NUS National University of Singapore | NUS Global Asia Institute

Newsletter

a quarterly report on projects from the NUS Global Asia Institute Issue 10 July 2015

Data-Driven Patient Scheduling in Emergency Departments

by Meilin ZHANG and Shuangchi HE

Emergency department (ED) crowding and its resulting delays are global issues that have received considerable attention from governments, the public, media, and academic communities. ED crowding compromises the quality of emergency care and puts patients at greater risk of treatment errors. Numerous studies have revealed an association between crowding and increased morbidity and mortality in EDs. For hospitals, ED crowding damages their public reputations and incurs revenue loss. Widespread crowding also impedes their ability to achieve national safety and quality goals, compromises the healthcare system, and limits the regional capacity for disaster response.

In Singapore, ED crowding has also become a critical issue because of increasing demands from the ageing population. To address this issue, the operational performance of hospitals' emergency care is closely monitored by the Ministry of Health (MOH), which publishes on its website, the daily median patient waiting time at each public hospital. According to the posted data, on 29 May 2015, median patient waiting times for admission to wards in the EDs of six public hospitals ranged from 1.0 hour to 8.1 hours. The latter median waiting time could imply serious crowding in the ED and an unpleasant experience for a large proportion of patients.

The common causes of ED crowding include an increasing demand for emergency care, insufficient capacity, operational inefficiencies, etc. The main theme of this article is effective patient flow management, which is expected to solve excessive patient delays without needing direct capacity expansion. To evaluate the timeliness and efficiency of emergency care, the National Quality Forum has endorsed three quality metrics: (i) length of stay (LOS), (ii) door-to-provider time (i.e., the time a patient spends in the ED before being seen by a healthcare provider), and (iii) leaving without being seen. Since the percentage of patients leaving without being seen is closely related to door-to-provider times, we regard the two time metrics as major performance concerns. In general, door-to-provider times should be kept below certain safety limits according to each patient's clinical urgency; LOS should also be kept below certain targets in order to reduce crowding.

The ED is a complex service network with prioritized customers (patients) and time-sensitive service requirements. In addition, the network is characterized by frequent customer returning routes; i.e., after initial consultations or treatment by a physician, most patients need to undergo medical tests and return to the same physician at least once before eventually being discharged or hospitalized. When there are multiple patients waiting to be seen in the ED, their respective physicians and the sequence of their consultations and treatments must be carefully scheduled in order to meet the stringent time requirements.

We use a quantitative research approach to schedule patients in the ED. The central element is patient flow data from the ED. We would like to maximize the percentage of patients whose door-to-provider times and LOS's are within the prescribed limits. Whenever a physician becomes available, the dynamic scheduling algorithm determines the next patient who should be seen, using data from existing patients. As recommendations need to be made in real time, the algorithm must be highly efficient in terms of computational time. This requirement turns out to be the greatest challenge because of the complex structure of the ED as well as time and routing constraints. This scheduling problem cannot be solved either by queueing theory, which relies heavily on unrealistic probabilistic assumptions, or by conventional stochastic optimization approaches, which are difficult to solve efficiently. To tackle this issue, we take a worst-case approach, solving the patient scheduling problem through robust optimization. This approach greatly improves the computational efficiency of the scheduling algorithm and is able to yield a near-optimal scheduling policy. Numerical experiments show that our approach can achieve much better performance than the Sample Average Approximation (SAA) method as well as existing patient scheduling algorithms.

Aside from its practical value, the quantitative model is methodologically novel. Although our approach is based on a robust formulation, it stems from a completely new perspective on solving the optimization problem under ambiguity. As opposed to conventional classic formulations where uncertainty sets are specified as fixed constraints, our approach assumes the uncertainty set to be adjustable, and searches for the scheduling policy that has the most probable uncertainty within the given constraints. This approach is highly computationally efficient and can be extended to more general scheduling problems arising from applications in healthcare, transportation, etc.