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 (Altman, Lewin, and others 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 most acute (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 scheduling has yet to be addressed by computational means.

Approach

We address the problem by modeling 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. We refer to these caregivers as the *expert panel* in this study.

Our machine learning-based approach, which we term as LEARNING-BASED SCHEDULING (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 a mixed integer linear program formulation of the optimization algorithm, each instance is optimally solved using a MIP solver (Gurobi Optimization 2018). 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.

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: Trained ranking model.
5: Use the resulting online scheduling policy.

The training data is extracted from the optimal solutions by identifying 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. 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 1.

Feature Vector		
Patient	severity (by 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 1: Combined Patient-Caregiver Feature Vector

For training, we use a deep neural network with an anti-symmetric shared weights architecture which is intrinsically reflexive and anti-symmetric, thus suitable to learn pair-wise ranking.

The neural network is used in the online algorithm to compare and select the suitable patient-to-caregiver assignment.

We evaluated our approach using real-world patient and caregiver data. The results proved our approach to be superior to the heuristics used by our expert panel both in terms of quality of care metrics and in terms of time metrics (e.g. patient length of stay in ED).

Future Work

We plan to extend this work, working with EDs, 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 allow for these complex allocations. Second, additional medical environments such as the online assignment of scans to radiologists will be investigated.

References

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

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

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

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.

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

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.

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.

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.

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