Andrews, Robert, Goel, Kanika, Corry, Paul, Burdett, Robert, Wynn, Moe Thandar, & Callow, Donna
(2022)
Process Data Analytics for Hospital Case-mix Planning.
Journal of Biomedical Informatics, *129*, Article number: 104056.

https://doi.org/10.1016/j.jbi.2022.104056

Process Data Analytics for Hospital Case-mix Planning

Robert Andrews[a,*], Kanika Goel[a], Paul Corry[b], Robert Burdett[b], Moe Thandar Wynn[a], Donna Callow[c]

[a] *School of Information Systems, Queensland University of Technology, Brisbane, Australia*
[b] *School of Mathematics, Queensland University of Technology, Brisbane, Australia*
[c] *Queensland Children's Hospital, Brisbane, Australia*

Abstract

The composition and volume of patients treated in a hospital, i.e., the patient case-mix, directly impacts resource utilisation. Despite advances in technology, existing case-mix planning approaches are mostly manual. In this paper, we report on a solution that was developed in collaboration with the Queensland Children's Hospital for supporting its case-mix planning using process mining. We investigated (1) How can process mining capabilities be used to inform hospital case-mix planning?, and (2) How can process data be used to assess hospital capacity assessment and inform hospital case-mix planning? The major contributions of this paper include (i) an automated workflow to support both process mining analysis, and capacity assessment, (ii) a process mining analysis designed to detect process performance and variations, and (iii) a novel capacity assessment model based on limiting-resource saturation.

Keywords: Hospital case-mix planning, Process Mining, Capacity assessment

*Please address correspondence to R.Andrews (r.andrews@qut.edu.au)
Email address: r.andrews@qut.edu.au (Robert Andrews)

1. Introduction

Hospitals are critical elements of health care systems. To provide the best care, hospitals must have sufficient resources and be managed and operated well [1]. Their productivity and utilisation are affected by many things, but it would be fair to say that the number of hospital resources employed (i.e., staff, treatment spaces, wards, and theatres), their time availability, and the shrewd allocation of patients' surgical and medical activities to those resources is highly influential [2, 3].

At present the corona virus pandemic has stretched the capacity of many hospitals around the world. The lack of capacity in hospitals, however, has been an issue for some time, due to population growth, aging populations and a rise in obesity related conditions [4]. To meet future demands better and to restrict further increases in the cost of health care, hospitals need to be expanded and reconfigured wisely where possible or else their existing capacity should be utilised more effectively via improved scheduling and resource allocation [4, 5, 6].

Capacity planning and capacity assessment [7, 5] activities are equally vital in this endeavour. The purpose of capacity planning is to determine the resources that are required to meet notional levels of demand. Capacity assessment in contrast, considers what can be achieved with a given set of resources. It is an enabler for capacity planning, which can not be performed without a means to evaluate alternative resource selections.

Despite the advancements in information technology (IT) over the last 30 years, planning and scheduling in most hospitals, to the best of our knowledge, is predominantly performed at a departmental level, in meetings, using experience and mental arithmetic, and Excel style spreadsheets. The existing planning approaches can often be myopic, and the consequences of decisions are seldom known from a long-term hospital wide perspective. Anecdotal evidence and site visits to local hospitals revealed few scientific methods and "on the shelf" IT technologies for hospital staff to use. A lack of appropriate visualisation tools exacerbates the situation.

Analysing hospital capacity and productivity is an important topic. Hospital planners and executives regularly contend with challenging capacity-related decisions [5, 7]. Decisions relating to prioritisation, allocation and sharing of resources within a single hospital or across many hospitals within a localised region, have a profound impact on productivity, efficiency, and patient outcomes.

Choosing a patient case-mix that can be treated efficiently is key to success [8]. The composition and volume of patients to be treated in a hospital, namely the patient case-mix, has a large impact on resource utilisation. Case-mix planning (CMP) is the name given to the problem of determining the ideal composition and volume of patients to be treated in a hospital. Case-mix planning is a task specifically relevant to hospitals and other health care facilities [9, 5]. This is the problem of identifying a patient cohort (a.k.a., case-mix) with a specific set of features deemed desirable or ideal [9]. Some case-mixes are favourable for some patient types and unfavourable for others. Second, the term "ideal" is subjective and can mean different things in a practical setting. A case-mix may be sought that is most equitable, for instance in the allocation and usage of hospital resources (e.g., [10]). A case-mix may also be sought that is most economical or financially viable to treat (e.g., [11]). From a utilisation and output-oriented perspective, a maximal cohort may also be sought. That cohort results in the greatest number of patients treated over time. A maximal cohort saturates the resources of the hospital and is a measure of the hospital's capacity. Identifying a case-mix that meets or exceeds specified demands or targets is also of significant interest

In recent years, there has been much research on hospital capacity planning and capacity assessment. Patient case-mix and care pathways are a key ingredient in the approaches developed [7, 5, 9, 12]. However, upon retrospection, those two things are difficult to obtain from current hospitals. Identifying the current case-mix as it were, is difficult. Understanding a hospital's case-mix is challenging. CMP is often made more complex by either a lack of precise information, a high volume of unrefined empirical data, and stochastic parameters.

Process mining is a specialised form of data-driven analysis that uses algorithms and data (in the form of event logs) to construct models that aim to provide insights into the behaviour of organisational processes [13]. Process mining takes a retrospective, bottom up approach to understanding process behaviour. The execution of individual process steps (events) are logged. Within the log, events are grouped into cases where each case represents a single, end-to-end, process execution instance. An event will minimally include attributes that identify the case (process execution instance), the action that was performed, and an attribute (usually a timestamp) that allows actions to be ordered within a case. Optionally, an event log record may contain other attributes, such as the resource that performed the action, the organisational unit associated with the action or resource, that provide additional information about the event.

Evidence of the application of process mining across multiple industry sectors can be found in studies such as [14] and [15] which review respectively 144 and 152 case studies. Both these studies found that process mining was most frequently applied in sectors such as public administration, finance, insurance, healthcare, manufacturing, and education. Rojas et al. [16] provides a detailed review of 74 published accounts of the application of process mining in the healthcare domain. Although there have been many applications of process mining in healthcare, to the best of our knowledge, existing process mining studies have not explored how process driven insights could be used to inform hospital case-mix planning.

Research Aims and Methodology. In this work we explore how process mining techniques can be used to inform the problem of hospital case-mix (capacity) planning, and report on a solution that was developed in collaboration with the Queensland Children's Hospital (QCH), a major Australian hospital, for supporting its case-mix planning. We investigated two interrelated analysis questions which emerged, after consultation with QCH planners, as being of interest to both the research team and the hospital planners.

1. How can process mining capabilities be used to detect process behaviour and performance variations to inform hospital case-mix planning?

2. How can process data be used to assess hospital capacity assessment and inform hospital case-mix planning?

The major contributions of this paper include (i) an automated workflow to support both process mining analysis, and capacity assessment to inform case-mix planning, (ii) a process mining analysis designed to detect process performance and variations, and (iii) a novel capacity assessment model based on resource saturation.

The rest of the paper is organised as follows. Section 2 discusses related work while Section 3 describes the case scenario of the hospital and the dataset. Section 4 details the overall approach proposed in the paper. Section 5 presents our key findings while Section 6 concludes this paper.

2. Related work

In this section, we consider related works in the areas of (i) process mining in healthcare, and (ii) capacity planning and assessment.

2.1. Process Mining in Healthcare

Healthcare processes are characterised as being complex with significant variations over time [17]; the variation being due to the patient-centric nature of treatment pathways and multiple sequences in which activities in the treatment pathways can be executed by resources (physicians, nurses, etc). Process mining, with its various techniques for discovering process models and analysing their performance, affords the exploitation of the wealth of information stored in hospital information systems (HIS) to properly understand and improve the quality and efficiency of delivered healthcare services.

The first published accounts of process mining being applied in the healthcare domain [18, 19] were exploratory studies of data pertaining to stroke patients [18], and to gynaecological oncology patients [19] which aimed to demonstrate the potential utility of process mining in the healthcare domain. These studies discuss some important considerations in applying process mining to healthcare processes including (i) event log preparation - identifying, extracting, abstracting, and pre-processing event data from hospital information systems (that are generally not 'process-aware') such that interpretable results can be obtained, (ii) 'spaghetti' process models - discovered process models with a high number of pathways due to the multiple ways in which a typical, patient-centric healthcare process may execute.

In an important contribution to the field, Mans et al. [20] addresses three issues affecting process mining in healthcare namely (i) data correlation from multiple systems, (ii) typical questions of interest for healthcare stakeholders, and (iii) identification of data quality issues.

In [21], the authors review 37 studies in which process mining was applied to clinical pathways. The studies are classified according to whether they attempt to (i) discover actual execution pathways of different clinical pathways (process discovery), (ii) analyse variants of execution pathways, or (iii) evaluate and improve clinical pathways. The authors conclude, that at the time of writing, challenges remain including improving process mining algorithms so that they are (i) efficient enough to deal with the unstructured processes (clinical pathways) and (ii) able to discover models from which a good explanation of the variants can be obtained.

More recently, process mining has been used to discover processes, analyse performance, and check conformance of medical treatment processes and healthcare organisation processes [16]. For instance, process mining has been widely applied to improve cancer care processes [22]. Another application of process mining is to compare processes between healthcare organisations. Partington et al. [23] describe approaches to performing comparative analysis using process mining for cohorts of patients suffering chest pains in four Australian hospitals. Process mining has also been applied in pre-hospital setting. Badakhshan and Alibabaei [24] apply discovery, conformance checking and performance analysis techniques in a case study involving ambulance services in Iran. Process mining has also been conducted to obtain additional valuable insights related to processes. For example, social network can be mined to understand interaction among health professionals [25].

Process mining thus enables data-driven process improvement in healthcare. However, there still remains limited uptake of process mining in healthcare organisations [26]. In particular, the potential of process mining for hospital capacity planning remained unexplored.

2.2. Capacity Planning

Quantitative research on hospital processes has increased significantly over the last decade. Hospital capacity assessment, however, has received comparatively less attention in the literature. There are a few well documented "stand alone" capacity assessment approaches like the deterministic approach of [7, 5, 9, 12]. Stochastic approaches like those of McRae et al. [27] now also exist.

A brief summary of the contributions of the aforementioned approaches is now provided. Burdett et al. [5] developed a mixed integer linear programming (MIP) model that determines the maximum number of activities that can be performed within a given duration of time, subject to some technical constraints. The definition of activity is unrestricted and open to interpretation. In that paper, it was defined as a surgical or medical activity within a patient care pathway. Burdett et al. [7] later continued their research and introduced multiple objectives. The multi-objective hospital capacity model (MOHCM) identifies non-dominated capacity solutions and provides a sensitivity analysis of patient case-mix and the effect on hospital capacity. In their numerical testing, 21 objectives were considered, one for each surgical specialty.

Freeman et al. [12] considered case-mix planning and developed a multi-phase approach to generate a pool of solutions. They used simulation to evaluate each solution. They also simulated the master surgical schedule (MSS). They reported that existing CMP approaches provide a single solution and exclude uncertain patient arrivals and operation times and the arrival of patients requiring urgent care.

McRae et al. [9] developed a non-linear mixed integer programming model for CMP. Their model incorporates economies of scale and permits an investigation of the effect of changes in the efficiency of resource use with increasing scale on the optimal case-mix. As their model is non-linear, piecewise linear functions were used, and an iterative approximation scheme (e-optimal solution methodology) was applied. They conclude that meaningful results depend upon the accuracy of input parameters. Demand for instance is difficult to obtain. They omitted uncertainty to keep their computations tractable. McRae et al. [27] presented a framework for evaluating stochastic aspects and different levels of aggregation on the performance of CMP in hospitals. Stochastic influences are categorised according to whether they relate to demand, resource consumption, and resource availability. Numerical testing of different options is performed. Capacity planning (a.k.a. resource capacity planning) has been considered far more frequently. Capacity planning however is a catchphrase for many types of resource planning, scheduling and forecasting activities. Hulshof et al. [28] have provided a comprehensive, structured overview of resource capacity planning and control in health care. Their review of the literature confirms prior observations that there are few contributions that incorporate complete hospital and health care system interactions.

Within the domain of capacity planning, a variety of models that determine how many beds are required for notional levels of demand have been developed [29]. Approaches to forecast hospital demand also exist. Jalalpour et al. [30] developed a MATLAB toolbox to forecast count data. Their toolbox uses the maximum likelihood method to estimate model parameters from data. A generalised auto-regressive moving average (GARMA) model was used because it can produce forecast models that outperform the traditional Gaussian models. An approach like this can be used to predict hospital demands and predict overall performance. It is however incapable of analysing the system if parametric or structural changes are made, as there is no longer data to analyse. Their approach does not consider concrete physical attributes of a hospital.

Network flow models have also been successfully developed in Akcali et al. [31]. Their generic approach produces bed capacity plans and incorporates facility and budgetary constraints over a finite planning horizon. They conclude however that alternative model formulations are needed when there are specialty-specific demand rates, length of stays and costs. Rechel et al. [32] report that bed numbers are still used for capacity planning in many countries. An empty bed, however, does not count as capacity if there is insufficient staff to care for a patient in that bed [33].

Chen et al. [34] consider patient flow scheduling and capacity planning within the context of a smart hospital and health care environment. In their viewpoint, a smart environment is one that is designed to facilitate people's experience that includes a set of devices and many intelligent supporting techniques. In response, a quantitative analysis and a dynamic scheduling policy are proposed. A formal algebraic modelling approach, an ordinary differential equation (ODE) based fluid flow analysis and simulation tools are implemented to facilitate the activities. Promising as their approach is, it has only been applied to a rheumatology department and not to an entire hospital.

Hospitals are located within specific geographical locations, and the effect of demographic factors on hospital management decisions is important. Li and Benton [35] have analysed the effect of hospital capacity resource management choices on cost control and quality using a structural equation modelling (SEM) methodology. Their goodness of fit statistical approach relates specified factors to observed/collected data. A data envelopment analysis (DEA) approach was developed by Valdmanis et al. [33] to determine state-wide hospital capacity. Each hospital in their study was assessed in terms of specialty capacity and general capacity. Their capacity metric is consistent with engineering practices and is the maximum rate of output per unit of time. Their approach was developed to facilitate emergency preparedness planning. Information on hospital capacity, patient characteristics of inpatient discharges, and financial performance were merged to perform their study.

Process mining has not, to date, been applied to capacity planning in hospitals. Further, to date, no end-to-end, automated, approach has been found that can pre-process and format data routinely collected process data extracted from hospital information systems, and perform analyses (such as process mining and capacity assessment) that are the precursors for capacity (case-mix) planning. In this paper, we explore how process data and process mining can be used to contribute to capacity assessment and capacity (case-mix) planning for a hospital. In particular, we aim to develop an integrated, automated workflow, using already available tools, to support these tasks.

3. Case Scenario

The Queensland Children's Hospital (QCH) provides general paediatric care within its catchment area of inner Brisbane (a major metropolitan city in Australia). QCH provides emergency, critical care, general and specialist paediatric services for patients from all over Queensland and northern New South Wales, which has a combined population in excess of 5 million people.

The dataset used for this study was provided by QCH and included all inpatient and outpatient episodes in the period July 1, 2019 to December 2020. The data included admission details (for inpatients), ward movements or location changes for any patient during a hospital encounter, imaging (scans and treatments), as well as details of all surgical procedures for each surgical admission.

The data was supplied in tabular form (4 x CSV files) with a unique

Table	**Description**	**Size**
Admissions	Details of an **inpatient** encounter for an individual patient.	57,472 records including 57,138 distinct encounters.
	NB An outpatient is a patient that consumes one or more hospital resources, such as a cubicle in the emergency department, or undergoes a day surgery, without being formally admitted to the hospital.	There were 54,831 outpatient encounters recorded in the study period.
Movements	Details of ward movements for a patient encounter (including ward/location name, time of allocation, and time of departure.	119,865 records.
Imaging	Details of request/orders as well as type of imaging/radiation treatment carried out for a patient during an encounter.	89,211 records
Surgery	Details of surgical procedures carried out for a patient during an encounter.	46,450 distinct procedures comprising 46,211 so-called surgeries.

Table 1: Summary of Queensland Children's Hospital data.

encounter identifier provided as a linking key. **NB** *Outpatients* are patients that consume some hospital resource, e.g., an emergency department cubicle, for some period of time without being properly admitted to the hospital.

Outpatients include emergency department presentations that are discharged from emergency, patients undergoing day surgery (no ward stay after the procedure), and patients undergoing some form of regular treatment such as radiotherapy. Table 1 provides more details of the QCH data.

Features of interest to the project stakeholders, and required by the hospital capacity assessment model include:

- a process model,
- descriptive statistics of ward and operating theatre usage over time,
- assignment of individual resources (wards, theatres, etc.) into functional resource groups,
- determination of patient groups around which capacity assessment and case-mix planning is done),
- determination of patient care pathways that capture the resources and utilisation of the resources by each patient group.

4. Research Approach

Figure 1 provides an overview of how process data was used for case-mix planning in this paper. The coloured boxes in the figure depict tasks that were conducted by the research team (and will be elaborated further in the later subsections) while the other three tasks (i.e., identification of analysis questions, data extraction, hospital case-mix planning) were led by hospital stakeholders.

The first step involved the identification of key analysis questions, following which an appropriate dataset was defined and extracted. We conducted an in-depth data quality assessment and pre-processed the data provided by

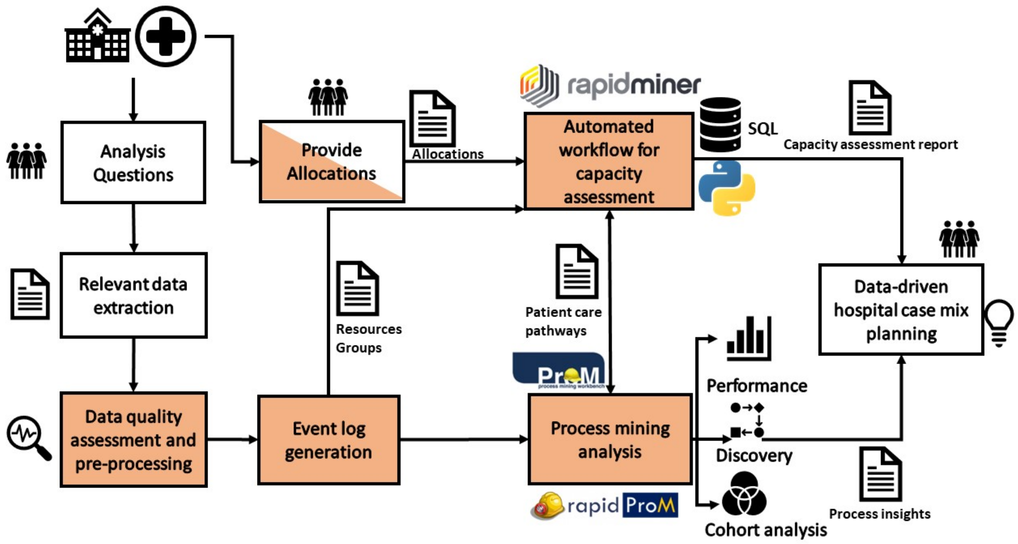

Figure 1: Overall approach demonstrating the use of process data for case-mix planning.

the hospital (described in 3) to extract an event log. The event log (process data) was used as input for process mining analysis, in particular, discovery, performance, and patient group analysis. The event log (particularly data related to patient groups and resources) was also the input to the automated workflow developed for hospital capacity assessment. Furthermore, allocations were provided by stakeholders and also analysed from dataset, which was another input for the automated workflow for capacity assessment. In addition, patient care pathways were discovered using process mining analysis, which was another input for the automated workflow for capacity assessment. The process insights and the capacity assessment report are then produced by this approach to assist case-mix planning at QCH. The following subsections describe the key steps of the proposed approach in detail.

4.1. Data Quality Assessment and Data-preprocessing

The impact of data quality on insights derived from model-based analysis techniques (e.g., process mining and capacity assessment) is well recognised. This is why considerable time was spent in assessing the quality of data and pre-processing it. The original dataset was extracted from hospital information systems and comprised tables regarding admission of a patient, movement of patients among wards and operation theatres, and surgical procedures conducted on the patients. Inspection of the dataset revealed multiple data quality issues such as missing timestamps, duplicate values, and incorrect event ordering due to manual editing of timestamps during case review. Examination of the organisational context revealed that many of the data quality issues resulted from a combination of human error and system configuration. For instance, when conducting post hoc review/edit of timestamped activities, rather than presenting the originally recorded date/time for the activity, the system defaulted to the current date/time. So, if the operator was not aware of this and changed only the time component, the edited date no longer reflected the date on which the activity was conducted. These issues were rectified in consultation with domain experts. Additionally, a lack of clarity regarding movement to wards and operation theatres was also observed. This information was also clarified through discussions with domain experts.

4.2. Event Log Generation

A unique encounter identifier was used in each of the Admissions, Movements, Imaging, and Surgery tables. This attribute was used as the case identifier.

The Admissions table included 2 attributes (admission date/time and discharge date/time) that could be cast as events in an event log. The Admissions table contained attributes such age, gender, DRG (Diagnostic Related Group), and (medical) specialty that could easily be cast as event log case attributes. We note that only patients that were actually admitted to hospital (i.e., inpatients) appeared in this table. Emergency patients not requiring admission for further treatment (i.e., discharged home from the emergency department) and outpatients (for instance, patients receiving regular chemotherapy treatment and not requiring an overnight stay) were not recorded in the Admissions table.

The Movements table included 2 attributes (allocation date/time and leaving date/time) that could be cast as events in an event log. The location against which the allocation was made could be cast as an event attribute and constituted the resource being consumed by the patient during the period between allocation to, and departure from, the location.

The Imaging table contained 3 attributes (order date/time, image start date/time, and image complete date/time) that could be cast as events. The name of the actual imaging or treatment procedure was included as an event attribute. We note that the imaging data did not include the location used for the imaging procedure. Hence, it was not possible to model the resource utilisation associated with imaging (or radiation oncology treatments).

The Surgery table contained 11 attributes that were converted to events. These events captured the milestones of a usual surgery and included the date/time the patient was (i) checked-in for surgery, (ii) pre-op begin and end, (iii) anaesthesia begin and end, (iv) surgery begin and end, and (v) post-operative recovery begin and end. Attributes such as the relevant surgical specialty, the theatre suite and actual operating theatre, as well as the actual surgical procedure carried out, were cast as event attributes. Here the surgical specialty was used to group patients, and the operating theatre was treated as a resource utilised by the patient.

For analysis, we chose the period 1-July-2019 to 31-December-2020 as representing a relatively stable and recent period. Ultimately, the event log consisted of:

- 109,501 cases (inpatients and outpatients)
- 901,431 events
- 19 activities
- 30 surgical specialties
- 70 locations (wards, operating theatres emergency, imaging)

Case numbers in each of the three 6 month segments were as follows:

- 2019-07-01 to 2019-12-31: 39,885 cases (avg duration: 11 days)
- 2020-01-01 to 2020-06-30: 31,930 cases (avg duration: 9 days)
- 2020-07-01 to 2020-12-31: 37,686 cases (avg duration: 3 days)

We note that the period 2020-01-01 to 2020-06-30 corresponded to the emergence of COVID-19 in the Australian community. During this period, there were disruptions to normal behaviour (including work from home, lockdowns, etc.) which may have had some influence on reduced case numbers.

4.3. Automated Workflow to Support Capacity Assessment

To generate insights related to the analysis questions, we used a combination of SQL queries to extract model inputs from the event log data and

RapidMiner [1] to develop an automated workflow which executed the SQL queries, and developed and visualised summary statistics relevant to resource utilisation. Using the RapidProM extension, the RapidMiner workflow also read the event log and visualised the log as a discovered process model using the Inductive visual Miner operator. The RapidMiner workflow executed the SQL queries that populated the capacity assessment model inputs (see Section 4.4). These were passed to the ExecutePython operator to execute the capacity assessment model (which had been implemented as a Python script). In this section, the automated workflow is explained (see Section 4.5).

For the first phase of the study, it was decided to focus on operating theatre utilisation, hence, cases in the event log were grouped according to the dominant (most frequent) surgical specialty.

A possibility for such characterisation was to simply use event log trace variants. However, for the initial capacity assessment modelling, it was sufficient to know only which resources were consumed, (and for how long). Hence, cases were grouped according to the case's dominant surgical specialty, and the set of hospital resources consumed by the case. Thus, we characterise a patient care pathways as vectors, with each vector having a care path identifier, an attribute for the relevant patient group, and attributes for each resource type and the average utilisation (in hours) of the resource type by patients belonging to the group. To determine hospital-level resource usage over time, the event log and patient-level resource usage

[1] https://rapidminer.com/

was combined to construct a dataset that gave an hour-by-hour occupancy of hospital locations (wards and operating theatres). For operating theatres, the relevant surgical specialty was also recorded. Lastly, to assist in 'typing' patient admission and pathways, a dataset was constructed that mapped each admission to the set of surgical specialties involved with the admission. The data transformation process is illustrated in Fig 2.

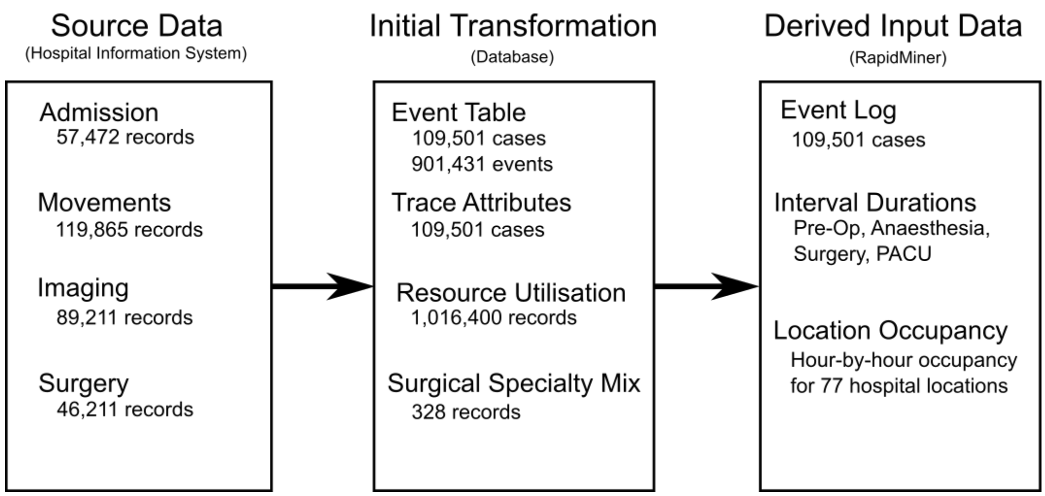

Figure 2: Data Transformation Summary

Figure 3 illustrates the RapidMiner process developed to address the process mining, capacity assessment modelling, and operating theatre usage summary statistics and visualisations. Inputs to the process include (i) event data (exported from the relational database, (ii) trace attributes (exported from the relational database), and (iii) resource usage (direct connection to the database).

The RapidMiner process is divided into four concurrent streams. The first stream filters the event data to extract events such that durations of key surgery activities can be extracted (Figure 3(a) - Filter Durations). These

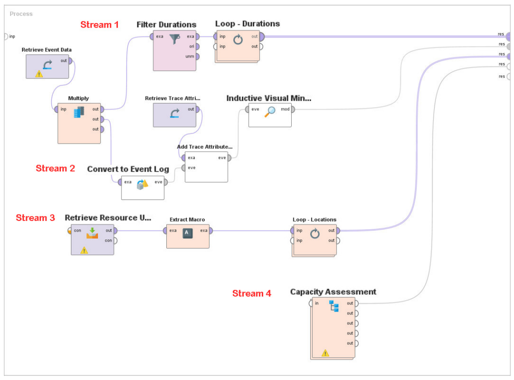

(a) High level RapidMiner workflow.

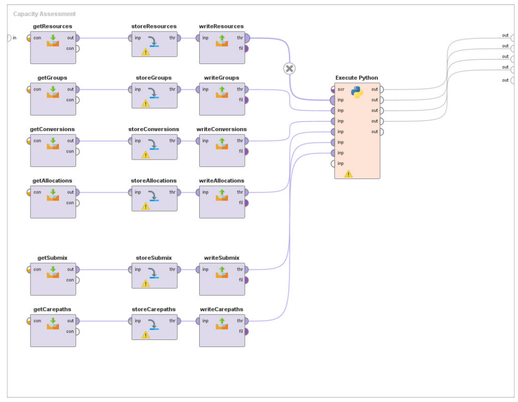

(b) Sub-process that implements the capacity assessment model.

Figure 3: RapidMiner Automated Workflow.

activities include PreOp, Anaesthesia, Surgery, Patient in OT, PostOp Care (PACU). The workflow then loops through each activity (Figure 3(a) - Loop - Durations) to allow visualisation of activity durations.

The second stream of the RapidMiner process converts the tabular event data to an event log suitable for process analysis (Figure 3(a) - Convert to Event Log). The Inductive Visual Miner plugin of the RapidProM extension of RapidMiner is used to discover a process model (Figure 3(a) - Inductive Visual Miner). Once the variations in distribution of resource usage across specialities is investigated, the developed approach enables the user to model and view process behaviour for different specialities. The discovered model can also be used to conformance checking purposes and identify specialities, which display the behaviour that is different from the expected behaviour.

The third stream of the RapidMiner workflow reads the hour-by-hour resource usage (wards and operating theatres) from the database. Here, we define resource usage as occupancy rate. That is, the number of patients in the ward or operating theatre at any given hour of any given day in the study period. Once the variations in distribution of resource usage across specialities is investigated, the developed approach enables the user to model and view process behaviour for different specialities.

The fourth stream is a sub-process (Figure 3(a) - Capacity Assessment expanded in (Figure 3(b)) that executes the SQL queries that extract the capacity assessment inputs from the data and executes the model Python code.

4.4. Capacity Assessment

Hospital capacity assessment is similar to bottleneck analysis in process mining. In process mining, timestamped activity information enables analyses involving the timing and frequency of events, including monitoring of resource utilisation, and discovery of bottlenecks (where a bottleneck is a resource having an arrival rate greater than the throughput rate).

In the hospital setting, different patient groups will have different limiting resources. For each patient group, a limiting resource is a resource that (i) has a fixed and finite capacity (i.e., patient occupancy), (ii) is used by all patients in the group. Examples of such limiting resources include, (i) the number of haemodialysis machines for renal dialysis patients, and (ii) the number of operating theatres for surgical patients requiring general anaesthetic. Each patient group will also, optionally, utilise other hospital resources. For example, some (not all) patients will be admitted to the hospital and will occupy a ward bed for some period of time. Some (not all) patients will require intensive care and occupy a bed in ICU for some time. This therefore means that the hospital's patients divided into mutually exclusive patient groups according to the set of hospital resources utilised.

To formalise the problem of hospital case-mix planing, we define R as the set of all **resource types**, G as the set of **patient groups** and P as the set of all **patient care pathways (PCPs)** where $P = \bigcup_{g \in G} P_g$, where P_g are mutually exclusive sets of PCPs within patient group $g \in G$. We denote $R_g \subseteq R$ as the set of resource types with dedicated capacity allocated to patient group $g \in G$ (with $r = 0$, $r \in R$ representing operating theatres).

Now we define T_r as total time capacity (per unit time) of resource type $r \in R$; $T_{r,g}$ as the total time capacity of resource type $r \in R_g$ allocated to patient group $g \in G$; $t_{r,g,p}$ as the total time required of resource type r for one patient from group g with PCP $p \in P_g$; $\mu^2_{g,p}$ as the case sub-mix, i.e. the proportion of group g with PCP $p \in P_g$. Lastly, we define $r_g \in R_g$ as the allocated resource type considered to be capacity limiting for group g, with $r_g = 0$ for surgical in-patients.

We calculate $n^2_{g,p}$ the number of patients from PCP $p \in P_g$ treated per unit time (see Eq. 1), and subsequently derive $\rho_{r,g}$ the utilisation of type r resource allocation for group g with $r \in R_g$ (see Eq. 2), and ρ_r the overall utilisation of type $r \in R$ resources (see Eq. 3).

4.4.1. Capacity Assessment Model

This model does not restrict patient throughput due to over-saturation of resources upstream or downstream of surgery. However, any capacity violations will show up with utilisation greater than 1. From a decision-maker's perspective this would be symptomatic of misalignment between allocated theatre capacity (e.g. operating theatres in the case of the Master Surgical

Schedule), and capacity of other parts of the hospital.

The number of surgical patients per unit time of PCP (g, p) that can pass through operating theatres ($r_g = 0$ and $T_{0,g} > 0$), or medical patients that can pass through allocated treatment space ($r_g > 0$ and $T_{0,g} = 0$):

$$n^2_{g,p} = \frac{\mu^2_{g,p} T_{r_g,g}}{\sum_{q \in P_g} \mu^2_{g,q} t_{r_g,g,q}} \qquad \forall g \in G, p \in P_g \tag{1}$$

The overall utilisation of resource type r is then:

$$\rho_r = \frac{\sum_{g \in G} \sum_{p \in P_g} t_{r,g,p} n^2_{g,p}}{T_r} \qquad \forall r \in R \tag{2}$$

And the utilisation of group g allocation for resource $r \in R_g$ is:

$$\rho_{r,g} = \frac{\sum_{p \in P_g} t_{r,g,p} n^2_{g,p}}{T_{r,g}} \qquad \forall g \in G, r \in R \tag{3}$$

4.4.2. Calculation of Parameters

Assume that ward beds correspond to resource type $r = 1$, then T_1 would be the total number of ward beds in the hospital, and $T_{1,g}$ would be the number of beds allocated to specialty g.

For operating theatre capacity, let τ_g be the number of theatre hours per week allocated to specialty g (in a weekly Master Surgical Schedule). Let $\bar{\rho}$ be the limiting capacity of operating theatres (e.g. $\bar{\rho} = 0.85$ representing a practical upper limit on utilisation of 85%). Then the total time capacity per unit time available (per week) to specialty g can be calculated as:

$$T_{0,g} = \bar{\rho}\tau_g \qquad \forall g \in G \tag{4}$$

Consider the following as an example of the way the model may be applied. If there are 2 specialties, with 2 PCP in specialty 1, and 1 PCP in specialty 2. Consider $R = \{0\}$ including operating theatres only. Assume the following parameter values with time units in hours:

$T_0 = (120, 80), t_0 = ((2, 2.5), (1.22)), \mu^2 = ((0.5834, 0.4166), (1))$

Let $\bar{t_1} = 0.5834 \times 2 + 0.4166 \times 2.5 = 2.2083$ and $\bar{t_2} = 1.22$

then, $n^2_{0,1,1} = \frac{0.5834 \times 120}{2.2083} = 31.7$ and $n^2_{0,1,2} = \frac{0.4166 \times 120}{2.2083} = 22.6$

and $n^2_{0,2,1} = \frac{80}{1.22} = 65.6$

Now for ward bed utilisation and assuming 20 beds shared between specialties we have $T_1 = 3360$, (total bed hours per week) and assume the following average post-op times $t_1 = ((12, 18), (7))$. We can calculate bed utilisation as follows:

$\rho_1 = \frac{12 \times 31.7 + 18 \times 22.6 + 7 \times 65.6}{3360} = 0.37$

Low utilisation suggests ward beds are over-supplied.

4.4.3. Capacity Assessment Model Input

The following describes the inputs and data used by the model.

Resources: these are the physical resources occupied by each patient during their stay in hospital. For instance, a bed in a ward, an operating theatre, or an ICU bed. There were 77 separate locations in the hospital to which patients were allocated. (The location is an event attribute for the 'Allocate Location' and 'Surgery Start' activities in the event log.) After consultation with hospital stakeholders, it was decided to exclude several of these locations on the basis that they were not limiting resources. That is, they did not impact on the hospital's ability to deal with inpatients. Such locations included so-called 'virtual' locations, i.e., used for tele-medicine, and 'waiting lounges'. Also excluded were some specialist locations catering for small numbers of patients, e.g., burns or renal patients. Finally, resources associated with emergency patients were excluded as, while in the emergency department, a patient is not considered an inpatient. The remaining 55 resources were then grouped according to 'type'. Resource types included in the model then were operating theatres (OT), ward beds (WBED), and ICU beds (ICU). Table 2 gives the number of resources in each resource type.

Resource Type	**Locations**	**Description**
ICU	3	intensive care wards
OT	25	dedicated operating theatre rooms
WBED	27	inpatient wards

Table 2: Resource types and number of locations by type

Groups: these are non-overlapping subsets of 109,501 cases included in the study. Each such subset is designated a 'group'. In this scenario, as the focus was on operating theatre capacity, groups were based on the surgical specialty under which each case (patient) is treated. If a patient used an operating theatre (underwent a surgical procedure), they were assigned to a group related to the surgical specialty performing the procedure(s). There were 30 such surgical specialties. A further group was created for patients who were admitted but who did not undergo a surgical procedure. The top 5 largest patient groups are shown in Table 3. Note that the limiting resource is resource that all patients in the group must use.

Patient Group	Num Patients	Limiting Resource
Non-Surgical	83,366	WBED
Paediatric Surgery SN	5,049	OT
Ear Nose and Throat SN	3,889	OT
Orthopaedic SN	3854	OT
Gastroenterology SN	1,493	OT

Table 3: Top 5 largest patient groups (SN = surgical)

Patient Care Pathways: are the set of resources and the quantity of each resource consumed. Each Group may have multiple patient care pathways (carepaths) depending on the locations (and hence resource types) used by patients in the group. The actual hospital locations were used to generate the carepaths, then the resources were aggregated by type. Carepaths were identified by querying the event log. Cases in each patient group were queried to determine the set of resources used and the time-based utilisation of each resource. These were aggregated into distinct sets, and the distinct locations were mapped to their respective resource type to form a carepath.

Table 4 shows the distribution of cases and patient care pathways across

Surgical Specialty	Cases	PCPs	Surgical Specialty	Cases	PCPs
Paediatric Surgery	5,279	808	Oncology	2,063	222
Orthopaedic	3,885	528	Gastroenterology	1,496	220
Ear Nose And Throat	3,897	335	Plastics	1,435	191
Radiology	2,799	250	Respiratory	419	150
Neurosurgery	518	237	Paediatric Medicine	233	148

Table 4: Surgical Specialties, Case Load, and Patient Care Pathways

the top ten largest surgical specialties. **NB** A *patient* may be admitted more than once. A *case* refers to an admittance. Multiple cases may relate to one patient. Table 3 refers to patients and Table 4 refers to cases.

Example of carepaths are shown in Table 5. Overall, 4,142 distinct carepaths were identified.

Patient Group	Carepth ID	Resource	Occupancy (hrs)
Paediatric Surgery SN	2780	OT	3.05
Paediatric Surgery SN	2780	WBED	29.05
Anaesthetic SN	10	OT	2.5
Anaesthetic SN	10	WBED	279.36

Table 5: Sample carepaths

Allocations: these are the planned allocation of hospital resources to the various patient groups. The hospital planners provided its Master Surgical Schedule, a 4 week roster specifying on any day, for each theatre, hours of theatre usage by surgical specialty. The Master Surgery Schedule is not prescriptive, however, and may be interrupted for various reasons (e.g. emergency cases or unavailability of the surgeon for whom the room had been reserved). As some surgical specialties require operating theatres only infrequently, not all specialties are included in the Master Surgery Schedule and will reserve time by negotiation. Similarly, the hospital tries to use certain wards and locations for particular dedicated functions. For our study, we used the Master Surgery Schedule to determine allocations for specialties included in the schedule, and then interrogated the event log to find the usual allocations for other patient groups. Allocations are recorded in terms of hrs/week for OT resources and physical space, i.e., number of beds for other resources. No such detailed information was available for resources such as ward beds. Allocations for ward bed and ICU bed usage was estimated by sampling and averaging the actual ward bed usage as recorded in the event log.

5. Findings

5.1. Data Analysis Through RapidMiner Interface

RapidMiner provided an interface where a workflow could be devised such that event data could be rapidly compiled and visualised. The process mining event log was the source data for all the analysis and subsequent visualisations. One aspect of interest to stakeholders was the actual operating theatre usage. A workflow was constructed that interrogated the event log for occurrences of the “Surgery End” activity. For this activity, event attributes included the surgical specialty, operating theatre, and surgery duration. Hence, it was possible to assess theatre usage by surgical specialty. Figure 4 exploits RapidMiner’s visualisation capabilities to show theatre usage across the study period for a selected theatre. The visualisation clearly shows that the theatre was mainly used by a single surgical specialty (Oph-

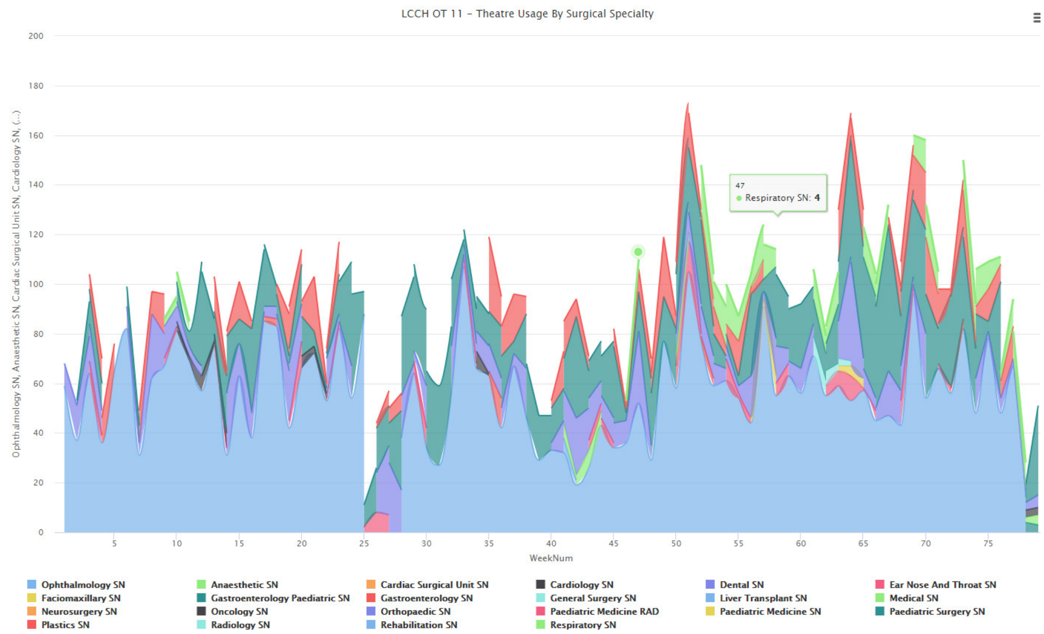

Figure 4: Theatre usage by Surgical Specialty Jul-2019 to Dec 2020 (Theatre - LCCH OT 11).

thalmology). At QCH, it is not uncommon for a single specialty to dominate a particular theatre, given that some theatres have specific, fixed equipment that's only used for that specialty. This is not to say they cannot be utilised for other cases, and indeed are, when the theatre is not occupied by the dominant specialty.

The workflow could also be used to analyse operating theatre time by surgical specialty. Figure 5 shows distinct differences between usual surgery durations across surgical specialties in a single operating theatre - LCCH OT 11.

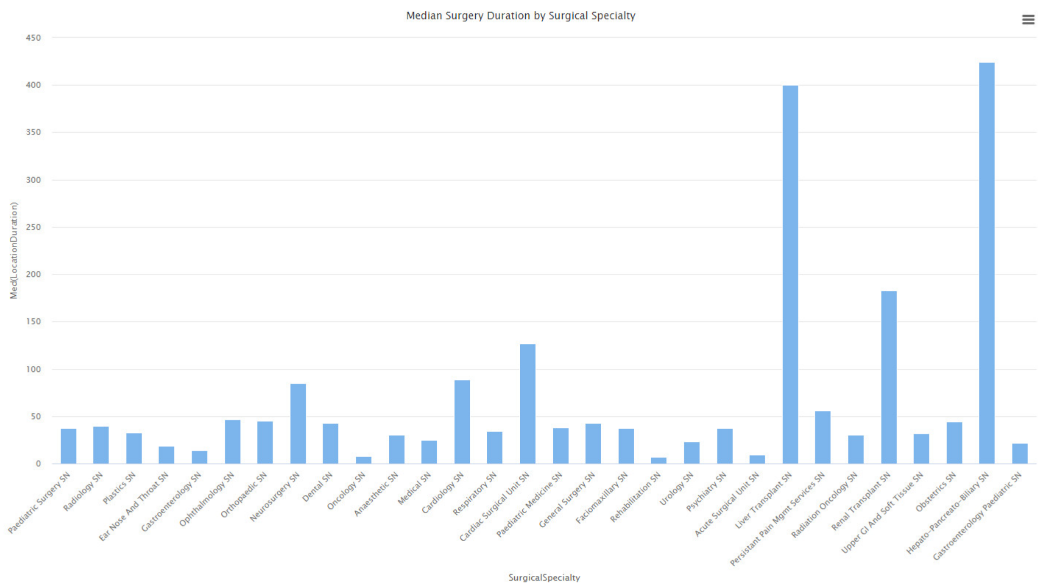

Figure 5: Median theatre duration per Surgery by Surgical Specialty Jul-2019 to Dec 2020 (Theatre - LCCH OT 11).

5.2. Proces Mining

The RapidProM extension of the RapidMiner suite was included in the automated workflow to generate a process model. RapidProM includes Inductive visual Miner which was used to generate various process models to highlight differences in process behaviour between various patient groups. We automatically discovered models for emergency and elective cases for the Ear Nose Throat (ENT) specialty.

Figure 6 shows the process model for emergency cases and Figure 7 displays the process model for elective cases. We note there are control flow differences between the ENT emergency and elective surgery process models. For elective cases, higher degree of parallelism is observed for surgery and anaesthesia than emergency cases.

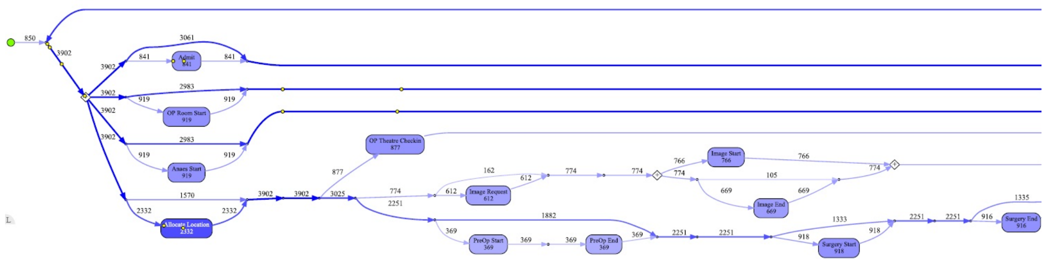

Figure 6: Process model derived from patient care data - Ear Nose Throat SN - Emergency cases.

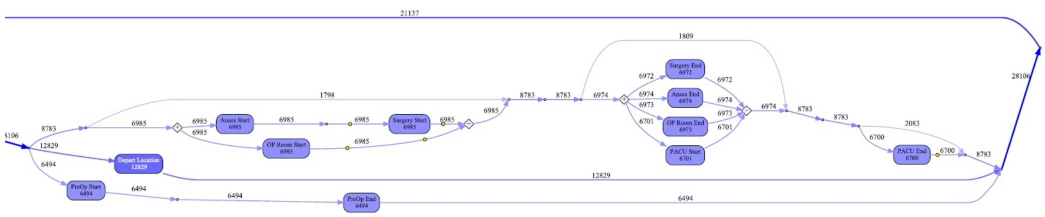

Figure 7: Process model derived from patient care data - Ear Nose Throat SN - Elective cases.

The process models are also able to highlight performance differences between the two patient groups. For instance, the average sojourn time (time in theatre) from surgery start to surgery end (26 mins 12 sec for emergency and 28 mins 59 sec for elective) as well as average PACU time (51 mins 08 sec for emergency and 53 mins 21 sec for elective) is similar for both cohorts. However, the average admission time (1 hr 46 min 58 sec for emergency and 31 sec for elective) is distinctly different. Such comparative analysis can be repeated for any specialty and patient group of interest.

5.3. Capacity Assessment

The capacity assessment model reports on utilisation of resources by (i) resource type, and (ii) patient group. Utilisation is as a fraction of the allocations made for the resource type and patient group. Overall resource utilisation is shown in Table 6. When applying the capacity assessment model, we specify a set of resources to saturate (in this case, the OT for surgical patients). The model then calculates the resulting utilisation of the other hospital resources. We note that the model is indicating ward bed utilisation lower than that expected by the hospital planners. Expectations were that saturating OT, the WBED would be over-saturated (i.e. utilisation in excess of 100%).There are several possible explanations for this, including, for example, modelling ward bed utilisation as actual hours of occupancy plus an allowance for a (minimum) bed change-over times (we did not include periods of time when a bed is not occupied by a patient but is nevertheless unavailable for use by another patient), moving towards bed-nights rather than bed-occupancy hours for characterising bed use in particular wards, or perhaps our representation of available hospital resources needs fine-tuning.

Resource Type	**Utilisation**
WBED	67.3%
OT	100.0%
ICU	40.8%

Table 6: Resource utilisation by resource type against allocations.

Utilisation by patient group shows utilisation of resources against the limiting resource. The limiting resource is the resource that governs the number of patients that may be handled by the hospital.

Table 7 presents some results for a small selection of patient groups. The capacity assessment model reports utilisation of resources against allocation where the limiting resource is fully utilised by the patient group.

Patient Group	**ICU**	**OT**	**WBED**
Paediatric Surgery SN	0.05%	100%	20.9%
Ear Nose And Throat SN	0.04%	100%	21.6%
Oncology SN	0.02%	100%	52.3%
Radiology SN		100%	24%
Non-surgical		0.0%	100%

Table 7: Resource utilisation against allocation by patient group.

5.4. Discussion

Our focus in this study was in using process data anaytics to support hospital capacity assessment and case-mix planning. As a starting point, process mining (discovery and comparative analysis) was used to provide insights into process behaviour and differences in performance between different patient groups. We noted that there were many similarities between event data/log and the inputs for capacity assessment. For instance, both utilise the notion of process (as patient journey through the hospital). Both utilise the notion of patient groups, and both rely on being able to attribute resources to patient care pathways. We therefore investigated whether process data could inform capacity assessment and case-mix planning.

We developed a basic capacity assessment model which took as inputs patient groups, resource groups, and care paths (sets of resources utilised by each patient group). We then developed techniques whereby these inputs could be automatically derived from the event log.

While planned usage of operating theatres was provided in the Master Surgical Schedule, planned allocations of ward beds, another major resource provided by the hospital, was not able to be provided by the hospital planners. Here we note some limitations affecting the results, particularly, relating to the allocations of resources to patient groups. The Master Surgery Schedule provided planned usage of the theatres for some, not all, of the surgical specialties. For surgical specialties not explicitly mentioned in the Master Surgical Schedule, a sampling approach was employed that derived average theatre hours over the 78 weeks of the study period for each such specialty. This was used as the allocation for these surgical specialties. For other resource types (WBED, ICU), sampling of occupancy of each location was taken at 4 times per day (6am, 9am, 3pm, 9pm) over the 78 weeks of the study. The average occupancy by patient group was taken as the allocation for that resource type and patient group. Actual occupancy durations derived from the event log were used in generating in-theatre hours and ward-bed occupancy times. The sampling approach is valid for large groups of patients however, allocations for small groups where individual patients' usage of the resource varies widely, are possibly not reflective of the actual allocation. Hence, utilisation against allocation for non-limiting resources for these groups may not accurately reflect the true resource usage.

We also observed, that when saturating OT usage, the utilisation of other resources, particularly ward beds, was different from that expected by the hospital planners, i.e., the expectation was that resource type WBED (ward bed) would be over-saturated (i.e. utilisation in excess of 100%). There are several possible explanations for this. For example, (i) ward bed occupation should include not only hours occupied, but include an additional allowance for bed change-over times, and (ii) moving towards bed-nights rather than hours for characterising bed use in particular wards. On-going, we will continue to work on this data/model in collaboration with QCH planners to improve the model.

6. Conclusion

The composition and volume of patients to be treated in a hospital, namely the patient case-mix, has a large impact on resource utilisation. In this paper, we reported on a project that used process data analytics, including process mining and a custom capacity assessment model, to inform hospital case-mix planning activities in a major Australian hospital (the Queensland Children's Hospital). We developed an integrated approach supported by an automated workflow. Our approach applied process mining to analyse historical process data to detect variations in process performance and behaviour across various patient groups. We also developed a novel capacity assessment model which calculated resource utilisation in the hospital where a selected resource (in our case, operating theatres) was saturated. We implemented the approach using the ProM process mining framework, and the RapidMiner scientific workflow environment (including the RapidProM environment).

We note that the approach can be enhanced by better use of process mining techniques to define patient groups and more fine-grained definition of resources. That is, rather than simply defining WBED (ward bed) as a resource, the actual ward, and individual bed in the ward, could be used. These enhancements will be the subject of future collaborative work with Queensland Children's Hospital. Refining the method for quantifying ward bed usage (for use by the capacity assessment model) that is more in line with hospital planners' conception of this particular resource is also an area for future work.

We believe that the approach, i.e., linking process mining and capacity assessment to inform case-mix planning, is a natural fit and promising. We recognise that the implemented approach was specific to Queensland Children's Hospital. However, all hospitals record, in some form or other, the data (admissions, ward movements, operating theatre usage, etc.) used as input in our implementation. We further note that Queensland is moving towards a 'digital hospital' environment across the larger hospitals in the state, with each hospital using the same hospital information systems. Hence, there is an opportunity to generalise our approach for use in multiple sites. This will be the basis for future work.

Acknowledgements. This research was funded by the Australian Research Council (ARC) Linkage Grant LP 180100542 and supported by the Princess Alexandra Hospital and the Queensland Children's Hospital in Brisbane, Australia.

References

[1] G. Ma, E. Demeulemeester, A multilevel integrative approach to hospital case mix and capacity planning, Computers & Operations Research 40 (9) (2013) 2198–2207.

[2] L. V. Green, Capacity planning and management in hospitals, in: Operations research and health care, Springer, 2005, pp. 15–41.

[3] N. Dellaert, E. Cayiroglu, J. Jeunet, Assessing and controlling the im-

pact of hospital capacity planning on the waiting time, International Journal of Production Research 54 (8) (2016) 2203–2214.

[4] B. Spratt, E. Kozan, An integrated rolling horizon approach to increase operating theatre efficiency, Journal of Scheduling (1808) 1–23.

[5] R. L. Burdett, E. Kozan, M. Sinnott, D. Cook, Y.-C. Tian, A mixed integer linear programing approach to perform hospital capacity assessments, Expert Systems with Applications 77 (2017) 170–188.

[6] B. Spratt, E. Kozan, A real-time reactive framework for the surgical case sequencing problem, Flexible Services and Manufacturing Journal (2019) 1–29.

[7] R. L. Burdett, E. Kozan, A multi-criteria approach for hospital capacity analysis, Eur. J. Oper. Res. 255 (2) (2016) 505–521.

[8] S. Hof, A. Fügener, J. Schoenfelder, J. O. Brunner, Case mix planning in hospitals: a review and future agenda, Health care management science 20 (2) (2017) 207–220.

[9] S. McRae, J. O. Brunner, J. F. Bard, Analyzing economies of scale and scope in hospitals by use of case mix planning, Health Care Management Science 23 (1) (2020) 80–101.

[10] F. Dexter, J. Ledolter, R. E. Wachtel, Tactical decision making for selective expansion of operating room resources incorporating financial criteria and uncertainty in subspecialties' future workloads, Anesthesia & Analgesia 100 (5) (2005) 1425–1432.

[11] M. S. Rauner, G. Schneider, K. Heidenberger, Reimbursement systems and regional inpatient allocation: a non-linear optimisation model, IMA Journal of Management Mathematics 16 (3) (2005) 217–237.

[12] N. Freeman, M. Zhao, S. Melouk, An iterative approach for case mix planning under uncertainty, Omega 76 (2018) 160–173.

[13] W. van der Aalst, Process Mining: Data Science in Action, Springer, 2016.

[14] M. Thiede, D. Fuerstenau, A. P. B. Barquet, How is process mining technology used by organizations? A systematic literature review of empirical studies, Business Process Management Journal 24 (4) (2018) 900–922.

[15] F. Emamjome, R. Andrews, A. H. ter Hofstede, A case study lens on process mining in practice, in: OTM Confederated International Conferences" On the Move to Meaningful Internet Systems", Springer, 2019, pp. 127–145.

[16] E. Rojas, J. Munoz-Gama, M. Sepúlveda, D. Capurro, Process mining in healthcare: A literature review, Journal of Biomedical Informatics 61 (2016) 224–236.

[17] P. Homayounfar, Process mining challenges in hospital information systems, in: 2012 Federated Conference on Computer Science and Information Systems (FedCSIS), IEEE, 2012, pp. 1135–1140.

[18] R. Mans, H. Schonenberg, G. Leonardi, S. Panzarasa, A. Cavallini,

S. Quaglini, W. Van Der Aalst, Process mining techniques: an application to stroke care, in: XXIst International Congress of the European Federation for Medical Informatics, 2008, pp. 573–578.

[19] R. S. Mans, M. Schonenberg, M. Song, W. M. van der Aalst, P. J. Bakker, Application of process mining in healthcare–a case study in a dutch hospital, in: International Joint Conference on Biomedical Engineering Systems and Technologies, Springer, 2008, pp. 425–438.

[20] R. S. Mans, W. M. van der Aalst, R. J. Vanwersch, A. J. Moleman, Process mining in healthcare: Data challenges when answering frequently posed questions, in: Process Support and Knowledge Representation in Health Care, Springer, 2012, pp. 140–153.

[21] W. Yang, Q. Su, Process mining for clinical pathway: Literature review and future directions, in: ICSSSM 2014, IEEE, 2014, pp. 1–5.

[22] A. P. Kurniati, O. Johnson, D. Hogg, G. Hall, Process mining in oncology: A literature review, in: ICICM 2016, IEEE, 2016, pp. 291–297.

[23] A. Partington, M. T. Wynn, S. Suriadi, C. Ouyang, J. Karnon, Process mining for clinical processes: A comparative analysis of four Australian hospitals, ACM Transactions on Management Information Systems (TMIS) 5 (4) (2015) 1–18.

[24] P. Badakhshan, A. Alibabaei, Using process mining for process analysis improvement in pre-hospital emergency, in: ICT for an Inclusive World, Springer, 2020, pp. 567–580.

[25] R. Mans, H. Reijers, M. van Genuchten, D. Wismeijer, Mining processes in dentistry, in: Proceedings of the 2nd ACM SIGHIT International Health Informatics Symposium, 2012, pp. 379–388.

[26] N. Martin, J. D. Weerdt, C. Fernández-Llatas, A. Gal, R. Gatta, G. Ibáñez, O. A. Johnson, F. Mannhardt, L. Marco-Ruiz, S. Mertens, J. Munoz-Gama, F. Seoane, J. Vanthienen, M. T. Wynn, D. B. Boilève, J. Bergs, M. Joosten-Melis, S. Schretlen, B. B. V. Acker, Recommendations for enhancing the usability and understandability of process mining in healthcare, Artif. Intell. Medicine 109 (2020) 101962.

[27] S. McRae, J. O. Brunner, Assessing the impact of uncertainty and the level of aggregation in case mix planning, Omega 97 (2020) 102086.

[28] P. J. Hulshof, N. Kortbeek, R. J. Boucherie, E. W. Hans, P. J. Bakker, Taxonomic classification of planning decisions in health care: a structured review of the state of the art in or/ms, Health systems 1 (2) (2012) 129–175.

[29] M. Utley, D. Worthington, R. Hall, Capacity Planning. Handbook of Healthcare System Scheduling, International Series in Operations Research and Management Science 168 (2012) 11–30.

[30] M. Jalalpour, Y. Gel, S. Levin, Forecasting demand for health services: Development of a publicly available toolbox, Operations research for health care 5 (2015) 1–9.

[31] E. Akcali, M. J. Coˆté, C. Lin, A network flow approach to optimizing

hospital bed capacity decisions, Health Care Management Science 9 (4) (2006) 391–404.

[32] B. Rechel, S. Wright, J. Barlow, M. McKee, Hospital capacity planning: from measuring stocks to modelling flows, Bulletin of the World Health Organization 88 (2010) 632–636.

[33] V. Valdmanis, P. Bernet, J. Moises, Hospital capacity, capability, and emergency preparedness, European Journal of Operational Research 207 (3) (2010) 1628–1634.

[34] X. Chen, L. Wang, J. Ding, N. Thomas, Patient flow scheduling and capacity planning in a smart hospital environment, IEEE Access 4 (2016) 135–148.

[35] L. Li, W. C. Benton, Hospital capacity management decisions: Emphasis on cost control and quality enhancement, Eur. J. Oper. Res. 146 (3) (2003) 596–614.