The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)

Emergency Department Online Patient-Caregiver Scheduling

Hanan Rosemarin, Ariel Rosenfeld, Sarit Kraus
Bar-Ilan University
Israel

Abstract

Emergency Departments (EDs) provide an imperative source of medical care. Central to the ED workflow is the patient-caregiver scheduling, directed at getting the right patient to the right caregiver at the right time. Unfortunately, common ED scheduling practices are based on ad-hoc heuristics which may not be aligned with the complex and partially conflicting ED's objectives. In this paper, we propose a novel online deep-learning scheduling approach for the automatic assignment and scheduling of medical personnel to arriving patients. Our approach allows for the optimization of explicit, hospital-specific multi-variate objectives and takes advantage of available data, without altering the existing workflow of the ED. In an extensive empirical evaluation, using real-world data, we show that our approach can significantly improve an ED's performance metrics.

Introduction

Nearly half of all US hospital-associated medical care is delivered by Emergency Departments (EDs, also known as emergency rooms), making EDs a major source of medical care, especially for vulnerable populations (Lewin and Altman 2000; Marcozzi et al. 2018). EDs are faced with a dynamic flow of patients who present a wide variety of conditions, ranging from severe multiple percussive injuries and drug overdoses to common colds and cuts and scrapes, all of which seek fast and quality medical attention. Due to the variability in patients' conditions, as well as the limited availability of medical resources and their own variability (i.e., attending physicians, interns, etc), an efficient patient-caregiver scheduling process is needed, a process which is often referred to as triage (Christ et al. 2010).

Patient-caregiver scheduling is directed at getting the right patient to the right caregiver at the right time, given the ED's constraints. Specifically, given a preliminary evaluation of the patient upon arrival (commonly done by a triage nurse) and the available medical staff, a decision has to be made as to *when* the patient should receive treatment and by *which caregiver*. Today, the patient-caregiver scheduling process focuses almost entirely on assigning each patient a severity level using triage scales (e.g., between 1 and 5, 1 being the highest (Gilboy et al. 2012)), which in turn translates into an *upper bound* on the desired patient's waiting time, leaving the decision as to *when and which* caregiver should provide the treatment entirely in the hands of the triage nurse(s). Unfortunately, due to the time-critical environment, the multiple partially-conflicting objectives of the ED (as discussed next) and multiple interruptions, decisions are often inadequately made and are mainly based on ad-hoc heuristics and experience which need not necessarily fully align with optimizing the ED's objectives, e.g., (Franklin et al. 2011; Tanabe et al. 2004; ENA 2017). Specifically, while EDs have been computationally investigated for over 70 years (Saghafian, Austin, and Traub 2015), mainly focusing on modeling the patient arrival flow and required staffing levels, to the best of our knowledge, the patient-caregiver scheduling problem has yet to be addressed by computational means.

In order to address this shortcoming, we first model the patient-caregiver scheduling process as a novel online scheduling problem. Deriving an efficient scheduling policy to the corresponding problem is hard, therefore we remedy this hardness by introducing a deep-learning-based pairwise ranking approach which relies on ED-provided objectives and leverages real-world data. Our approach provides the ED with an effective and efficient scheduling policy targeted at optimizing the hospital-specific objectives given the hospital's available resources and expected patient flow.

In an extensive empirical evaluation, using real-world data and medical experts' input, we show that our proposed approach can significantly improve the patient-caregiver scheduling process, which can translate into better ED care for the greater good.

To ensure the validity of our approach and evaluation from a medical perspective, we recruited 4 medical caregivers (who did not co-author this paper) to follow this study: a triage nurse, a physician's assistant, an attending physician and an ED director, from three large hospitals in Israel. We refer to these caregivers as the *expert panel* in this study.

The ED Workflow and Purpose

While different EDs deploy slightly different modes of operations, a basic work-flow is common to most modern EDs (Sinreich and Marmor 2005; Bedoya-Valencia and Kirac 2016), as depicted in Figure 1. In words, when a patient arrives at the ED, her first stop would be the *triage sta-*

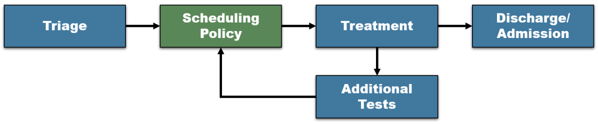

Figure 1: The common ED work-flow

tion, where she would receive a *severity rank*. Based on a **scheduling policy**, which is in the focus of this study, the patient would then continue to a (first) treatment/examination by a caregiver. If the condition is appropriately diagnosed and/or treated, the patient may be released or admitted to a hospital ward. Otherwise, additional lab tests (e.g., CT, bloodwork) would be needed which in turn would require a re-evaluation of the patient's severity rank and a second treatment/examination, possibly by a different caregiver. According to our expert panel, it is extremely rare for a patient to have more than 2 cycles of treatment before she is discharged or admitted.

The principal purpose of the ED is to ensure that patients receive the level and quality of care appropriate to their clinical needs and that the ED resources are most usefully applied to this end (FitzGerald et al. 2010). Unfortunately, explicitly quantifying the above purpose is highly complex (Schuur et al. 2013), which often leads hospitals and governmental agencies to define multiple, often partially conflicting, objectives (Welch et al. 2011). These objectives primarily focus on minimizing the following measures: 1) risk of adverse consequences to patients (Calder et al. 2015) (e.g., misdiagnosis, inappropriate medication); 2) ED overcrowdedness (Cowan and Trzeciak 2004); 3) interruptions to caregivers (Westbrook et al. 2010); 4) patients' wait times (Mowen, Licata, and McPhail 1993); and 5) patients' length of stay in the ED (Trzeciak and Rivers 2003). In this work, we focus on these measures.

Patient-Caregiver Scheduling

We start by modeling the two main sets of actors in the ED, patients and caregivers, and their interaction. Our modeling is based on existing literature and common clinical practices as prescribed by the expert panel.

Patients. A patient $p_i \in P$ is represented as a pair $\langle Severity, Injury \rangle$ where $Severity$ is defined using a common triage scale (such as the popular Emergency Severity Index (ESI) (Tanabe et al. 2004))[1] and $Injury$, which defines the type of injury or condition based on the patients' symptoms (e.g., *Orthopedic*, *Internal*, etc.). According to clinical guidelines, p_i is associated with the maximal time she is permitted to wait for the initial caregiver's treatment, denoted $\overline{t^1_{p_i}}$ and, if needed, the second caregiver's treatment $\overline{t^2_{p_i}}$. Patients may arrive at time t to the ED based on an estimated distribution $\mathcal{D}_{p_i}(t)$, commonly assumed to follow a non-homogeneous Poisson process (Whitt and Zhang 2017). We assume that upon arrival to the ED, p_i's characteristics are correctly identified by the triage nurse. After leaving the triage station (or when p_i's test results arrive, see Figure 1), she is scheduled to meet one of the caregivers c_j.

Caregivers. A caregiver $c_j \in C$ is represented as a pair $\langle Seniority, Specialty \rangle$: $Seniority$ is defined based on the caregiver's qualifications over a discrete set[2], and a $Specialty$ which indicates if the caregiver has "special training" in a specific injury type defined over the same set of injury types which characterize the patients (*NONE* otherwise). As a result, different caregivers may have different required treatment time and varying levels of care quality. The set of available caregivers, as well as their characteristics, is assumed to be known in advance and does not change during a shift.

Scheduling Objective. We model the primary ED objectives as follows:

1) **Minimizing risk of adverse consequences.** Each examination ($e \in \{1, 2\}$) performed by c_j on p_i has some risk of an adverse consequence. For example, a physician's assistant with no specific specialty may be very well equipped to perform a first examination of minor orthopedic injuries with only a minimal risk of an adverse consequence, while severe head injuries should be examined by a qualified physician. We use $AC(p_i, c_j, e)$ as an indicator of whether p_i was examined by c_j in her e^{th} examination. $risk_{i,j,e}$ denotes the risk of adverse consequences associated with such an examination.

2) **Minimizing patients' waiting time**. Each patient p_i, has to wait for her first (and second) examination for $WT(p_i, e)$ minutes. Given p_i's $\overline{t^1_{p_i}}$ and $\overline{t^2_{p_i}}$ (as defined by the triage scale), the ED seeks to minimize the wait time and avoid exceeding the wait time limits. The penalty for exceeding the limits is provided by exc_{i,e,δ_t} where δ_t is the excess wait time.

3) **Minimizing patients' length of stay**. Each patient p_i spends $LOS(p_i)$ minutes from the time she arrives at the ED to the time she is discharged or admitted to a hospital ward. This $LOS(p_i)$ includes the time p_i waits for examinations, the treatment time needed by the caregiver, denoted $CE(c_j) \cdot TT(p_i, e)$, and (if needed) lab test time $LT(p_i)$ between the two examinations. $TT(p_i, e)$ denotes the nominal treatment time and $CE(c_j) \geq 1$ denotes the caregiver's time efficiency factor, capturing the relative "examination speed" which varies between caregivers.

4) **Minimizing over-crowdedness**. At any point in time, one can measure the number of patients currently waiting and being treated in the ED, denoted $\kappa(t)$, where t indicates continuous time.

5) **Minimizing interruption to caregivers.** Unfortunately, in some (extreme) cases a caregiver may be asked to stop the treatment of one patient in order to treat another. This *preemption* may be very costly. The number of preemptions during p_i's e^{th} examination is denoted as $PC(p_i, e)$. The penalty for each interruption is given by $pre_{i,e}$.

[1]Using the ESI, each patient is assigned a number between 1 and 5 representing the acuity of her condition (with 1 being the most acute).

[2]In Israel (as in most countries), c's $Seniority$ is classified to one the following (from lowest to highest): physician's assistant, intern, resident and attending physician.

Notation	Meaning
t	Time.
$p_i \in P$	Patient.
$c_j \in C$	Caregiver.
$AC(p_i, c_j, e)$	Indicator whether c_j is assigned to p_i's e^{th} examination.
$WT(p_i, e)$	p_i's waiting time for her e^{th} examination.
$TT(p_i, e)$	Time (nominal) required for the e^{th} examination of patient p_i.
$LT(p_i)$	Time required for lab tests of p_i.
$CE(c_j)$	A type-c caregiver time efficiency factor.
$\kappa(t)$	#patients in the ED at time t.
$PC(p_i, e)$	#interruptions to p_i's e^{th} examination.

Table 1: Summary of key notations.

The ED must choose, for each patient p_i, which caregiver c_j will provide the examination/treatment e and at what time t. Let $y_j^i[e,t]$ be indicator decision variables denoting that patient p_i is assigned to caregiver c_j for her e^{th} examination at time t. We assume that the ED is evaluated based on some metric, defined by stakeholders and governmental agencies, over the above 5 objectives, e.g., using a linear objective which summarizes the weighted objectives over all patients, examinations and time as proposed in Equation 1. Note however, that the objective need not be linear and, in the interest of generality, we do not assume it to be in the following sections.

$$\begin{aligned} &\underset{y_j^i(e,t)}{Minimize} \sum_{p_i} \sum_{e\in\{1,2\}} \sum_{c_j} (\int_t y_j^i[e,t] \\ &(\alpha_1 AC(p_i,c_j,e) risk_{i,j,e} + \alpha_5 PC(p_i,e) pre_{i,e})\ dt + \\ &\alpha_2 WT(p_i,e) exc_{i,e,\delta_t} + \alpha_3 LOS(p_i)) + \alpha_4 \int_t \kappa(t)\ dt) \end{aligned} \tag{1}$$

Table 1 summarizes the paper's notations.

The patient-caregiver scheduling problem is akin to the well studied job shop scheduling problem of unrelated machines with preemption, with the analogy of caregivers to machines and patients to incoming jobs, as discussed later in this paper. This problem is known to be NP-hard (Du and Leung 1991). Thus, in the next section, we propose a novel machine-learning based approximation algorithm, tailored to our setting.

Learning-Based Scheduling

In order to tackle the patient-caregiver scheduling challenge outlined above, we propose a novel machine learning-based approach we term as LEARNING-BASED SCHEDULING (LBS). LBS is aimed at approximating the idealized optimal offline schedule, which is informed of the entire flow of patients and their characteristics in advance.

LBS works as follows: First, LBS creates a set of offline patient-caregiver scheduling problems based on past data or patient arrival models learned from actual data (e.g., (Whitt and Zhang 2017)). Then, using an appropriate optimization algorithm, each instance is optimally solved. The optimized solution set is then used to generate a set of training examples to train a deep-learning ranking model which is used in the online setting. See Algorithm 1. We first discuss how the optimal solution for each offline instance is obtained. Then, we present the training of our machine learning algorithm and its architecture which, in turn, translates into an online scheduling algorithm.

Algorithm 1 The Learning-Based Scheduling Process

1: Create a set of patient-caregiver offline optimization problems based on past patient flow data or patient flow distribution.
2: Solve the offline optimization problems.
3: Translate each assignment in each schedule into training instances.
4: Train ranking model.
5: Use the resulting online scheduling policy.

Optimal Offline Scheduling

When the patients' arrival times and characteristics are known in advance, optimal patient-caregiver scheduling can be derived over a discrete finite horizon $t = 0, \dots, T$ using the following Mathematical Problem (MP):

$$\underset{y_j^i[e,t]}{Minimize}\ \text{ED Objective (i.e., Eq. 1)} \tag{2}$$

$$\text{s.t} \sum_i \sum_e y_j^i[e,t] \le 1 \quad \forall j,t \tag{3}$$

$$\sum_j y_j^i[e,t] \le 1 \quad \forall i,e,t \tag{4}$$

$$\rho_{i,j,e,t} = y_j^i[e,t](\neg y_j^i[e,t-1]) \quad \forall i,j,e,t \ge 1 \tag{5}$$

$$\chi_{i,e,t} = t \cdot \sum_j \rho_{i,j,e,t} \quad \forall i,e,t \ge 1 \tag{6}$$

$$st_{i,e} = \min_t \{\chi_{i,e,t}\} \quad \forall i,e \tag{7}$$

$$\phi_{i,e,t} = t \sum_j (\neg y_j^i[e,t]) y_j^i[e,t-1] \quad \forall i,e,t \ge 1 \tag{8}$$

$$et_{i,e} = \max_t \{\phi_{i,e,t}\} \quad \forall i,e \tag{9}$$

$$st_{i,e} < et_{i,e} \quad , \quad st_{i,2} \ge et_{i,1} \quad \forall i,e \tag{10}$$

$$\text{arrival}_{i,2} = et_{i,1} + LT(p_i) \quad \forall i \tag{11}$$

$$st_{i,e} \ge \text{arrival}_{i,e} \quad \forall i,e \tag{12}$$

$$LOS(p_i) = et_{i,2} - \text{arrival}_{i,1} \quad \forall i \tag{13}$$

$$WT(p_i,e) = st_{i,e} - \text{arrival}_{i,e} \quad \forall i,e \tag{14}$$

$$\sum_j \sum_t y_j^i[e,t]/CE[j] = TT[p_i,e] \quad \forall i,e \tag{15}$$

$$PC(p_i,e) = \sum_t \sum_i \rho_{i,j,e,t} - 1 \quad \forall i,e \tag{16}$$

$$AC(i,j,e) = I(\sum_t \rho_{i,j,e,t} > 0) \quad \forall i,j,e \tag{17}$$

The MP consists of an ED-specific objective function (e.g., Eq. 1) and the following constraints: Eq. 3 and 4 enforce that at most one patient is treated at a time by each caregiver and, similarly, at most one caregiver can treat a patient at a given time. Eq. 5, 6 and 7 extract the treatment start time and Eq. 8 and 9 extract the treatment end time. Eq. 10 and 12 enforce a valid treatment duration while Eq. 11 extracts the time in which a patient becomes available for her second treatment. *Note that* arrival$_{i,1}$ *is assumed to be given in the offline setting*. Eq. 13 and 14 extract each patient's LOS and WT, respectively. Next, Eq. 15 makes sure that the time caregivers are assigned to a patient is appropriate and Eq. 16 extracts the preemptions that took place. For simplicity, the above MP assumes each patient is treated twice. Our implementation, which allows for up to two treatments per patient, is available as a Mixed Integer Linear Program (MILP) at https://goo.gl/rXaBRh. The MILP was instantiated with the help of the expert panel.

Learning to Schedule

The main task of an online scheduling algorithm in our setting is to select, or rank, a scheduling pair (patient-caregiver matching) over all other possible pairs. The task of learning a function to select the best match is also common in information retrieval and is called "Learning to Rank" (Liu 2009), where documents are ranked based on their relevance to a given query.

Therefore, we draw on the intuition from information retrieval literature and adapt and extend the "Learning to Rank" approach to the online ED scheduling setting.

Using the set of optimized solutions generated offline, we identify the times in which a new patient arrives or when treatment of a patient is completed. For each such case, we create all scheduling pairs consisting of the selected assignments according to the optimized solution ($\langle p_i^\star, c_j^\star\rangle$ or $\langle p_i^\star, WaitRoom\rangle$) coupled with any other assignment option which was not selected (i.e., $\langle p_i^\star, c_j\rangle$ or $\langle p_i, c_j^\star\rangle$ and $WaitRoom$ options). For simplicity, from this point onwards, we will consider the assignment to the $WaitRoom$ as a dummy caregiver which can support an infinite number of patients but does not provide any treatment. The resulting pairs are used as training data for a supervised ranking machine learning algorithm as we will discuss next. In other words, we use the set of optimized solutions to generalize and mimic the optimal decisions made in the offline settings.

With the help of the expert panel, we define a feature vector that combines a description of the patient and the caregiver's current state as shown in Table 2.

Feature Vector		
Patient	severity (Following ESI)	$1/5, 2/5, ..5/5$
	injury	one-hot vector
	remaining treatment time	in minutes
	wait time	in minutes
Caregiver	seniority	$1/4, 2/4, ..4/4$
	specialization	one-hot vector
	status	0-idle; severity of patient
	idle time	in minutes

Table 2: Combined Patient-Caregiver Feature Vector

Inspired by the neural network ranking approach in information retrieval (Rigutini et al. 2011), we develop a new Deep Neural Network (DNN) architecture targeted at learning to rank among assignments based on the created dataset of pairs discussed above.

Specifically, our DNN is composed of two identical subnetworks, with shared weights. The network architecture is shown in Figure 2.

The anti-symmetric nature of the network is built by sharing weights, as can be demonstrated for the connection between the input and the first hidden layer:

$$\overrightarrow{w}^{\,1}_{i,1} = w(\overrightarrow{X} \to \overrightarrow{H_{1,1}}) = w(\overrightarrow{Y} \to \overrightarrow{H_{1,2}})$$
$$\overrightarrow{w}^{\,2}_{i,1} = w(\overrightarrow{X} \to \overrightarrow{H_{1,2}}) = w(\overrightarrow{Y} \to \overrightarrow{H_{1,1}})$$

The bias term of both parts of the first hidden layer is also shared. Thus, the two output vectors of the first hidden layer are:

$$\overrightarrow{H_{1,1}} = \tanh(\vec{w}^{1}_{i,1} \cdot \overrightarrow{X} + \vec{w}^{2}_{i,1} \cdot \overrightarrow{Y} + \overrightarrow{b_1})$$
$$\overrightarrow{H_{1,2}} = \tanh(\vec{w}^{1}_{i,1} \cdot \overrightarrow{Y} + \vec{w}^{2}_{i,1} \cdot \overrightarrow{X} + \overrightarrow{b_1})$$

The rest of the layers share weights and connections in a similar fashion with their appropriate activation functions. Complete technical details are available in our code.

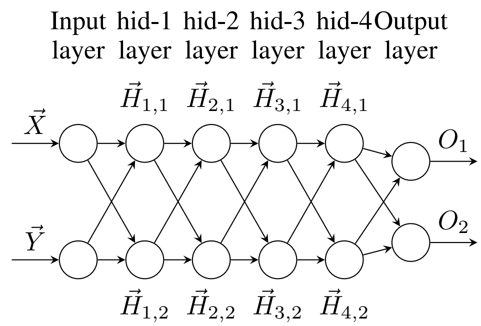

Figure 2: DNN scheduling network.

This architecture has the following properties:

1. *Reflexivity*: for identical input vectors, the network produces identical outputs; and
2. *Antisymmetry*: for input vectors x, y, if $x \succ y$ (reads "x is preferred over y according to the DNN") then for input vectors y, x, we get $y \prec x$ and vice versa.

These properties make the network well suited to learn pairwise ranking functions.

The trained DNN is used in our online scheduling algorithm as follows:
At each scheduling event, i.e. patient arrival, or caregiver completion of treatment, all possible assignments are compared to each other. Each comparison is worth one point to the higher ranking assignment. Using a majority vote with random tie breaking, the caregiver-patient assignment is selected.

Evaluation

Setup

Case study. We consider one of the largest hospitals in Israel: *Rambam* hospital in Haifa, which published extensive data on its ED, specifically its patient flow and caregivers' statistics. As a result, Rambam hospital has been the subject of various research efforts aimed at modeling the patient arrival process and caregivers' and lab tests' required time, among others. Using existing literature, and with the help of the expert panel (who are familiar with Rambam hospital's practices and Israel's guidelines), we instantiate all functions and constants of our model. Specifically, a linear objective function as proposed in Equation 1 was adopted and instantiated by our expert panel.

Rambam hospital works in a 3-shift workday. For this evaluation, we focus on perhaps the most challenging shift – the *night shift*. The night shift takes place from 23:00 to 7:00 during which only 2 to 4 caregivers, of various seniority and specialties, man the ED, making the assignment extremely complex. Complete technical details as to the objective function's parameters, expected treatment times, etc. are provided in the code.

We examine two scenarios – *normal* patient flow and *heavy* patient flow. The normal patient flow is provided in the literature (Whitt and Zhang 2017) whereas the heavy patient flow is derived by multiplying the distribution parameters, resulting in twice the number of patients on expectancy. We train two LBS DNNs, one for the normal patient flow case, denoted LBS_N, and one for the heavy patient flow case, denoted LBS_H. we randomly generate 500 scheduling instances for each condition, which in turn were optimally solved using the Gurobi solver (Gurobi Optimization 2018). The results are translated into two training sets for LBS_N and LBS_H. The resulting scheduling policies, alongside the baseline FCFSwU policy, will be evaluated next.

Baseline. In order to evaluate our approach compared to existing ED practices, we will compare our approach to the First-Come-First-Served-with-Urgencies (*FCFSwU*) heuristic which, according to our expert panel, is the backbone of most ED scheduling decisions, including those at Rambam hospital. FCFSwU works as follows: patients of severity levels 3-5 are treated as a single "non-urgent" type and are admitted in a first-come-first-served fashion to a caregiver who specializes in the relevant injury type or to a caregiver with no specialty. Specifically, a patient would not be assigned to a specialist who specialized in a different injury type. Patients of severity levels 1 or 2 are treated as a single "urgent" type and, upon arrival, the most senior specialized caregiver who is not already treating another urgent patient is called (or interrupted) in order to provide the needed treatment.

Results

We evaluate both the LBS_N and LBS_H against the FCFSwU heuristic on a series of 100 eight-hour scheduling instances sampled according to the parameters discussed earlier. Both approaches were evaluated using the expert panel objective function available in our code.

Interestingly, for all 100 sampled instances, the LBS approach brings about a better scheduling performance compared to the FCFSwU heuristic. The difference is statistically significant, for both the normal and heavy patient flow conditions, using paired samples t-test, $p < 0.05$.

We further evaluate the results based on the five major ED objectives. We encounter the following results:

1. *Risk of Adverse Consequences*: the average risk of adverse consequences is reduced by 10% (normal flow) and 15% (heavy flow) compared to FCFSwU.
2. *Wait times*: the average wait time was slightly reduced by an average of 20 seconds per patient, across both patient flow conditions, compared to FCFSwU.
3. *Length of stay*: the average length of stay was reduced by 5% (normal flow) and 11% (heavy flow) compared to FCFSwU.
4. *Crowdedness*: no significant differences.
5. Interruptions: the most prominent difference was measured in the number and cost of preemptions. Specifically, a treatment is approximately **10 times more likely to be interrupted** using the FCFSwU compared to LBS. By weighting the interruptions by their associated penalties, we see a penalty that is **20 times** higher per shift.

Table 3 summarizes the results.

Criteria	Normal Load	Heavy Load
Risk	10%	15%
Wait times	marginal	marginal
Length of stay	5%	11%
Crowdedness	-	-
Interruptions	95%	90%

Table 3: Marginal improvement of the LBS approach compared to FCFSwU.

Contrary to what the authors initially expected, *there is no apparent trade-off between the LBS approach and the FCFSwU heuristic*. Specifically, based on the results, the use of LBS improved 4 out of the 5 performance metrics while having no impact on the fifth. It is important to note that the crowdedness measure was assigned a very low priority by the expert panel, possibly explaining this result.

Discussion

The results for Rambam hospital indicate that the use of the LBS approach encompasses significant benefits compared to the common practice. Specifically, by leveraging real-world data and an explicitly defined ED objective, the LBS approach brings about a better suited scheduling policy to the ED environment.

By explicitly considering the ED-specific objective and characteristics (e.g., patient flow, available caregivers), the LBS further allows stakeholders to experiment and investigate different patient flows, capacity planning, room and staff shortages, etc. which are at the core of ED research and practice (Saghafian, Austin, and Traub 2015). These may also include "softer" objectives (not usually specified by medical metrics) such as the fairness of a caregiver's workload. Such an investigation could be accomplished by simply changing the modular functions and constants of our model and re-running the process described in Algorithm 1.

However, when presenting a new approach, such as the LBS, it is worth discussing limitations. First, the results demonstrate an interesting trade-off between scheduling quality and development time. While the LBS allows for better scheduling performance, arriving at the LBS policy requires the construction of the scheduling instances, solving them and training a DNN, which in turn takes significantly more time compared to the easy-to-deploy heuristic commonly applied today. It is important to note that the LBS training is performed offline, thus in deployment of the resulting policy, no runtime differences are encountered. Second, deploying an automated scheduling policy such as the one proposed in this study may encounter deployment challenges or even resistance from the medical staff. We plan to focus on translating the LBS approach into a real-world application in future work.

It is important to note that the evidence that the ED patient arrival process can be modeled (e.g (Whitt and Zhang 2017)) provides an assurance that a scheduling policy can be derived - but the model itself is not constructive. Namely, the model cannot be directly mapped to a scheduling policy. In this study we adopted a machine-learning approach to tackle this challenge.

Related Notions in Scheduling

The patient-caregiver scheduling problem can be seen as a special type of Job Shop Scheduling (JSS), which is a fundamental problem in computer science and operations research (Pinedo 2015). The scheduling of jobs to machines in JSS is analogous to the allocation of patients to caregivers in the ED setting. However, the unique characteristics of the ED environment, such as the complex and partially conflicting objectives of the ED and the heterogeneous patients and caregivers available, push us to address a novel version of the JSS problem, most relevant for medical scheduling.

In JSS, considering the value of completing a job is usually much simpler than in the ED environment. For example, Zheng and Shroff (2016) address the scheduling of computer jobs to a cloud cluster in a setting where tasks arrive online but give some partial value for partial execution. Naturally, in the ED setting, all patients must be *fully* treated and therefore we do not allow for partial treatment, except for the preemption case which only occurs in extreme cases. The notion of preemption was also investigated in Doucette et al. (2016), addressing the assignment of tasks to agents in an online fashion. Neither of the above studies addresses the possible differences in valuation and the completion rate between the agents/machines in contrast to how our model is capable of differing between caregivers (i.e., seniority and specialty). From a machine learning perspective, task scheduling using classification was recently investigated (e.g., (Tripathy, Dash, and Padhy 2015; Panda, Mohapatra, and Panigrahi 2015)). Most work in this realm focuses on the *offline* scheduling of tasks with dependencies (e.g., temporal dependencies) and deadlines, while we focus on the online scheduling of independent tasks (patients) where there is no strict deadline but rather a desired upper bound on the patient's waiting time. More relevant to this work is (Gombolay et al. 2016) which takes a Learning from Demonstration (LfD) approach—that is, learning human-quality heuristics based on demonstrations—to a scheduling problem without varying values for tasks. Similar to our proposed approach, the authors use a pairwise ranking function, however, while the authors try to mimic a human-quality scheduling policy, we follow optimal solutions, thereby overcoming the inherent suboptimality of existing human-generated ED policies.

From the medical scheduling perspective, a few works have addressed patient scheduling concerns. Notably, Peretz et al. (2013) focus on the nuclear medicine domain, and take a two-stage stochastic integer programming approach to scheduling patients that require multi-step tests, e.g., a patient arrives with three tests to be performed that have to be performed sequentially, with the restriction that any individual task cannot be paused once it has begun. In the proposed model, once a patient's tasks are scheduled (in the future), they cannot be changed or interrupted, a constraint we do not have in the ED setting. There are also various techniques for scheduling a caregiver to shifts under different constraints (Erhard et al. 2017). To the best of our knowledge, none have addressed the patient-caregiver scheduling problem to date.

Conclusions

This paper introduces a novel framework and solution for the patient-caregiver scheduling problem in EDs. First, based on existing literature, real-world data and 4 medical experts (including an ED director), we model the ED workflow, characteristics and objectives in terms of a new online scheduling problem. Next, we propose the Learning-Based Scheduling (LBS) approach, which leverages optimal offline solutions to sampled scheduling instances, in order to learn an efficient online scheduling policy. In an extensive empirical evaluation, we demonstrate the benefits of the LBS approach compared to existing ED practices. All code and data used in this work are available in https://goo.gl/rXaBRh.

We plan to extend this work in two main directions: First, since many hospitals also operate as training centers, there may also be an added value for assigning multiple caregivers of different seniority to treat the same patient. Therefore, we plan to extend our model to incorporate these complex allocation objectives. Second, additional medical environments such as the online assignment of scans to radiologists will be investigated.

We hope that this study will encourage other researchers to tackle the important and challenging task of promoting quality and timely medical care.

Acknowledgements

This research was supported by a grant from the Ministry of Science & Technology, Israel & the Ministry of Education, Science, Research and Sport of the Slovak Republic.

References

Bedoya-Valencia, L., and Kirac, E. 2016. Evaluating alternative resource allocation in an emergency department using discrete event simulation. *Simulation* 92(12):1041–1051.

Calder, L.; Pozgay, A.; Riff, S.; Rothwell, D.; Youngson, E.; Mojaverian, N.; Cwinn, A.; and Forster, A. 2015. Adverse events in patients with return emergency department visits. *BMJ Qual Saf* 24(2):142–148.

Christ, M.; Grossmann, F.; Winter, D.; Bingisser, R.; and Platz, E. 2010. Modern triage in the emergency department. *Deutsches Ärzteblatt International* 107(50):892.

Cowan, R. M., and Trzeciak, S. 2004. Clinical review: emergency department overcrowding and the potential impact on the critically ill. *Critical care* 9(3):291.

Doucette, J. A.; Pinhey, G.; and Cohen, R. 2016. Multiagent resource allocation for dynamic task arrivals with preemption. *ACM Transactions on Intelligent Systems and Technology (TIST)* 8(1):3.

Du, J., and Leung, J. Y.-T. 1991. Minimizing the number of late jobs on unrelated machines. *Operations Research Letters* 10(3):153–158.

ENA. 2017. Triage qualifications and competency (position paper). *Emergency Nurses Association*.

Erhard, M.; Schoenfelder, J.; Fügener, A.; and Brunner, J. O. 2017. State of the art in physician scheduling. *European Journal of Operational Research*.

FitzGerald, G.; Jelinek, G. A.; Scott, D.; and Gerdtz, M. F. 2010. Emergency department triage revisited. *Postgraduate medical journal* 86(1018):502–508.

Franklin, A.; Liu, Y.; Li, Z.; Nguyen, V.; Johnson, T. R.; Robinson, D.; Okafor, N.; King, B.; Patel, V. L.; and Zhang, J. 2011. Opportunistic decision making and complexity in emergency care. *Journal of biomedical informatics* 44(3):469–476.

Gilboy, N.; Tanabe, P.; Travers, D.; Rosenau, A. M.; et al. 2012. Emergency severity index (esi): a triage tool for emergency department care, version 4. *Implementation handbook* 12–0014.

Gombolay, M.; Jensen, R.; Stigile, J.; Son, S.-H.; and Shah, J. 2016. Apprenticeship scheduling: Learning to schedule from human experts. AAAI Press/International Joint Conferences on Artificial Intelligence.

Gurobi Optimization, L. 2018. Gurobi optimizer reference manual.

Lewin, M. E., and Altman, S. 2000. *America's health care safety net: intact but endangered*. National Academies Press.

Liu, T.-Y. 2009. Learning to rank for information retrieval. *Information Retrieval* 3(3):225–331.

Marcozzi, D.; Carr, B.; Liferidge, A.; Baehr, N.; and Browne, B. 2018. Trends in the contribution of emergency departments to the provision of hospital-associated health care in the usa. *International Journal of Health Services* 48(2):267–288.

Mowen, J. C.; Licata, J. W.; and McPhail, J. 1993. Waiting in the emergency room: how to improve patient satisfaction. *Marketing Health Services* 13(2):26.

Panda, S.; Mohapatra, P. K.; and Panigrahi, S. P. 2015. A new training scheme for neural networks and application in non-linear channel equalization. *Applied Soft Computing* 27:47 – 52.

Pérez, E.; Ntaimo, L.; Malavé, C. O.; Bailey, C.; and McCormack, P. 2013. Stochastic online appointment scheduling of multi-step sequential procedures in nuclear medicine. *Health Care Management Science* 16(4):281–299.

Pinedo, M. 2015. *Scheduling*. Springer.

Rigutini, L.; Papini, T.; Maggini, M.; and Scarselli, F. 2011. SortNet: Learning to rank by a neural preference function. *IEEE Transactions on Neural Networks* 22:1368–1380.

Saghafian, S.; Austin, G.; and Traub, S. J. 2015. Operations research/management contributions to emergency department patient flow optimization: Review and research prospects. *IIE Transactions on Healthcare Systems Engineering* 5(2):101–123.

Schuur, J. D.; Hsia, R. Y.; Burstin, H.; Schull, M. J.; and Pines, J. M. 2013. Quality measurement in the emergency department: past and future. *Health Affairs* 32(12):2129–2138.

Sinreich, D., and Marmor, Y. 2005. Emergency department operations: the basis for developing a simulation tool. *IIE transactions* 37(3):233–245.

Tanabe, P.; Gimbel, R.; Yarnold, P. R.; Kyriacou, D. N.; and Adams, J. G. 2004. Reliability and validity of scores on the emergency severity index version 3. *Academic emergency medicine* 11(1):59–65.

Tripathy, B.; Dash, S.; and Padhy, S. K. 2015. Dynamic task scheduling using a directed neural network. *Journal of Parallel and Distributed Computing* 75:101 – 106.

Trzeciak, S., and Rivers, E. P. 2003. Emergency department overcrowding in the united states: an emerging threat to patient safety and public health. *Emergency medicine journal* 20(5):402–405.

Welch, S. J.; Asplin, B. R.; Stone-Griffith, S.; Davidson, S. J.; Augustine, J.; Schuur, J.; et al. 2011. Emergency department operational metrics, measures and definitions: results of the second performance measures and benchmarking summit. *Annals of emergency medicine* 58(1):33–40.

Westbrook, J. I.; Woods, A.; Rob, M. I.; Dunsmuir, W. T.; and Day, R. O. 2010. Association of interruptions with an increased risk and severity of medication administration errors. *Archives of Internal medicine* 170(8):683–690.

Whitt, W., and Zhang, X. 2017. A data-driven model of an emergency department. *Operations Research for Health Care* 12:1–15.

Zheng, Z., and Shroff, N. B. 2016. Online multi-resource allocation for deadline sensitive jobs with partial values in the cloud. In *IEEE International Conference on Computer Communications (INFOCOM)*.