Analytical Approaches to Operating Room Management

Projects at Lucile Packard Children's Hospital Stanford

David Scheinker and Margaret L. Brandeau

Abstract In recent decades, healthcare has become increasingly expensive, creating pressure on healthcare providers to cut costs while maintaining or improving quality. Operations research can play an important role in supporting such efforts. A key challenge faced by hospital planners is scheduling and management of operating rooms, as operating rooms typically provide highly specialized care, require significant resources, and contribute significantly to a hospital's bottom line. We describe recent work on hospital operating room management at Lucile Packard Children's Hospital Stanford. We describe preliminary outcomes of three projects aimed at improving the efficiency of the hospital's operating rooms: machine learning to improve surgical case length estimation; queuing analysis to improve operational efficiency; and integer programming to schedule cases to reduce surgical delays.

Keywords Healthcare · Operations management · Optimization
Machine learning · Queueing

1 Introduction

In recent decades, healthcare has become increasingly expensive [17], creating pressure on healthcare providers to cut costs while maintaining or improving quality. The tools of operations research can play a key role in helping to improve the efficiency and effectiveness of healthcare services. Operations research analyses can be used to support high-level decisions such as facility planning (capacity, location, layout and design), public health planning, planning for population health needs, and human resource planning; tactical decisions such as capacity planning, case mix plan-

D. Scheinker
Department of Management Science and Engineering, Stanford University, Lucile Packard Children's Hospital Stanford, Stanford, CA, USA
e-mail: dscheink@stanford.edu

M. L. Brandeau (✉)
Department of Management Science and Engineering, Stanford University, Stanford, CA, USA
e-mail: brandeau@stanford.edu

P. Cappanera et al. (eds.), *Health Care Systems Engineering*, Springer Proceedings in Mathematics & Statistics 210, https://doi.org/10.1007/978-3-319-66146-9_2

ning, resource management, patient and resource scheduling, staffing assignment, and quality control and management; and operational decisions such as management of patient flows, waitlists, and staffing levels [1].

A key area of focus in many hospitals is planning and management of the perioperative environment, consisting of the operating rooms and their supporting facilities such as pre- and post-procedure units. Operating rooms typically provide highly specialized care, require significant resources, and contribute significantly to a hospital's bottom line. Even relatively minor delays in an operating room can have a significant impact on quality of patient care, staff satisfaction, and hospital financial stability. Indeed, the average cost of operating room time in the US is approximately $4000 per hour [15, 19] and it is estimated that each procedure that must be cancelled due to operating room delays reduces hospital revenue by approximately $1500 per hour [8].

For these reasons, many operations researchers have focused on developing models to improve the performance of hospital operating rooms. Extensive work has been carried out in areas such as case mix planning and patient and staff scheduling. Work on surgical procedure scheduling for the pediatric environment has combined optimization and simulation (e.g., [2, 4, 5, 23]). For a recent review, see [6]. We note, though, that many such planning models have not been implemented in practice. Additionally, some models that have been implemented in practice are quite specific to the hospital where they were developed and thus cannot be applied in other hospitals. Our goal is to develop methods that can be implemented at our hospital but also generalized to other hospitals.

In this paper we describe preliminary results from three projects in the perioperative environment that we are currently carrying out at Lucile Packard Children's Hospital Stanford (LPCH): using machine learning techniques to improve surgical case length prediction [27]; using queuing analysis to improve the operational efficiency of the perioperative process [10]; and using integer programming to schedule cases to minimize surgical delays [12].

LPCH is a 312-bed hospital that is part of the Stanford University healthcare system. The hospital has 7 operating rooms that are used to perform more than 6000 surgical procedures annually for 23 different services (e.g., cardiology, orthopedics). An expansion that will be completed in 2017 will add 149 beds and 6 additional operating rooms.

Planners at LPCH recently focused their attention on reducing delays in the operating rooms. The elective surgery process can be broken down into three stages: the surgeon sees the patient in clinic and schedules the procedure; the patient prepares for surgery at home, is prepared for surgery in the pre-operating room areas, and has the surgery; and the patient recovers from the procedure in a specialized unit. Three common causes of operating room delays and cancellations associated with these stages at LPCH (and other pediatric hospitals) are: surgical cases are mis-scheduled [3]; surgical preparation resources and processes are managed sub-optimally [13, 26]; and recovery beds for surgical patients are not immediately available [25]. We developed projects to systematically improve performance in each of these stages, as we now describe.

2 Surgical Case Length Prediction

In order to schedule procedures in an operating room, an estimate of the time needed to perform each procedure is required. Creating such estimates is particularly challenging in pediatric hospitals because pediatric patient populations tend to have widely variable needs, even for the same type of procedure [20]. In theory, the operating room time needed for each surgical case at LPCH is estimated by the surgeon after the surgeon examines the patient in clinic. Interviews with surgeons and non-clinician schedulers working at the surgical clinic reveal that, in practice, the scheduler estimates the time needed for surgery based on guidance from the surgeon or based on a historical average. Estimates from the various clinics are then used to manually create a schedule. This schedule is the basis for managing downstream patient flow (e.g., in the post-anesthesia recovery unit).

Predicted procedure durations are often very different from actual durations. When surgeries take less time than expected, operating rooms will be idle and patients may have to wait in the operating room for a recovery bed to become available. When surgeries take longer than expected, delays are incurred for subsequent surgeries, overtime may be required, and in some cases procedures must be cancelled. LPCH planners believed that better estimates of surgical case lengths would lead to improved operating room utilization, fewer delays, and higher patient and staff satisfaction.

We undertook a project to improve the prediction of surgical case lengths. Previous approaches to estimating surgical case length have included not only expert opinion, as in the case of LPCH, but also various types of statistical analysis of historical data (e.g., [11, 14, 21, 22, 24]). We used a prediction approach based on supervised learning, as we describe below, and a classification approach based on support vector machines that we do not describe here. Further details of our models can be found in [27].

We developed tree-based automated models to predict surgical case length: three automated models that use only patient and procedure characteristics and three semi-automated models that additionally use surgeon case length prediction as a feature. The simplest automated model, which we denote by DTR, is a single decision tree regressor. We also use a random forest regressor, denoted by RFR, and a set of gradient-boosted regression trees, denoted by GBR.

We designed and compared the models based on an operationally relevant loss function: the percentage of cases that are significantly mis-scheduled relative to their scheduled duration. Interviews with operating room staff and surgeons revealed that relatively minor differences between actual and scheduled case length do not cause significant disruptions. The impact of mis-scheduling depends on the scheduled length of the procedure. Consider a room in which 10 cases are scheduled, each 1 h long, and a room in which 2 cases are scheduled, each 5 h long. Scheduling errors of 15 min will significantly disrupt the performance of the first room but not of the second. We define a case as mis-scheduled if the actual duration differs from the scheduled duration by more than 25% of the scheduled case length or 15 min.

The automated models develop predictions based on the following features of patients and procedures: sex, weight, age, American Society of Anesthesiologists Physical Status Score (a score ranging from 1–6, indicating a range of health status from normal good health to brain dead), identity of the primary surgeon performing the procedure, location of the procedure (in an operating room or an ambulatory procedures unit), patient class (inpatient or outpatient), and procedure name.

The semi-automated prediction models use the above features and, in addition, use the surgeon's case length estimate as a feature. We denote these semi-automated models as DTR-S, RFR-S, and GBR-S, respectively, corresponding to the automated approach of models DTR, RFR, and GBR.

We tested the prediction models on the 10 most common procedures performed at LPCH from May 2014 through January 2015. The data set had a total of 3426 observations. We divided the data set into a training set and a tuning set of roughly equal size: 1640 observations in the training set and 1846 observations in the tuning set. For each procedure type, we compared the performance of our six prediction methods to two benchmarks: the historical average duration for that procedure and the expert estimate of the procedure duration (i.e., the value currently used when developing the operating room schedules).

As described in [27], our simple DTR model was not better than either benchmark. The other two automated methods, RFR and GBR, outperformed both benchmarks, with GBR performing better than RFR in most cases. The semi-automated prediction models DTR-S and RFR-S performed better than their automated counterparts, while GBR-S had approximately the same performance as GBR. These results suggest that the automated GBR method could be used as an adjunct to expert opinion when estimating surgical case length.

In partnership with LPCH's analytics provider, Qventus, we are now developing a system to implement the results of the work. To minimize disruption, surgical schedulers will continue to submit their case length estimates to EPIC, the LPCH electronic medical record, as they currently do. In real time, the estimate and all relevant information will be transmitted to Qventus for analysis by a variant of our algorithm. If the resulting predicted time differs significantly from the scheduled time, then Qventus will text-message and email an alert to the scheduler with a suggested time. The scheduler can then consult with the surgeon and modify the time appropriately. We will measure the performance of this system using the loss function described above (the percentage of cases that are significantly mis-scheduled relative to their scheduled duration).

Our model is readily generalizable to other hospitals and surgical centers where the relevant patient data, or at least the subset of the most useful features, are collected. For the GBR method, the most important features for prediction were procedure name, patient weight, and primary surgeon identity. For the GBR-S method, the most important features were primary surgeon identity and the surgeon's case length estimate, followed by procedure name and patient weight. The model can be implemented with minimal disruption to operating practices through automated notifications of potential case length mis-estimates.

3 Improving Operational Efficiency of the Perioperative Process

Our second project focuses on the stage between when a surgical procedure is scheduled and when the patient completes surgery and goes to the post-anesthesia care unit (PACU). Numerous interrelated factors contribute to operational inefficiency in this process. These include, for example, variability in patient preferences and clinical needs, lack of patient adherence to guidelines, equipment and supply availability, variability in provider practice, scheduling errors, inefficient resource allocation, and communication errors. In such a complex system, it may be difficult to identify the areas of the process where interventions would most reduce delays. For example, suppose that historical time stamps show that a patient entered the operating room later than scheduled. This may be because the patient came late to the hospital, or came on time and was delayed during preparation, or was prepared on time but delayed by the previous case running late.

We faced three major challenges in this project: identifying the factors that have the most significant impact on efficiency; determining how to address those problems without adverse impact on the wider system; and finding time for staff to implement change while they continue to operate in a very busy environment. We partnered with perioperative leaders, clinicians, and staff to address these challenges systematically.

To identify the factors that have the most significant impact on efficiency we created a detailed queuing representation of patient flows in the system, estimated the utilization and capacity of each step of the process using historical time stamp data, and measured the frequency and magnitude of delays associated with each process. We identified bottlenecks to determine how to make improvements that minimize disruption to the broader system. To achieve change without unduly taxing the perioperative staff, we augmented the current process with several automated notifications powered by data already in the hospital's electronic medical record. Below, we describe the design, implementation, and results of the project. Further details are provided in [10].

We first created a flow chart that maps the process starting from the days before surgery, through the activities in the pre-operative area, to the completion of surgery and the transfer of the patient to the PACU. The process flow is as follows: In the days leading up to surgery, dedicated nurses, nurse practitioners, and physicians from the hospital contact the patient's family to collect relevant information (e.g., allergies to medications) and instruct them to prepare for surgery (e.g., explain NPO guidelines and when to arrive to the hospital). After the patient is admitted and checked in, a nursing assistant takes the patient's vital signs, height and weight, and then brings the patient to a consult room where the patient sees a nurse practitioner and answers a number of questions. Then, depending on various circumstances such as the scheduled time of surgery, the availability of a nurse, or whether the operating room is running late, the nursing team will decide whether the patient needs to be brought back to the waiting area or can immediately see a nurse in the consult room. In both cases, the patient is eventually taken to the holding area, so that the nurse can complete the exam if necessary, and prepare the patient. At some point, the operating room sends a notification that they expect the patient to be ready to see a physician in the holding area within the next 20 min. The patient is then seen by an anesthesiologist and a surgeon for final preparations and then taken to the operating room. When the surgery nears completion, a nurse from the operating room contacts the PACU to reserve a bed.

We interviewed nurses, physicians, and staff members to identify problems and opportunities for improvement in the process flow. We identified several major causes for delays: patients do not follow guidelines to not eat before surgery; on the morning of surgery patients are missing needed paperwork; patients need an interpreter but no interpreter is available; no nurse is available for room turnover; or the PACU is full. In order to quantify the impact of each of these delays and to identify other sources of delays we used a queuing representation. We used two years of historical time stamps to estimate the capacity of each set of resources in the process, the rate of patient arrivals during busy periods, and the frequency of associated delays.

The detailed process mapping revealed that we could relieve the burden of nonclinical work on perioperative staff by using automated communication. We therefore implemented automated text-message reminders to patients to assist with the communication of surgical guidelines such as not eating the morning of a procedure. We redesigned the process for completing the pre-surgical documentation to be electronic rather than paper-based and to allow for automated alerts to notify physicians and staff when documentation was missing. Additionally, we implemented several other, similar interventions, as described in [10].

The queuing representation of the system revealed numerous days in which the PACU was the bottleneck causing surgical delays. Since neither decreasing the number of arrivals to the PACU nor increasing PACU capacity were feasible options, we considered ways to increase the rate of PACU service. A detailed study of the time stamps revealed that the notification from the operating room to the PACU, intended to be made 20 min before the patient is ready to exit the operating room, was frequently premature. When the PACU receives such a notification, a bed is reserved for the patient. Premature notifications thus effectively increase how long a patient occupies a PACU bed.

We implemented our recommended just-in-time operating room notifications to the PACU at the end of June 2016. LPCH's surgical caseload is highest during the summer months, as parents schedule procedures when children are not in school. In the two months following the policy change, the percentage of patients who arrived in the PACU more than 20 min after the notification fell from approximately 54% to approximately 22%. The number of cases with a PACU hold fell from 45 with an average length of 27 min in June, to 6 with an average length of 15 min in July and 16 with an average length of 14 min in August. The improvements were not related to surgical volume, as the average number of weekday cases using the PACU remained constant at approximately 30 per weekday in June, July, and August.

The transition to just-in-time bed requests increased the effective capacity of the PACU. Other institutions that track bed request time stamps and patient arrivals could reproduce our analysis to determine whether a similar intervention is appropriate in their setting, and could explore the potential for implementing automated notifications, using an approach similar to the one we have described.

4 Post-Anesthesia Care Unit Scheduling

Our third project focuses on reducing operating room delays caused by patients waiting for a bed to become available in the PACU. Such delays are a common, extensively studied problem. Research in this area can be categorized into projects that estimate the resources, such as beds or staff, necessary to minimize such delays [7, 18] and projects that develop methods to adjust the order of cases to reduce delays [9, 16]. Results of projects that involve adjusting the order of surgical cases have been largely negative, yielding conclusions such as, "Although effective, such methods can be impractical because of large organizational change required and limited equipment or personnel availability" [9] and "The uncoordinated decision-making of multiple surgeons working in different operating rooms can result in a sufficiently uniform rate of admission of patients into the PACU and holding that the independent sequencing of each surgeon's list of cases would not reduce the incidence of delays in admission or staffing requirements" [16].

We undertook a project to develop an easily implementable surgical procedure scheduling decision support tool that would create a level load of PACU bed and staff utilization. We tested its performance to estimate the resulting improvements and are in the process of implementing it. Below, we describe the current scheduling system at LPCH, the design of the decision support tool, and its implementation.

At LPCH, after a surgical procedure in the operating room, patients are sent to one of 10 recovery beds in the PACU. A patient cannot be assigned to the PACU unless a bed is free and the appropriate staff are available to supervise the patient's recovery. If a bed and needed staff in the PACU are not available when the patient's surgery finishes, the patient must wait in the operating room for a PACU assignment. This means that the next surgical procedure cannot begin and the next patient scheduled for surgery must continue to wait in the pre-operative area. These delays lead to inefficient use of operating rooms and staff as well as lowered patient satisfaction.

The current process to reduce PACU holds is as follows. Each day, starting in the morning, a scheduler 'builds' the operating room case schedule for the following day. Since each operating room is typically reserved for cases performed by a given surgical service, building the schedule consists primarily of determining the order of the cases in each room. The scheduler accounts for special considerations (e.g., patient characteristics, specialized equipment needs, or the need for more than one surgeon for a case) that may require certain cases to be performed at specific times. Each afternoon, by which time a preliminary schedule has been created, a meeting is held to estimate the corresponding demand for PACU and other beds, make changes, and finalize the schedule for the following day. If estimates based on the preliminary schedule suggest that the PACU will reach capacity at a given time of day, then the order of the procedures is shuffled to reduce the number of patients sent to the

PACU during that time. After this process is complete, schedulers call patients to notify them of their surgery time.

Our model uses as input the cases scheduled for the following day, patient information relevant to forecasting PACU length of stay, and the patient and surgeon information relevant to constraining when certain cases must be scheduled. The model uses a random-forest-based method to estimate the likely duration of each patient recovery in the PACU. An integer program is then used to determine the order in which the procedures should be scheduled in each operating room so as to minimize maximum overall PACU occupancy.

We used a discrete event simulation model to test the performance of the optimization. We validated the simulation by reproducing 6 months of historical PACU occupancy based on scheduled order of procedures, procedure durations, and recovery durations. After validation, we used the simulation to compare the historical PACU occupancy to that which would have resulted from scheduling with the optimization. We found that 60% of operating room days finished earlier with the optimized schedule compared to the actual schedule, suggesting that significant operational improvements can be achieved with the optimized scheduling system. We are now in the process of implementing the system at LPCH. For full details of the design, testing, and implementation, see [12].

This model has the potential to improve on the current process in several ways. First, the output of the model is automated; it does not require a scheduler to spend time creating the preliminary schedule. Second, the order of procedures is determined based on case-specific estimates of PACU occupancy whereas the current process assumes equal PACU recovery lengths for all cases. Third, the arrangement of cases is optimal for minimizing maximum PACU occupancy, a combinatorial result not readily achievable with the current manual process. Additionally, implementation of the model is minimally disruptive, as it produces as output a preliminary schedule that can be reviewed and revised at the current afternoon meeting.

Our model is readily generalizable to other healthcare institutions that finalize their operating room schedule after the majority of cases are scheduled. The data used to generate the forecast of PACU length of stay and to determine the order of procedures are routinely tracked by institutions with an electronic medical record. The preliminary schedule is easily modifiable by perioperative staff who can make changes to satisfy ad hoc constraints that are not captured in the model. We are currently working with Stanford Health Care (the adult hospital at Stanford) to explore implementation of our optimization model in their operating rooms.

5 Discussion

Many opportunities exist to improve the efficiency and effectiveness of healthcare services, and operations research can play an important role in supporting such efforts. The projects we are carrying out at Lucile Packard Children's Hospital Stanford demonstrate that a systematic, analytical approach to problems in hospital operating room management can help planners achieve significant operational improvements without expanding resources or unduly taxing hospital staff.

An important goal of our projects is to develop solution approaches that not only can be implemented in the specific setting under study, but that can also be generalized to other hospitals. The models we developed rely on data that are available in almost any electronic medical record system. With the recent expansion of electronic medical record systems, our models have potential usefulness in many settings.

The primary limitation of this work is that the projects described were designed and implemented at a single pediatric hospital. To ensure that these tools are generalizable, future work should implement the tools at a second hospital and report the necessary modifications. We are currently exploring this possibility at the Stanford adult hospital.

Another promising area for further research on improving operating room management is to determine the days on which elective surgery procedures are scheduled so that surgical bed occupancy is balanced. If one knew exactly what the demand for elective surgeries over time would be, then this problem could be solved as an integer program: an assignment problem with the goal of minimizing deviations from an average surgical bed occupancy level. However, future demands for elective surgery cannot all be known when assignments are being made. Thus, the challenge is to develop a prospective algorithm that achieves solutions close to those that would be found with perfect knowledge of future demand and the use of an optimization model.

References

1. Aleman, D., Brandeau, M.L., Carter, M.W., Scheinker, D.: (Draft) Healthcare systems engineering: an analytical approach. Springer Publishers, New York
2. Banditori, C., Cappanera, P., Visintin, F.: A combined optimization-simulation approach to the master surgical scheduling problem. IMA J. Manage. Math. **24**(2), 155–187 (2013)
3. Bravo, F., Levi, R., Ferrari, L.R., McManus, M.L.: The nature and sources of variability in pediatric surgical case duration. Paediatr. Anaesth. **25**(10), 999–1006 (2015)
4. Cappanera, P., Visintin, F., Banditori, C.: Comparing resource balancing criteria in master surgical scheduling: a combined optimisation-simulation approach. Int. J. Prod. Econ. **158**, 179–196 (2014)
5. Cappanera, P., Visintin, F., Banditori, C.: Addressing conflicting stake-holders priorities in surgical scheduling by goal programming. Flex. Serv. Manuf. J. Epub. (2016) (ahead of print)
6. Cardoen, B., Demeulemeester, E., Beliën, J.: Operating room planning and scheduling: a literature review. Eur. J. Oper. Res. **201**(3), 921–932 (2010)
7. Dexter, F., Epstein, R.H., Penning, D.H.: Statistical analysis of postanesthesia care unit staffing at a surgical suite with frequent delays in admission from the operating room: a case study. Anesth. Analg. **92**(4), 947–949 (2001)
8. Dexter, F., Blake, J.T., Penning, D.H., Lubarsky, D.A.: Calculating a potential increase in hospital margin for elective surgery by changing operating room time allocations or increasing nursing staffing to permit completion of more cases: a case study. Anesth. Analg. **94**(1), 138–142 (2002)
9. Dexter, F., Epstein, R.H., Marcon, E., de Matta, R.: Strategies to reduce delays in admission into a postanesthesia care unit from operating rooms. J. Perianesth. Nurs. **20**(2), 92–102 (2005)

10. Durand, A., Kim, H., Pei, F., Petersen, K.: A generalizable, systematic approach to improving perioperative efficiency. Working Paper (2017)
11. Eijkemans, M.J., van Houdenhoven, M., Nguyen, T., Boersma, E., Steyerberg, E.W., Kazemier, G.: Predicting the unpredictable: a new prediction model for operating room times using individual characteristics and the surgeon's estimate. Anesthesiology **112**(1), 41–49 (2010)
12. Fairley, M.C., Scheinker, D., Caruso, T.J., Brandeau, M.L.: Improving the efficiency of the operating room environment with a generalizable optimization and machine learning model. Working Paper (2017)
13. Hiltrop, J.: Modeling neuroscience patient flow and inpatient bed management. Ph.D. Thesis, Massachusetts Institute of Technology (2014)
14. Kayış, E., Khaniyev, T.T., Suermondt, J., Sylvester, K.: A robust estimation model for surgery durations with temporal, operational, and surgery team effects. Health Care Manag. Sci. **18**(3), 222–233 (2015)
15. Macario, A.: What does one minute of operating room time cost? J. Clin. Anesth. **22**(4), 233–236 (2010)
16. Marcon, E., Dexter, F.: An observational study of surgeons' sequencing of cases and its impact on postanesthesia care unit and holding area staffing requirements at hospitals. Anesth. Analg. **105**(1), 119–126 (2007)
17. Organisation for Economic Co-operation and Development (OECD): Focus on health spending: OECD health statistics 2015. https://www.oecd.org/health/health-systems/Focus-Health-Spending-2015.pdf (2016)
18. Schoenmeyr, T., Dunn, P.F., Gamarnik, D., Levi, R., Berger, D.L., Daily, B.J., Levine, W.C., Sandberg, W.S.: A model for understanding the impacts of demand and capacity on waiting time to enter a congested recovery room. Anesthesiology **110**(6), 1293–1304 (2009)
19. Shippert, R.D.: A study of time-dependent operating room fees and how to save $100 000 by using time-saving products. Am. J. Cosmetic. Surg. **22**(1), 25–34 (2005)
20. Smallman, B., Dexter, F.: Optimizing the arrival, waiting, and npo times of children on the day of pediatric endoscopy procedures. Anesth. Analg. **110**(3), 879–887 (2010)
21. Stepaniak, P.S., Heij, C., Mannaerts, G.H., de Quelerij, M., de Vries, G.: Modeling procedure and surgical times for current procedural terminology-anesthesia-surgeon combinations and evaluation in terms of case-duration prediction and operating room efficiency: a multicenter study. Anesth. Analg. **109**(4), 1232–1245 (2009)
22. Strum, D.P., Sampson, A.R., May, J.H., Vargas, L.G.: Surgeon and type of anesthesia predict variability in surgical procedure times. Anesthesiology **92**(5), 1454–1466 (2000)
23. Visintin, F., Cappanera, P., Banditori, C.: Evaluating the impact of flexible practices on the master surgical scheduling process: an empirical analysis. Flex. Serv. Manuf. J. **28**(1–2), 182–205 (2016)
24. Wright, I.H., Kooperberg, C., Bonar, B.A., Bashein, G.: Statistical modeling to predict elective surgery time. Comparison with a computer scheduling system and surgeon-provided estimates. Anesthesiology **85**(6), 1235–1245 (1996)
25. Zenteno, A.C., Carnes, T., Levi, R., Daily, B.J., Dunn, P.F.: Systematic OR block allocation at a large academic medical center: comprehensive review of a data-driven surgical scheduling strategy. Ann. Surg. **264**(6), 973–981 (2016)
26. Zenteno, A.C., Carnes, T., Levi, R., Daily, B.J., Price, D., Moss, S.C., Dunn, P.F.: Pooled open blocks shorten wait times for nonelective surgical cases. Ann. Surg. **262**(1), 60–67 (2015)
27. Zhou, Z., Miller, D., Master, N., Scheinker, D., Bambos, N., Glynn, P.: Detecting inaccurate predictions of pediatric surgical durations. In: Data Science and Advanced Analytics (DSAA), 2016 IEEE International Conference. IEEE, pp. 452–457 (2016)