

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Mining the patient flow through an emergency department to deal with overcrowding

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1662826> since 2018-04-05T00:26:46Z

Publisher:

Springer New York LLC

Published version:

DOI:10.1007/978-3-319-66146-9_5

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)



UNIVERSITÀ DEGLI STUDI DI TORINO

This is an author version of the contribution published on:

Duma D., Aringhieri R.

Mining the patient flow through an emergency department to deal with overcrowding.

In 3rd International Conference on Health Care Systems Engineering, volume 210 of Springer Proceedings in Mathematics and Statistics, pages 49-59. Springer, Cham, 2017. Advance online publication 30 January 2018.

DOI: 10.1007/978-3-319-66146-9_5

When citing, please refer to the published version available at:

https://link.springer.com/chapter/10.1007%2F978-3-319-66146-9_5

Mining the patient flow through an Emergency Department to deal with overcrowding

Davide Duma and Roberto Aringhieri

Abstract The Emergency Department (ED) management presents a really high complexity due to the admissions of patients with a wide variety of diseases and different urgency, which require the execution of different activities involving human and medical resources. This have an impact on ED overcrowding that may affect the quality and access of health care. In this paper we apply Process Mining techniques to a real case study: from the ED database, discovery techniques identify the possible paths of a patient on the basis of the information available at the triage. Our purpose is to obtain precise process models for replicating and predicting the patient paths.

1 Introduction

The Emergency Department (ED) is a sub-unit of the hospital that operates 24 hours per day, 365 days per year, providing immediate treatments to patients. Although a substantial number of patients is self-referred, but does not need emergency care, EDs must ensure treatments to all of them. The ED management presents a really high complexity due to the admissions of patients with a wide variety of diseases and different urgency, which require the execution of different activities involving human and medical resources. The uncertainty determined by the high heterogeneity of cases brings about different problems, such as the extension of waiting times and the inefficient use of the available resources. For all these reasons, the ED overcrowding represents an international phenomenon [8], which may affect the quality

Davide Duma
Dipartimento di Informatica, Università degli Studi di Torino, Corso Svizzera 185, I-10149 Torino, Italy, e-mail: davide.duma@unito.it

Roberto Aringhieri
Dipartimento di Informatica, Università degli Studi di Torino, Corso Svizzera 185, I-10149 Torino, Italy, e-mail: roberto.aringhieri@unito.it

and access of health care because of medical errors, delays in treatments, risks to patient safety and poor patient outcomes [3], but also high levels of stress, impairment of the staff morale and increased costs [1].

The ED overcrowding is manifested through an excessive number of patients in the ED, patients being treated in hallways, ambulance diversions, long patient waiting times and patients leaving without treatment [5]. Since the perception of the crowding level from the staff is subjective [9] and because of the need to adequately prevent the phenomenon, several indices for the real-time measurement of overcrowding have been introduced and studied. They are based on different indices about the current operating status: the amount of available resources, the number of patients in the ED involved in some activities or waiting for a resource, their waiting times, the patient outcome and the predicted arrivals. However, the analysis in [4] shown that none of most popular overcrowding measures is capable of providing an adequate forewarning. Simulation has been widely used to understand causes of the overcrowding and to analyse the impact of several interventions (patient flow analysis, the bottlenecks detection, resource allocation) to alleviate its effects.

After interviewing the ED-staff operating on a real case study, we observed how the patient paths can be different and intricate, since they depend on specific needs that are not easily identifiable at the time of their arrival, taking into account many variables that affect medical conditions, shifts of the ED staff, crowding of the ED but also availability of beds for hospitalisation within the other wards. Although the flowchart of the ED process can be easily designed interviewing the management staff, rules defined by physicians are usually subjective and not sufficient precise due to their complexity, the practical sense of physicians taking decisions and the flexibility required by urgent actions. Then the replication of the patient flow is difficult without making significant assumptions, which do not allow us to act in real time on the bottlenecks to prevent or relieve overcrowding.

Nowadays huge amounts of data are collected by EDs, recording diagnosis and treatments of patients. Process Mining can exploit such data and provide an accurate view on health care processes [7], ensuring their understanding in order to generate benefits associated with efficiency [10]. In literature there are several discovery techniques that use specialised data-mining algorithms to extract knowledge from datasets, creating a process model that takes into account dependency, order and frequency of events, but also decision criteria and durations. However, the ED process we would to mine has the characteristics of a *Spaghetti process*, that is an unstructured process in which the huge variety of sequences of events affects the trade-off between simplicity and precision discovering the process.

In this paper we deal with the discovery of the ED process in a real case study. We use several Process Mining techniques to identify the possible paths of each patient on the basis of the only information known at the access of the patient. Our purpose is to obtain precise process models for replicating and predicting patient paths. An accurate replication of the patients paths allows us to simulate the process, generating a resource demand that depends on decisions taken in the past activities. Then the prediction of the future patient activities can be used to implement online approaches for the resource allocation to act on bottlenecks.

The paper is structured as follows. In Section 2 we describe a case study of a medium-size Italian ED, providing an analysis of the patient population. The process mining approach is presented in Section 3, discussing some preliminary results. Section 4 closes the paper.

2 A case study

We present a real case study concerning the ED sited at *Ospedale Sant'Antonio Abate di Cantù*, which is a medium size hospital in the region of Lombardy, Italy. The ED serves about 30000 patients per year, the urgency of which is classified by a code from 1 (most urgent) to 5 (less urgent) assigned by the triage-nurse at the time of their arrival, in according with Table 1. After the triage, the patient is visited in one of the visit rooms by a physician, which can prescribe X-ray examinations, several laboratory tests or therapies and a Short-Stay Observation (SSO). Certain patients are visited in other special rooms, such as the shock-room that is properly equipped for severely urgent interventions, and the Minor Codes Ambulatory (MCA), provided by the ED from Monday to Friday in the time slot 8:00–16:00 for adult patients with low urgency codes and good ambulation ability. In addition to all these ED tasks, there are activities that can be performed outside the ED, that is several specialist visits and the paediatric visit that is provided for non-urgent young patients instead of the medical visit. After examinations, treatments and specialist visits, the patient is revalued again by a physician of the ED, which establish how to continue the treatments, the need of hospitalisation or the discharge. In case of hospitalisation, patients are observed in the SSO units until a bed is available within the assigned hospital ward.

Table 1: Urgency codes: description and frequency over 2013–2015

number	color	description	frequency
1	red	immediate danger of death	1.5%
2	yellow	need of a timely medical visit	15.8%
3	green	need of treatments or investigations	61.6%
4	blue	symptoms that could be treated as primary care	13.8%
5	white		7.3%

The resources available within the ED are: 4 beds for the medical visits placed in 3 different visit rooms, in addition to one bed within the shock-room and another one in the MCA, one X-ray machine, 5 SSO units, 10 stretchers and 10 wheelchairs to transport patients with walking difficulties. The medical staff is composed of 4–6 nurses and 1–3 physician(s), depending on the time of day and the day of week, in addition to the X-ray technician.

Patient population

Thanks to the collaboration with the ED, we have available data concerning all accesses made in the years 2013–2015. Such data contains sex (male 52.7% or female 47.3%) and age of the patient, type of access (autonomously 79.9% or with a rescue vehicle 20.1%), the urgency code (1–5), the main symptom (undefined 35.2%, trauma 30.7%, abdominal pain 6.9%, flue 4.5%, chest pain 3.8%, dyspnea 3.4%, and other 25 options), timestamps and resources used during the activities, and type of discharge (ordinary 82.0%, hospitalisation 8.0%, abandonment 6.9%, transfer to another facility 2.6%, death 0.3% or hospitalisation refusal 0.2%).

The patient population is quite uniformly distributed across the different ages, with slight peaks for the age groups 5–9 and 35–54. To motivate this fact we compared the access frequencies of the five-year age classes with the demographic distribution. As shown in Figure 1, the almost uniform distribution of accesses among the age classes is due to the balance between the lower percentage of children and older people in the territorial area and the higher percentage of adults, which have a lower number of accesses per person. For the comparison, we used ISTAT data about 2014 in the province of Como, in which Cantù is located, observing that Lombardy Region and Italian territory have very similar distributions.




Fig. 1 Comparison between territorial and patient age distributions

Retrospective analysis

The ED of Cantù performed a retrospective analysis using the National Emergency Department Overcrowding Scale (NEDOCS) [12] in the aftermath of several management changes, such as the introduction of the MCA or a new staff rostering. In addition to inadequacy of this measure, proved in [4], the analysis performed by the ED of Cantù has been affected by the lack of several information that has been dealt with approximations. For all these reasons, we omit the NEDOCS results, focusing on a brief retrospective analysis that describes the variability of demand over time.

The accesses have different fluctuations over the day, among the days of the week and among the seasons, but also among the urgency classes. The higher arrival rate

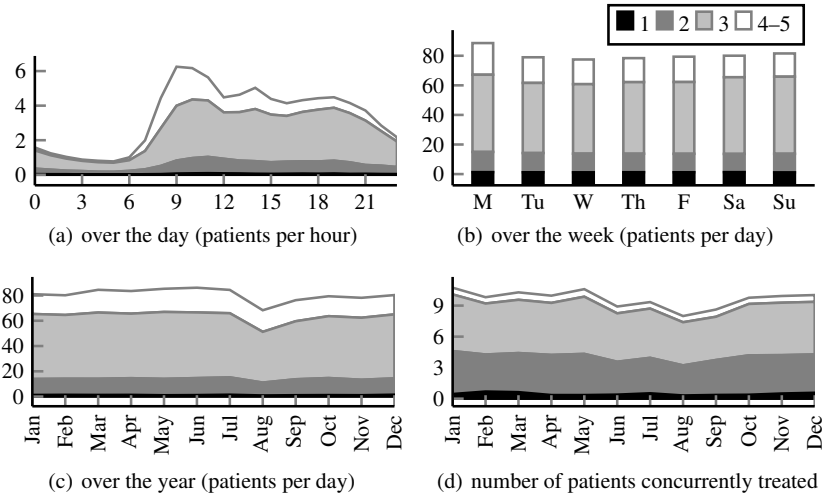


Fig. 2 Patient accesses divided for urgency code

fluctuations occur during the business hours of the day, as shown Figure 2(a), especially for the minor codes, which usually go to the ED instead of relying on primary care. For the same reason, a higher number of non-urgent arrivals has been registered on Monday, as shown in 2(b). Conversely, the urgency class 1 has the highest coefficient of variation among the different months of the week, because of medical and epidemiological reasons that causes more arrivals in winter. Nevertheless, from Figure 2(c) a uniform workload over the year (except for August) could be deduced, in fact, the workload do not depend directly of the number of accesses. Then, we report in Figure 2(d) the average number of patients concurrently treated (including all the activities between the first visit and the discharge), that is a more consistent indicator with respect to the ED staff perception. The statistics in Table 2 justify this fact, indeed more urgent patients have a longer average EDLOS (Emergency Department Length-Of-Stay). Such a difference is due to the higher frequency of SSO for patients with urgency codes 1 and 2, caused by a higher percentage of hospitalisations. The average waiting times confirm us that the priority among urgency codes is respected. Finally, lower urgency codes also have an higher rate of patients Leaving Without Been Seen (LWBS).

Table 2: Waiting times, LWBS and statistics on the treatment of patients

urgency code	average wait time	percentage of LWBS	average EDLOS	percentage of SSO	average SSO duration	percentage of hospitalisation
1	15 min	0.2%	9 hours	28.4%	19 hours	60.3%
2	34 min	0.5%	7 hours	19.5%	20 hours	28.9%
3	65 min	3.3%	2.5 hours	5.4%	19 hours	7.8%
4-5	68 min	11.1%	1 hour	0.5%	17 hours	1.1%

3 Process Modelling

In order to use discovery mining techniques, we need to preprocess the ED database in order to create an event log, that consists of a set of traces (i.e. ordered sequences of events of a single case), their multiplicity and other information about the single events, such as timestamps and/or durations, resources, case attributes and event attributes. In our case the events correspond to the activities concerning the patient treatments recorded for each access within the ED of Cantù's database, while each trace identifies a patient path. Then, our objective is to discover a simple process model that allows us: (i) to simulate patient paths, resource allocation and workload accurately, and (ii) to predict the next activities and the required resources of patients on the basis of their characteristics and their partial paths. All the process mining techniques cited in this Section have been used as plug-ins of *ProM 6.6*.

3.1 Preprocessing

The event log has been generated taking into account the accesses of the year 2015, removing all the accesses interrupted by death or a voluntary abandonment, because in these cases we do not know what would be the continuation of the treatments. However, such data can be used to study the patient LWBS phenomenon and their impact on performance. Each case of the event log consists in an access and events consists in activities, which has been classified into 15 event classes: *T* (triage), *V* (medical visit), *K* (shock-room treatment), *A* (MCA visit), *P* (paediatric visit), *S* (specialist visit), *R* (reevaluation visit), *X* (X-ray examinations), *E* (laboratory examinations), *Y* (therapy), *O* (SSO), *H* (hospitalisation), *Z* (hospitalisation refusal), *F* (transfer to facility) and *D* (discharge). Consecutive events of the same event class were merged because of the irrelevance from a control-flow perspective and because this allows us to simplify the process models.

The event log is composed by 141 202 events regarding by 27 039 cases, which generated 3 986 different traces. On average traces are formed by 5 events belonging to distinct classes, but the number of events per trace ranges between 3 and 25. The high number of traces with a low frequency is caused by three different factors, that is medical reasons (i.e. patients need very different treatments), incorrect recordings (i.e. noise) and incomplete data. For instance: some activities have been recorded later and not always immediately at the end of their execution because of the need to deal with an urgency; for technical reasons, reevaluation visits are not recorded every time; some traces contain both medical visit and paediatric visit but they refer to the same activity.

3.2 Process discovery

The huge number of traces suggests the use of discovery techniques that deal with low frequent behaviour and noise. We used two different process miners, the *Heuristic Miner* (HM) [11] and the *Inductive Miner – infrequent* (IMi) [6], both based on the control-flow perspective. The HM takes into account the order and the causal dependencies among the events within a trace, generating a model that uses the Heuristic Net (HN) notation, which is flexible because it can be easily converted in other notations, for instance a Petri Net (PN). The IMi is an extension of the *Inductive Miner* (IM), that is a divide-and-conquer approach based on dividing the events into disjoint sets taking into account their consecutiveness within traces, then the event log is splitted into sub-logs using these sets. The IMi uses the same approach but filters a fixed percentage of traces representing infrequent behaviour to create a PN. Both the techniques require low computational time, that is an important requirement because of the dimension of our event log. We discovered from the event log two different process models using the HM and the IMi, that are shown in Figures 3 and 4, called \mathcal{H} and \mathcal{I} .

The model \mathcal{H} has been generated varying the parameters *dependency* and *relative-to-best* of the HM in such a way to reach the best fitness, that is an index of the capacity to reproduce the behaviour recorded in the event log, equal to 62%. The obtained model \mathcal{H} , as well all the other generated varying the parameters, is a so-called *Spaghetti process* that is not sufficient simple to understand the whole process.

The model \mathcal{I} has been obtained varying the noise parameter of the IMi in order to have a good precision avoiding or limiting infrequent behaviour. However, we observed very slight deviations among the models ranging the noise percentage, that has been fixed to 20%. Contrariwise to \mathcal{H} , this model is very simple but not precise: the parallelisms among activities (represented by the grey boxes in the Figure 4) allowed by \mathcal{I} implies additional behaviour that is not present in the event log.




Fig. 3 Process model mined with the HM: model \mathcal{H} (heuristic net).




Fig. 4 Process model mined with the IMi: model \mathcal{S} (Petri net).

3.3 Patient clustering

In order to have a better balance of the four most important quality criteria of the process modelling (i.e. fitness, precision, generalisation and simplicity), we use the *Decision-Tree Miner* (DTM) described in [2], which analyses data flow to find rules explaining why individual cases take a particular path. Although the approach that is most straightforward and common in literature is the *Trace Clustering* (TC) [10], we opt for the DTM because of its computational efficiency and the ability to determine the most influential variables by itself. We execute the DTM using the event log, which contains the patient attributes, and the PN of the model \mathcal{S} that is easier to handle of \mathcal{H} . We selected only the variables known at the access of the patient, because we would classify the patients at their arrival and to predict their paths on the basis of his/her cluster.

Several guards have been determined on a subset of the transitions of \mathcal{S} . Such guards are predicates that indicates a criterion to fire a transition, using the F_1 -score as accuracy index. We denoted with white block transitions the performing of a certain activity, while black block transitions fire when activities represented by their parallel transitions are not executed.


The guards detected by the DTM are reported in Table 3. We remark that such guards are more objective and accurate than the indications given by the medical staff, for instance we are now able to determine exact boundary ages to classify the patient behaviour. Furthermore, the guards involve the most relevant variables from a control-flow perspective.

We clustered the patients in according with the characteristics determined by the guards. We started from the guard G_1 that have the highest F_1 -score and we split the set of all the patients into 2 clusters on the basis of the guard satisfaction. We repeated this operation on the clusters for all the other guards in decreasing order of their F_1 -score, avoiding the splits when one of the two sub-cluster had less than 270 cases, that is 1% of all patients, in order to ensure significant samples.

Table 3: Guards detected by the DTM on \mathcal{I} (variables: age a , sex x , main symptom s , arrival mode m and urgency code c)

id	transition	predicate	F_1 -score
G_1	Paediatric visit	$a \leq 19$	0.954
G_2	X-rays	$s = \text{trauma} \vee (a \geq 66 \wedge m = 118\text{-ambulance})$	0.703
G_3	Examinations	$s = \text{abdominal-pain} \vee (c \leq 2 \wedge a \geq 61)$	0.655
G_4	Therapy	$(c \leq 2 \wedge a \geq 19) \vee (c = 3 \wedge a \in [19, 70] \wedge s \neq \text{trauma} \wedge x = F)$	0.624
G_5	Medical visit	$a \geq 19$	0.562

The patient clustering determined 10 clusters that we used to create the same number of event logs. The cases of the new event logs ranges between 355 and 6493, that is between the 1.3% and 24% of the initial event log. The number of different traces within the same event log is significantly decreased, obtaining on average 585 and maximum 1128 different traces per cluster. All the corresponding process models mined through the HM are simpler than that in Figure 3 because they have less transactions between events, but they are also more precise than the model \mathcal{I} in Figure 4 because of the absence of parallelisms. We report in Figure 5 the model \mathcal{C} obtained using the event log of a single cluster and the HM. We observe that the model \mathcal{C} is considerably simpler than \mathcal{H} but gives information about the behaviour that is not deductible by \mathcal{H} , such as sequential order of activities or frequencies of the transactions between the same couple of events in different clusters.

Fig. 5 Model \mathcal{C} (heuristic net) for the cluster that satisfies $G_1 \wedge G_2$

The patient clustering may be further continued converting the new process models in PNs and run the DTM to discover new guards. Additional event attributes can be taken into account in order to discover different decisions that depends on several time variables (e.g. MCA visits are never performed over the weekend) or the system state (e.g. degree of crowding). Furthermore, the HM compute frequencies of transition between two events, that could be very useful to estimate the probability of a patient to follow a certain path and consequently to compute the expected waiting and execution times.

4 Conclusions

In this paper we apply Process Mining techniques to a real case study in order to obtain precise process models for replicating and predicting the patient paths. This is an important starting point on the perspective of implementing online algorithms for the patient flow optimisation with the aim of reducing overcrowding. Finally, the computational inefficiency of some procedures, such as the TC or the ILP-Based Process Discovery, suggests ad-hoc implementations of the process modelling techniques.

Acknowledgements The authors wish to thank Alessandra Farina, Elena Scola and Filippo Marconcini of the ED at *Ospedale Sant'Antonio Abate di Cantù* for the fruitful collaboration and for providing us the data set and allowing their use in this paper.

References

1. Anantharaman, V., Seth, P.: Emergency department overcrowding. In: *Emergency Department Leadership and Management: Best Principles and Practice*, pp. 257–269. Cambridge University Press (2014)
2. De Leoni, M., Van Der Aalst, W.: Data-aware process mining: Discovering decisions in processes using alignments. *Proceedings of the ACM Symposium on Applied Computing* pp. 1454–1461 (2013)
3. George, F., Evridiki, K.: The effect of emergency department crowding on patient outcomes. *Health Science Journal* **9**(1), 1–6 (2015)
4. Hoot, N., Zhou, C., Jones, I., Aronsky, D.: Measuring and forecasting emergency department crowding in real time. *Annals of Emergency Medicine* **49**(6), 747–755 (2007)
5. Hwang, U., Concato, J.: Care in the emergency department: How crowded is overcrowded? *Academic Emergency Medicine* **11**(10), 1097–1101 (2004)
6. Leemans, S., Fahland, D., van der Aalst, W.: Discovering block-structured process models from event logs containing infrequent behaviour. *Lecture Notes in Business Information Processing* **171 LNBIP**, 66–78 (2014)
7. Mans, R., Van Der Aalst, W., Vanwersch, R., Moleman, A.: Process mining in healthcare: Data challenges when answering frequently posed questions. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **7738 LNAI**, 140–153 (2013)
8. Paul, S., Reddy, M., Deflitch, C.: A systematic review of simulation studies investigating emergency department overcrowding. *Simulation* **86**(8-9), 559–571 (2010)
9. Reeder, T., Bursleson, D., Garrison, H.: The overcrowded emergency department: A comparison of staff perceptions. *Academic Emergency Medicine* **10**(10), 1059–1064 (2003)
10. Rojas, E., Munoz-Gama, J., Seplveda, M., Capurro, D.: Process mining in healthcare: A literature review. *Journal of Biomedical Informatics* **61**, 224–236 (2016)
11. Weijters, A., Ribeiro, J.: Flexible heuristics miner (FHM). *IEEE SSCI 2011: Symposium Series on Computational Intelligence - CIDM 2011: 2011 IEEE Symposium on Computational Intelligence and Data Mining* pp. 310–317 (2011)
12. Weiss, S., Derlet, R., Arndahl, J., Ernst, A., Richards, J., Fernandez-Frankelton, M., Schwab, R., Stair, T., Vicellio, P., Levy, D., Brautigam, M., Johnson, A., Nick, T.: Estimating the degree of emergency department overcrowding in academic medical centers: Results of the national ed overcrowding study (NEDOCS). *Academic Emergency Medicine* **11**(1), 38–50 (2004)