}$ for $\forall r \in R$.

n_r Number of night shifts for resident r, $n_r = \sum_{d\in D}\sum_{s\in N} x_{rds}$ for $\forall r \in R$.

u_r Number of bad sleep patterns for resident r, $u_r = \sum_{d\in D}\sum_{u\in U} y_{rdu}$ for $\forall r \in R$.

p_r Number of postcontinuity clinic shifts for resident r, $p_r = \sum_{d\in C_r}\sum_{s\in P} x_{rds}$ for $\forall r \in R$.

i Total number of intern residents assigned to intern-undersirable shifts in the planning horizon, $i = \sum_{r\in\{r:l_r=\text{interns}\}}\sum_{d\in D}\sum_{s\in I} x_{rds}$.

e Total number of covered optional shifts in the planning horizon, $e = \sum_{r\in R}\sum_{(d,s)\in E} x_{rds}$.

f Total number of uncovered flex shifts in the planning horizon, $f = |D| - \sum_{r\in R}\sum_{d\in D}\sum_{s\in F} x_{rds}$.

Metric Constraints

$$y_{rds} \le x_{rij} \qquad \forall r \in R, \forall d \in D, \forall u \in U, \forall (i,j) \in U_{(d,u)} \tag{A.11}$$

$$y_{rds} + |u| \ge \sum_{(d,s)\in U_{(d,u)}} x_{rds} + 1 \qquad \forall r \in R, \forall d \in D, \forall u \in U \tag{A.12}$$

$$z_{rd} \ge x_{rds} \qquad \forall r \in R, \forall d \in C_r, \forall s \in P \tag{A.13}$$

$$z_{rd} \le \sum_{s\in P} x_{rds} \qquad \forall r \in R, \forall d \in C_r \tag{A.14}$$

$$s_r = \sum_{d\in D}\sum_{s\in S} x_{rds} \qquad \forall r \in R \tag{A.15}$$

$$n_r = \sum_{d\in D}\sum_{s\in N} x_{rds} \qquad \forall r \in R \tag{A.16}$$

$$u_r = \sum_{d\in D}\sum_{u\in U} y_{rdu} \qquad \forall r \in R \tag{A.17}$$

$$p_r = \sum_{d\in C_r}\sum_{s\in P} x_{rds} \qquad \forall r \in R \tag{A.18}$$

$$i = \sum_{r\in R}\sum_{d\in D}\sum_{s\in I} x_{rds} \tag{A.19}$$

$$e = \sum_{r\in R}\sum_{(d,s)\in OPT} x_{rds} \tag{A.20}$$

$$f = |D| - \sum_{r\in R}\sum_{d\in D}\sum_{s\in F} x_{rds} \tag{A.21}$$

$$\underline{S}_g \le s_r \le \overline{S}_g \qquad \forall g \in G, \forall r \in g \tag{A.22}$$

$$\underline{N}_g \le n_r \le \overline{N}_g \qquad \forall g \in G, \forall r \in g \tag{A.23}$$

$$\underline{U}_g \le u_r \le \overline{U}_g \qquad \forall g \in G, \forall r \in g \tag{A.24}$$

$$\underline{P}_g \le p_r \le \overline{P}_g \qquad \forall g \in G, \forall r \in g \tag{A.25}$$

$$\underline{I} \le i \le \overline{I} \tag{A.26}$$

$$\underline{E} \le e \le \overline{E} \tag{A.27}$$

$$\underline{F} \le f \le \overline{F} \tag{A.28}$$

$$y_{rdu} \in \{0,1\} \qquad \forall r \in R, \forall d \in D, \forall u \in U \tag{A.29}$$

$$z_{rd} \in \{0,1\} \qquad \forall r \in R, \forall d \in C_r \tag{A.30}$$

Constraints (A.11) and (A.12) link decision variables x_{rds} and auxiliary variables y_{rds} to count the number of bad sleep patterns. Constraints (A.13) and (A.14) link decision variables x_{rds} and auxiliary variables z_{rd} to count the number of postcontinuity clinic shifts.

Metric constraints (A.22) and (A.23) control the upper and lower bounds on the number of shifts and night shifts that are assigned to each resident r in a group $g \in G$. Metric constraints (A.24) control the number of bad sleep patterns that are assigned to each resident r in a group $g \in G$. Metric constraints (A.25) control that the number of postcontinuity clinic shifts that are assigned to each resident r in a group $g \in G$. Metric constraints (A.26) control the total number of interns assigned to intern residents assigned to intern-undersirable shifts in the planning horizon. Metric constraints (A.27) control the total number of covered optional shifts in the planning horizon. Metric constraints (A.28) control the total number of uncovered flex shifts in the planning horizon.

References

Bard JF, Shu Z, Leykum L (2013) Monthly clinic assignments for internal medicine housestaff. *IIE Trans. Healthcare Systems Engrg.* 3(4):207–239.

Bard JF, Shu Z, Morrice DJ, Leykum LK (2016a) Annual block scheduling for internal medicine residents with 4+ 1 templates. *J. Oper. Res. Soc.* 67(7):911–927.

Bard JF, Shu Z, Morrice DJ, Leykum LK (2017) Constructing block schedules for internal medicine residents. *IISE Trans. Healthcare Systems Engrg.* 7(1):1–14.

Bard JF, Shu Z, Morrice DJ, Leykum LK, Poursani R (2016b) Annual block scheduling for family medicine residency programs with continuity clinic considerations. *IIE Trans.* 48(9):797–811.

Beaulieu H, Ferland JA, Gendron B, Michelon P (2000) A mathematical programming approach for scheduling physicians in the emergency room. *Healthcare Management Sci.* 3(3):193–200.

Brucker P, Qu R, Burke E (2011) Personnel scheduling: Models and complexity. *Eur. J. Oper. Res.* 210(3):467–473.

Burke EK, De Causmaecker P, Berghe GV, Van Landeghem H (2004) The state of the art of nurse rostering. *J. Scheduling* 7(6):441–449.

Caramia M, Dell'Olmo P (2008) *Multi-objective Management in Freight Logistics: Increasing Capacity, Service Level and Safety with Optimization Algorithms* (Springer Science & Business Media, London).

Cardoen B, Demeulemeester E, Beliën J (2010) Operating room planning and scheduling: A literature review. *Eur. J. Oper. Res.* 201(3):921–932.

Carter MW, Lapierre SD (2001) Scheduling emergency room physicians. *Healthcare Management Sci.* 4(4):347–360.

Cheang B, Li H, Lim A, Rodrigues B (2003) Nurse rostering problems: A bibliographic survey. *Eur. J. Oper. Res.* 151(3):447–460.

Cohn A, Root S, Kymissis C, Esses J, Westmoreland N (2009) Scheduling medical residents at Boston University School of Medicine. *Interfaces* 39(3):186–195.

Edie LC (1954) Traffic delays at toll booths. *J. Oper. Res. Soc. Amer.* 2(2):107–138.

Erhard M, Schoenfelder J, Fügener A, Brunner JO (2017) State of the art in physician scheduling. *Eur. J. Oper. Res.* 265(1):1–18.

Ernst AT, Jiang H, Krishnamoorthy M, Sier D (2004a) Staff scheduling and rostering: A review of applications, methods and models. *Eur. J. Oper. Res.* 153(1):3–27.

Ernst AT, Jiang H, Krishnamoorthy M, Nott H, Sier D (2001) An integrated optimization model for train crew management. *Ann. Oper. Res.* 108(1–4):211–224.

Ernst AT, Jiang H, Krishnamoorthy M, Owens B, Sier D (2004b) An annotated bibliography of personnel scheduling and rostering. *Ann. Oper. Res.* 127(1–4):21–144.

Franz LS, Miller JL (1993) Scheduling medical residents to rotations: Solving the large-scale multiperiod staff assignment problem. *Oper. Res.* 41(2):269–279.

Fry MJ, Magazine MJ, Rao US (2006) Firefighter staffing including temporary absences and wastage. *Oper. Res.* 54(2):353–365.

Gans N, Koole G, Mandelbaum A (2003) Telephone call centers: Tutorial, review, and research prospects. *Manufacturing Service Oper. Management* 5(2):79–141.

Gendreau M, Ferland J, Gendron B, Hail N, Jaumard B, Lapierre S, Pesant G, Soriano P (2006) Physician scheduling in emergency rooms. Burke EK, Rudová H, eds., *Practice and Theory of Automated Timetabling VI. PATAT 2006*, Lecture Notes in Computer Science, vol. 3867 (Springer, Berlin Heidelberg), 53–66.

Goldman ML, Shah RN, Bernstein CA (2015) Depression and suicide among physician trainees: Recommendations for a national response. *JAMA Psychiatry* 72(5):411–412.

Graves SC (1981) A review of production scheduling. *Oper. Res.* 29(4): 646–675.

Guler MG (2013) A hierarchical goal programming model for scheduling the outpatient clinics. *Expert Systems Appl.* 40(12):4906–4914.

Guo J, Morrison DR, Jacobson SH, Jokela JA (2014) Complexity results for the basic residency scheduling problem. *J. Scheduling* 17(3):211–223.

Li Y, Kozan E (2009) Rostering ambulance services. *Indust. Engrg. Management Soc.*, 795–801.

Lockley SW, Barger LK, Ayas NT, Rothschild JM, Czeisler CA, Landrigan CP (2007) Effects of healthcare provider work hours and sleep deprivation on safety and performance. *Joint Commission J. Quality Patient Safety* 33(Suppl. 11):7–18.

Ozkarahan I (1994) A scheduling model for hospital residents. *J. Medical Systems* 18(5):251–265.

Perelstein E, Rose A, Hong YC, Cohn A, Long MT (2016) Automation improves schedule quality and increases scheduling efficiency for residents. *J. Graduate Medical Ed.* 8(1):45–49.

Pinedo ML (2016) *Scheduling* (Springer International Publishing, Cham, CH).

Proano RA, Agarwal A (2018) Scheduling internal medicine resident rotations to ensure fairness and facilitate continuity of care. *Healthcare Management Sci.* 21(4):461–474.

Raff S (1983) Routing and scheduling of vehicles and crews: The state of the art. *Comput. Oper. Res.* 10(2):63–211.

Rogers AE (2008) The effects of fatigue and sleepiness on nurse performance and patient safety. Hughes RG, ed. *Patient Safety and Quality: An Evidence-Based Handbook for Nurses* (Agency for Healthcare and Research Quality, Rockville, MD)

Saaty TL (1988) What is the analytic hierarchy process? Mitra G, Greenberg HJ, Lootsma FA, Rijkaert MJ, Zimmermann HJ, eds., *Mathematical Models for Decision Support.* NATO ASI Series (Series F: Computer and Systems Sciences), vol. 48 (Springer, Berlin Heidelberg), 109–121.

Schwenk TL, Davis L, Wimsatt LA (2010) Depression, stigma, and suicidal ideation in medical students. *JAMA* 304(11):1181–1190.

Sen S, Kranzler HR, Krystal JH, Speller H, Chan G, Gelernter J, Guille C (2010) A prospective cohort study investigating factors associated with depression during medical internship. *Arch. General Psychiatry* 67(6):557–565.

Shanafelt T, Habermann T (2002) Medical residents' emotional well-being. *JAMA* 288(15):1846–1847.

Shanafelt TD, Bradley KA, Wipf JE, Back AL (2002) Burnout and self-reported patient care in an internal medicine residency program. *Ann. Internal Medicine* 136(5):358–367.

Sherali HD, Ramahi MH, Saifee QJ (2002) Hospital resident scheduling problem. *Production Planning Control* 13(2):220–233.

Taylor PE, Huxley SJ (1989) A break from tradition for the San Francisco police: Patrol officer scheduling using an optimization-based decision support system. *Interfaces* 19(1):4–24.

Topaloglu S (2006) A multi-objective programming model for scheduling emergency medicine residents. *Comput. Indust. Engrg.* 51(3): 375–388.

Topaloglu S (2009) A shift scheduling model for employees with different seniority levels and an application in healthcare. *Eur. J. Oper. Res.* 198(3):943–957.

Turner J, Kim K, Mehrotra S, DaRosa DA, Daskin MS, Rodriguez HE (2013) Using optimization models to demonstrate the need for structural changes in training programs for surgical medical residents. *Healthcare Management Sci.* 16(3):217–227.

Vance PH, Barnhart C, Johnson EL, Nemhauser GL (1997) Airline crew scheduling: A new formulation and decomposition algorithm. *Oper. Res.* 45(2):188–200.

Van Den Bergh J, Beliën J, De Bruecker P, Demeulemeester E, De Boeck L (2013) Personnel scheduling: A literature review. *Eur. J. Oper. Res.* 226(3):367–385.

Verification Letter

Heather L. Burrows, MD PhD, Associate Chair of Education, Program Director, Pediatric Residency Program, Clinical Associate Professor, Department of Pediatrics, University of Michigan Medical School, Ann Arbor, Michigan 48109-5718, writes:

"As Program Director for the Pediatrics Residency Program at the University of Michigan, I have had the opportunity to work closely with Dr. Amy Cohn and her team as they have developed the process to create computer-generated schedules for our residents described in this article. I am very grateful for this work. Using automated tools to staff the emergency department has benefited our institution tremendously by providing substantially higher-quality schedules at significantly less effort.

"The computer-generated schedules provide for better balance (total and night shift equity), less resident fatigue (avoiding bad sleep patterns and postcontinuity clinic shifts), and greater patient care (intern-undesirable, covered optional, and uncovered flex shifts) than we could reasonably expect to create when building the schedule by hand. Over the last several years, our residents and patients have enjoyed safer care thanks to the high-quality schedules made using this tool. Moreover, attaining these better schedules requires little time involvement from our chief residents. Each month, the chief residents spend only a few hours in preparing inputs and reviewing the computer-generated schedule, whereas they used to spend nearly 30 hours each month to build the schedule by hand.

"I am excited that they have the opportunity to share this work with others, and I know it will be of great interest."

Young-Chae Hong holds a PhD degree in industrial and operations engineering from the University of Michigan. He is interested in combinatorial optimization and machine learning. His current research focuses on supply chain optimization using graph theory.

Amy Cohn is a professor in the Department of Industrial and Operations Engineering at the University of Michigan. She is also the Associate Director of the Center for Healthcare Engineering and Patient Safety. She holds an AB in applied mathematics from Harvard University and a PhD in operations research from the Massachusetts Institute of Technology. Her research interests focus on applied combinatorial optimization problems with multiple, ill-defined objective functions.

Stephen Gorga is currently transitioning from fellow to faculty in pediatric critical care medicine at C.S. Mott Children's Hospital at the University of Michigan, Ann Arbor, Michigan. He obtained his MD from the College of Human Medicine at Michigan State University. His area of medical education research focuses on advanced technologies to facilitate advanced trainee adult learning of high-risk, low-interaction situations.

Edmond O'Brien is a cardiothoracic anesthesiology fellow at the University of Michigan, where he also completed general pediatrics and anesthesiology residencies. He received his MD from Michigan State University.

William Pozehl is a researcher in the Center for Healthcare Engineering and Patient Safety at the University of Michigan. He received his BSE and MSE degrees at the University of Michigan, both in industrial and operations engineering. He designs and executes models for scheduling healthcare providers in clinical settings to better meet patient and provider needs.

Jennifer Zank is an assistant professor of pediatrics in the newborn medicine division at the University of Pittsburgh. She graduated from Loyola University Stritch School of Medicine with a medical degree. Her current research includes quality improvement in the neonatal intensive care unit with a focus on team communication and improved family inclusion in the medical decision-making process.