

FACILITATING DECISION SUPPORT IN HOSPITAL EMERGENCY DEPARTMENTS: A PROCESS-ORIENTED PERSPECTIVE

Ceglowski, Red; Churilov, Leonid; School of Business Systems Monash University
Melbourne, VIC, Australia, 3800, {red.ceglowski, leonid.churilov}@infotech.monash.edu

Jeff Wassertheil, Faculty of Medicine Monash University Melbourne, VIC, Australia, 3800,
Jeff.Wasserthiel@infotech.monash.edu.au

Abstract

Information systems exist for emergency departments (EDIS'), but even the most sophisticated ones concentrate on relatively simple coordination, resource allocation and documentation aspects of emergency department operations. There is little emphasis on management of the treatment process or optimization of resource use because definitive models do not exist for patient treatment processes. This paper outlines the identification of emergency department treatment processes and discusses how this treatment process perspective assists in framing optimization of resource utilisation, clinical decision making, training and emergency department layout.

Keywords: Decision support system, Emergency department, Data Mining, Process Modelling, EDIS.

1 INTRODUCTION

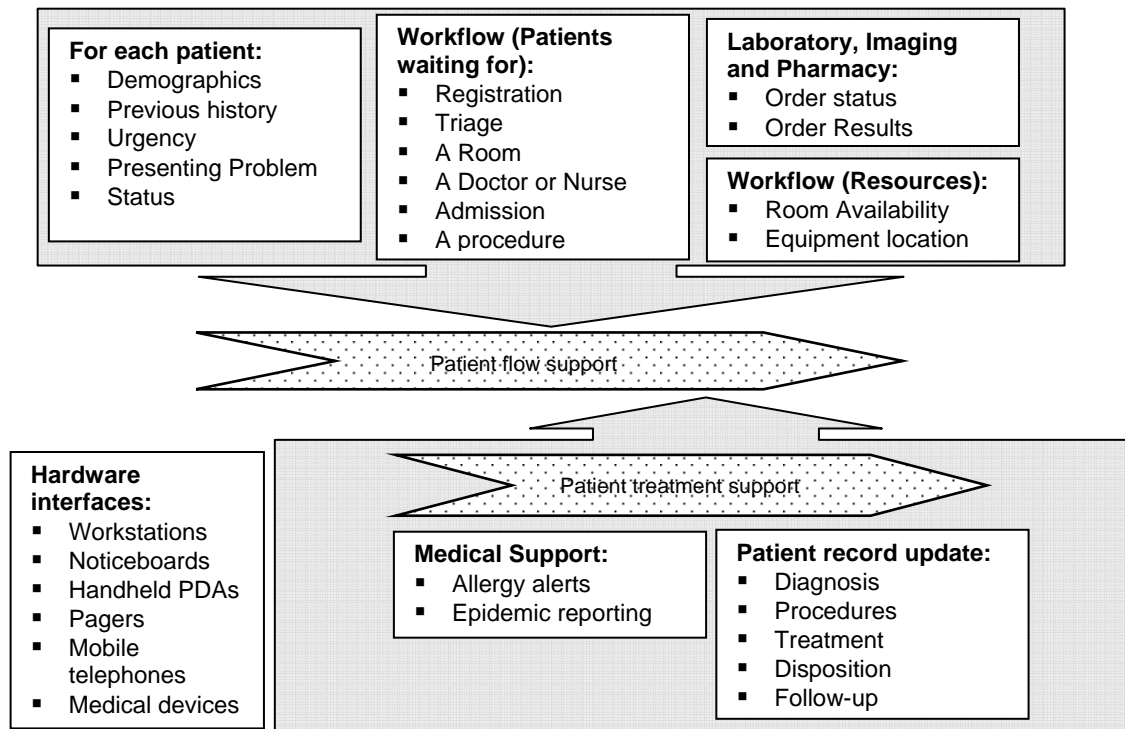
Hospital emergency departments are available at all hours of the day for people who need medical attention. Demand varies widely and every patient in the emergency department receives tailored treatment that must be provided within reasonable time. When combined with substantial increases in number of people presenting at emergency departments (23% in Victorian hospitals between 1998 and 2001 (Acute Health Division 2001)), this environment has resulted in instances where ambulances have been turned away and patient wait has been excessive. The health and political impact of these instigated efforts to ensure patient waiting and treatment times were minimised and to eliminate ambulance “bypass”. These efforts typically incorporated industrial engineering principles (Acute Health Division 2001; Salvendy 2001; Cooke et al. 2002; Department of Health 2003). While such efforts have met with some success, there are still large gaps in the understanding of emergency department operations. Overviews of patient flows that may facilitate decision support activities remain elusive.

The emergency department described in this work is typical in setting and complexity (Coleridge et al. 1993). There is a constant stream of patients into the emergency department with a range of ailments and urgencies. While the number of patients arriving each hour is reasonably well characterised, an investigation of the patient presentations for this hospital has shown that patients of any urgency may arrive at any time of day on any day of the year. They may be male or female, of any age and be injured or ill. It has not been possible to detect patterns in the arrival of ailments (presentations). Similar presentations may require vastly different reaction. For example, of two patients who have been stung by bees, one may show allergic reactions and require immediate response, while the other may merely be uncomfortable and may be treated subject to more urgent cases.

Some emergency departments are incorporating emergency department information systems (EDIS) to leverage off technology for efficiency and effectiveness improvements (The American College of Emergency Physicians 2003). EDIS provide a range of support from patient workflow management to electronic patient record facilities. While the level of support provided by EDIS is impressive, they focus on relatively simple coordination, resource allocation and documentation aspects of emergency department operations. EDIS design has been based on functions such as triage, bed or room allocation, nurse and doctor assessments, laboratory, drug and imaging ordering and coordination, and discharge-related documentation (Figure 1).

EDIS systems typically provide for data entry for triage, nursing assessment, doctor assessment and prescription management. Patient management is usually facilitated by provision of workflow modules that list patients awaiting treatment, their presenting problem and their severity. EDIS currently assist with patient tracking, workload management and record handling. Additional functionality may be provided through supplementary, particularly wireless, hardware that facilitates patient, patient record, test result and resource tracking. The EDIS may be interfaced with handheld tablets that can display patient records or accept nurse and doctor documentation, orders for prescriptions and follow-up instructions.

In theory an integrated EDIS can incorporate every step in the patient care process. Human handoffs can be automated. Each step can automatically be logged and tracked. Timing of steps can be determined and acceptable variation in timing and sequencing specified. Human interactions with



networked electronic devices such as personal computers, CT scanners, lab systems, telephones, IV pumps, and wireless patient tracking tags can be linked to the information system for automation of process control. Physicians entering orders, nurses bar-coding medications, clerks registering patients, and surgeons scheduling surgery can be linked and coordinated automatically as they perform their own specialties. Workflow engines sequence, monitor, track, alert, and reroute any step in each of the patient care processes (Rucker 2003). While such a vision of EDIS is laudable and possibly achievable in the future, current EDIS fall well short of this image. EDIS do promote reduction of wait, transfer and rework times by improving coordination of staff and resources, but the systems do not optimise patient flow and resource use, nor do they address decision support for clinical aspects of patient treatment. The modelling of emergency department patient flows is a problem at the centre of

Figure 1 EDIS support for emergency department patient flow

these inability. Models of emergency department treatment processes do not exist, so EDIS cannot predict what the next step in any process is likely to be, greatly handicapping their ability to actively support decision making.

It is the *objective* of this paper to present a new approach for the modelling of patient flow in emergency departments and to show how this may be incorporated into decision support systems that enhance existing emergency department information systems. We describe how knowledge of patient treatment and subsequent patient flow may be incorporated into systems that extend management of, and decision support for, patient treatment. We also discuss how this enhancement may provide information and evidence for knowledge management and decision support advances in emergency departments.

In the past, models of emergency departments have attempted to group patients according to demographic variables (Jelinek 1995a; Bond et al. 1998) or on high level patient flows (Walley 2003) using simulation, industrial engineering and medical casemix concepts (Cameron et al. 1990; Jelinek 1995b; Averill et al. 1998; Walley et al. 2001). Such views of emergency department patients are difficult because of the complexity of symptoms, range of severities and variety of medical specialisations involved in treatment. Each patient is different in seemingly unpredictable ways so treatment has to be individually customized. However, without some grouping of patients or

classification of treatment it is difficult to move EDIS beyond the generic ordering, recording and monitoring support for individual patients indicated in Figure 1.

In a process-centric representation, emergency department operations may be viewed as a series of value-adding functions (Figure 2). These functions describe the flow of patients through the emergency department from arrival to departure. While many of these functions align to those in EDIS, this representation portrays the inherent sequence of activities. EDIS do not support patient treatment processes because grouping of emergency department patients according to process similarity (iso-process grouping) (Visser 2002) does not exist. Every patient treatment is viewed as unique. Any attempt to manage the complexity of the emergency department by applying additional coordination mechanisms to cater for the variety of presentations, treatment locations, staff and resources may lead to a situation where overhead costs for the additional control systems surpass the benefit of efficient coordination of every variant (Becker et al. 2003). Iso-process grouping of patients has been particularly difficult in the case of emergency department patients because of the broad range of demographics and presentations (Jelinek 1995b; Walley et al. 2001), and this has undoubtedly limited EDIS advances in this regard. A new technique called self-organised process mining has allowed iso-process grouping of emergency department patients (Ceglowski et al. 2004b). This technique will be described in the next section and some results presented. Applications of these treatment groups are discussed in Section 3 and some conclusions close the paper in Section 4.

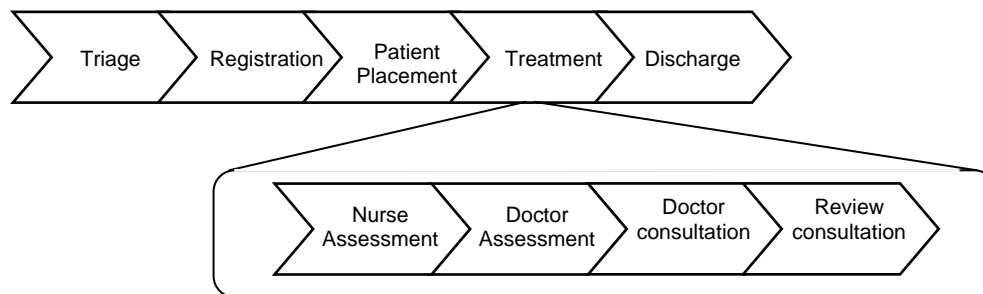


Figure 2 Core emergency department value adding functions (Djorhan and Churilov 2003)

2 SELF ORGANISED PROCESS MINING

This section describes the use of a new method called self organised process mining for detection of process from data.

Even though there is a wide range of patients and presentations, much of the work in the emergency department is based on applications of a short list of medical procedures. Patient observation, drug orders, laboratory and imaging investigations are examples of such procedures. Just 36 procedures account for 99% of all procedures in Victorian hospitals. Within this almost 17% are classed as “Other”, which includes observation of patients by medical staff; 6% are “No procedures”; some 10% are drug administration and over 9% X-ray imaging. Other significant procedures are venipuncture, intravenous catheter access in preparation for infusion of fluid or drugs, and echocardiogram diagnostics (figures derived from Victorian Emergency Medical database for 2002).

The hypothesis employed in this research was that variance in patient treatment most often related to the combination and timing of particular clinical procedures. The intention was to group patients on similarity of combination and timing of clinical procedures. This approach suggested analysis of the procedures that patients underwent on each visit in an attempt to determine common clusters of treatment. These “similarity of treatment” clusters were expected to indicate the primary patient treatments undertaken in the emergency department. The method that was used for identification of “similarity of treatment” clusters was purely data-driven. Patient data was explored for non-obvious insights using non-parametric clustering techniques which make no assumptions about the distribution of patient characteristics or inter-relationships. Analysis was restricted to cases where two or more procedures were recorded since analysis of activities associated with single-procedure patients could be done separately. Random samples of around 10 000 cases were extracted of patients who were

treated in the emergency department and whose records indicated two or more procedures involved in treatment.

The data was cleaned prior to sampling of obvious noise and inconsistencies that related to dates, residence times in the emergency department and errors such as letters in numeric fields. Software limitations forced reduction of the number of input variables to less than 50. Since the data carried 57 medical procedures, ten least common procedures were eliminated. This involved less than 1% of all records. The data was then comprised of a case identifier and 47 procedures. The procedures were recorded as integer counts, with zero indicating absence of a procedure. It was possible for a patient to receive repeated applications of a procedure. In practice this was not the case, except for a generic “observation” procedure which was often repeated. Thus each row of data had an identifier followed by essentially a binary string interjected by the counts between 1 and 5 for the “Observation” procedure variable. This did have the effect of weighting “Observation” over other variables, but the large number of variables dampened this effect so it was not considered beneficial to normalise the data.

| SUT Cluster for patients treated and discharged | | |
|---|---|------|
| Matching records (%) | | 5.35 |
| CT | CT Scan | 0.01 |
| DRS | Dressing | 0.82 |
| DRUG | Oral, sublingual, topical, rectal drug administration | 0.49 |
| ECG | 12 lead ECG | 0.02 |
| ECGM | ECG monitoring | - |
| FWT | Full ward test urine | 0.01 |
| HIO | Head injury observation | - |
| INF | Infusion IV fluid (ex blood products) | 0.01 |
| IV | Peripheral IV catheter | 0.04 |
| IVI | IV drug infusion | - |
| NEB | Nebulised medication | - |
| O | Observation | 0.39 |
| POP | Plaster-of-Paris | - |
| RBG | Random blood glucose | - |
| SUT | Sutures | 1.00 |
| ULS | Ultrasound | - |
| VB | Venipuncture | 0.03 |
| XRAY | X-ray imaging | 0.14 |

Table 1 Procedures in treatment of emergency department patients requiring sutures

The clustering method used in this work, Self Organised Mapping, has been used extensively on this and similar emergency department datasets (Ceglowski et al. 2004a; Ceglowski et al. 2004b). The Self Organising Map (SOM) nonparametric method is algorithm-driven and relies on data, rather than domain-specific expertise. The method seeks to minimise diversity within groups and to maximise between group differences. These differences are determined using a distance metric to compute relative distance of cases from one another. Self-Organizing Maps provide a visual understanding of patterns in data through a two dimensional representation of all variables. Cases that have similar characteristics are adjacent in the map, and dissimilar cases are situated at a distance determined by degree of dissimilarity. The SOM algorithm repeatedly repositions cases in the map until a classification error function is minimised. The method employs large datasets, works well with many input variables and produces “arbitrarily complex models unlimited by human comprehension” (Kennedy et al. 1998). Viscovery SOMine, the software tool used in this analysis, employs a variant of Kohonen’s Batch-SOM (Kohonen 1995), enhanced with a scaling technique for speeding up the learning process (Viscovery SOMine 1999). SOM-Ward clustering was used (Ward Systems Group 1996). This clustering method combines the local order information of the map with the classical hierarchical cluster algorithm of Ward (Ward 1963).

Self Organised Mapping generated 18 clusters that accounted for the vast majority of patient treatment with acceptably small quantisation error and map distortion. While many clusters appeared in retrospect to be simplistic grouping of procedures (such as the linking of X-rays to fractures and plaster of Paris), other less obvious groups of procedures were identified. One example of a core treatment process is provided in Table 1. This process is characterised by patients who receive sutures accompanied by a specific set of medical procedures. From a business process perspective patients receiving similar treatment are identical, despite having different medical designations. The treatment clusters describe treatment process and so create an opportunity for analysis and prediction of emergency department operations. Some of the insights generated from these clusters will be discussed in the next section.

3 TREATMENT CLUSTERS ENHANCE DECISION SUPPORT

In the Introduction it was discussed how information systems currently targeted at emergency departments fail to proactively support patient flow and resource allocation because they approach emergency operations from a clinical rather than process perspective. In other words patients tend to be grouped according to arrival sequence, urgency, demographic variables and diagnosis. Such groupings do not facilitate decision support activities because they have no predictive function, but they do provide insight into patient severity. Knowing a patient's urgency or age does not help determine their treatment or diagnosis, for instance, but may give an indication of resource requirements. Treatment clusters do have predictive properties that may be combined with traditional clinical groupings to provide a high level of decision support.

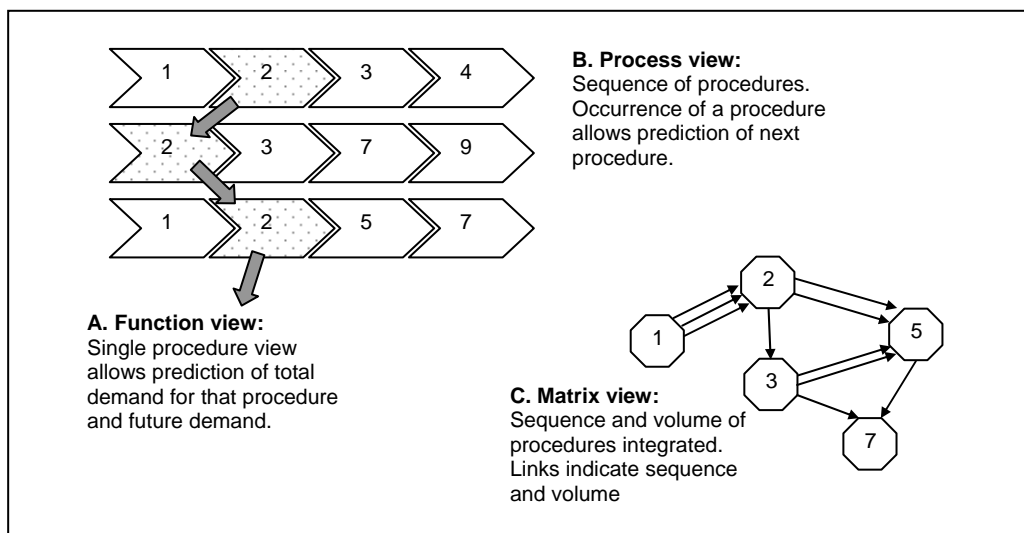


Figure 3: Three co-ordination views

Treatment clusters derived from scrupulously maintained emergency department records provide a picture of actual treatment being performed in emergency departments. Classifying patients according to the treatment they receive reduces the uncertainty associated with emergency department operations. Where demand (patient arrival) was indeterminate with respect to type of presentation (ailment) a host of methods now come available to predict what treatment the next patient is likely to need and what procedure any patient is likely to need next. Previously the view was that treatment needed to be tailored to every patient. It may now be possible to place patients into “treatment classes” that can be measured and managed using the idea of the “focused factory” (De Vries et al. 1999). A decision support system loaded with information about the pathways patients will follow will be able to provide several views of emergency department operations (Figure 3).

3.1 The Function View

A “Function” view analyses operations according to individual procedures and yields the workload due to a specific procedure at any given time. Two useful applications arise from this view. The first application is demand forecasting. Recall that the number of patient presentations is well

characterised and can be predicted for any time of day with a reasonable level of confidence. Patient presentation forecasts could be combined with data about procedure use to derive a predicted demand for each procedure. It is likely that such a model would have a sliding window of accuracy (that is it is likely to be more accurate for short-term forecasts) and would have to be tuned to some extent. Ultimately, however, it's reasonable to expect that the system could be combined with resource or other constraint information that enable it to detect whether the emergency department is susceptible to becoming overloaded. Such models would be a great advance on the measures currently in place that rely on holding hospital occupancy below a threshold figure (Lane et al. 2000). This forecasting application of the function view could also be used to give material requirements and trigger material ordering reminders.

Secondly, consider current thinking that it is impossible to predict what ailment each next patient into the emergency department is likely to present. The number of patients may be predictable, but the procedure demand is unknown, so empirical thinking has to be employed in order to make staffing decisions. The health-critical nature of work performed in emergency departments means that these empirical decisions have to include a large margin of safety, but there is no means of measuring the risk associated with any decision. The function view can help in planning and in directing training to reduce risk associated with a shortage of a specific skill set. The skills essential for each procedure can be listed. This can be linked to personnel who have the necessary skills to give an overview of the emergency department knowledge base, used to inform a training schedule or used in staff scheduling.

3.2 The Process View

Remember that there is a fairly limited set of procedures recorded in the emergency department. The existing set accounts for most treatment activities (although the "observation" or "other" class may need further elucidation). This short list of procedures provides an avenue for a basic EDIS documentation function where the triage nurse, attending physician or other qualified person accesses a list of procedures and selects the ones appropriate to management and documentation of each patient. Such a Case Handing approach (van der Aalst et al. 2003) fails to provide advice regarding treatment, nor intelligence about future demands on resources. All support is retrospective. Diagnostic tests are ordered when needed, and procedures are ordered or documented after the fact. There is no support for semantic analyses that may be provided by a sound ontological ordering of documentation.

A progression from iterative selection of individual procedures and retrospective support may be achieved by combining procedures into treatment sequences. The treatment clusters described in Section 2 give the likelihood of a patient undergoing procedures. The "Process" view in Figure 3 represents this sequence of procedures. Unfortunately, treatment clusters do not at this time carry sequence information, although it's expected that rules can be built up over time (x-rays always precede plaster of Paris, for example). Even without knowing how procedures are related in time it's possible to use treatment clusters to derive the set of procedures most likely to succeed the current one. As more procedures are completed this becomes easier.

One practical implementation of treatment cluster information is indicated in Figure 4. Triage nurses often order interventions that should be performed by nurses on duty once the patient is transferred to




Figure 4: Patients may be streamed into treatment clusters as more data is gathered

a cubicle (Djorhan and Churilov 2003). As such they have a good idea of the procedures that will be performed on a patient. It should be possible for patients to be allocated to a subgroup of treatment clusters at triage. When this information is integrated with other patient data in a decision support system it will be possible to isolate a discrete number of pathways each patient is likely to follow in the course of their treatment. Resource requirements, transport needs and potential problems can start to be identified at this stage. As patient treatment proceeds the set of possible pathways becomes reduced until the patient is regarded as an instance of a single treatment pathway. Markov or Bayesian models for patient state transitions, both within clusters and through the emergency department could facilitate optimisation of patient placement and movement activities.

Decision support systems need to be configurable (Shim et al. 2002). Configurable systems arise from a desire to build general purpose systems that can readily be specifically tailored for individual needs. This issue is most readily apparent in the implementation of generic enterprise systems, where the most challenging aspect is often adaptation of the basic system to specific needs (Johnston 2002). Enterprise systems are often described through reference models that describe the business processes and structure of the system through use of function, data and organisation artefacts. Application reference models depict all possible system capabilities but do not guide system configuration for particular instances (Rosemann and van der Aalst 2003). The concept of Figure 4 provides an avenue for configuring EDIS to not only support the particular treatment practices at an emergency department, but to evolve over time as more is learnt about specific treatment pathways. Such configurable decision support is inconceivable in the context of existing EDIS.

3.3 The Matrix View

The “Matrix” view in Figure 3 is a combination of the function and process views. It carries information about the quantity and type of procedures currently in progress and models of how the picture will change in the near future. This view is the “Holy Grail” of decision support in the emergency department. Ideally it would be able to “pre-book” resources, optimise patient placement and transfers and maximise resource use, as well as providing a range of likely scenarios to assist with staffing and other management decisions. But, even with perfect information about treatment process, it can be seen that such optimisation problems would be complex to solve in a single instance, let alone “on the fly” to keep up with the dynamic situation in the emergency department.

It’s naïve, however, to think that solution of the matrix view is impossible. Consider a restaurant analogue for the emergency department. Number of customers is uncertain and the mix of demand unknown (the meals that the customers will choose). Restaurant staff have a range of specialist roles, such as reception, waiting on tables, recommending wine, clearing tables, preparing specialities of food, and so on, in much the same way the emergency departments have a range of specialists and support staff. Restaurants have tables or stations whose use needs to be optimised, similar to patient beds and certain resources such as x-ray machines that service the emergency department. Food courses should be served in particular order, just as procedures have specific sequence during treatment. The co-ordination problem has successfully been addressed for restaurants. Treatment clusters start to make similar solutions available for hospital emergency departments.

3.4 Extensions and other issues

The function, process and matrix views provided a framework for discussion of several applications in the sections above. There are several extensions, and more research that must be done in order to fully realise the decision support potential of this work.

Hospital administrators will see the value that a process-based view has for Activity Based Costing (ABC). Whereas proxies such as “doctors time” have been used as cost drivers in ABC models, treatment processes now provide true “activity” cost drivers for costing models. A costing model has been built that uses Monte Carlo sampling of the range of procedure costs combined with information about procedures in clusters and patient presentations. This is being modified as a more accurate picture of procedure costs, procedure synergies and treatment-specific resource requirements become apparent.

Treatment based analyses of emergency department operations have the potential to alter the design and layout of emergency departments (Ceglowski et al. 2004c). The existence of process-based pathways through the emergency department allows industrial engineering approaches to be fully implemented. Optimal location of equipment and materials, modified staff travel pathways and behaviour and other typical production-type problems all become addressable.

While the vision for future applications of treatment clusters appears bright, there are several important issues that need to be addressed. In comparing treatment clusters across a number of different emergency departments it was seen that the selection and ratio of procedures was sometimes different (Ceglowski et al. 2004a). This may be due to clustering, coding or treatment issues. Clustering differences can arise because specific characteristics in the data drive the clustering optimisation algorithm towards separate solutions for different data sets. These can usually be detected and overcome by adjustment of parameters. Coding issues relate to interpretation of procedure definitions. These may vary from hospital to hospital but can be addressed by clearer, agreed guidelines. Treatment differences may arise because of different patient profiles or innovation in treatment at certain campuses. Cross-campus comparisons provide the opportunity for knowledge management and knowledge transfer across campuses so that clinical decision support systems can be constantly updated.

Associated with this knowledge management concept is the idea that treatment can be “scored” and consistency of treatment measured for training and performance management (Gunning and Rowan 1999). Scoring systems often include variables such as age and severity, so casemix classifications may be integrated with treatment clusters during development of such systems. The treatment cluster work described in this paper is currently being extended so that symptoms, diagnosis (where available) and treatment may be compared. It is expected that valuable information will become available in the course of this analysis that further elucidates emergency department treatment processes.

Another significant issue pertaining to this work is that treatment clusters have not been verified in situ. Time has to be spent following patients in an actual emergency department to detect how the recorded procedure data matches actual treatment practises. At present, drug administration is separated only by method of administration (oral, sublingual, intravenous, nebulised and so on), not by type of drug. Treatment clusters should carry more detail about the type of drug, but this may place an unacceptable data-entry burden on staff. There is currently a single large procedure named either “observation” or “other”, depending on the database. The component activities need to be isolated and data captured separately about them so that treatment clusters can be further clarified.

Possibly the greatest impediment to process-based decision support, however, is the belief in medical circles that definition of treatment process, prediction of resource requirements and delivery specifications are not possible. While exact identification of these elements may not yet be possible, treatment clusters move in the correct direction and it may be that the objections raised by clinicians are “a cover for not having to accept guidelines and protocols to defend medical autonomy” (De Vries et al. 1999).

4 CONCLUSIONS

This paper proposed to provide insight into how existing emergency department information systems might be extended to provide active decision support. It placed the emergency department in context in an uncertain and complex environment and went on to review the level of support currently provided by specialist information systems, including those integrated with networked, wireless and mobile technologies. It was seen that opportunities for decision support are not being exploited because process-based models do not exist for the core activity of emergency departments – patient treatment.

A new method for extraction of process information was introduced and the resulting treatment clusters of this self organised process mining described as well as some of the other findings that resulted from this new perspective. Three views arose from the process nature of the treatment clusters. Decision support applications of these function, process and matrix views were postulated. While some of these applications may immediately be realised, others, such as integrated optimisation of equipment, people and materials, may be some way off. Training, performance and emergency department layout consequences of the treatment process approach were touched upon.

It may be concluded that treatment clusters present an invigorating opportunity for enhancements to EDIS. It must be remembered, however, that the healthcare industry is necessarily cautious about implementing changes, and emergency departments among the most cautious of all, as may be expected from the imperative to care for patients above all else. There are two essential elements of further research required before testing the theories in a live emergency department setting. The first is the confirmation of treatment clusters at hospitals outside the state and country where this research was conducted. While initial explorations across multiple hospital campuses have confirmed the viability of the process approach, in-depth studies have yet to be conducted, and the researchers are actively seeking collaborators in other countries who wish to pursue similar explorations.

The second area in which this research needs to be extended is in the development and testing of a prototype system. A unique simulation model has been built to support prototype development (Ceglowski et al. 2005). It is planned to integrate this simulation model with a decision support system built around the ideas contained in this paper. This will permit the exploration of decision support scenarios in a safe, simulated environment prior to testing in the emergency department. It is expected that the strategies for emergency department roll-out will be similar to those employed by San Pedro et al (2005), where nurses have the ability to overrule the system at any time and learn from Sanchez' (2004) clinical IT implementation lessons.

It must also be realised that Information Technology change has been “more rapid outside healthcare than within healthcare. Only recently, for example, did healthcare organisations begin to develop e-health sites for access by their patients. IS researchers can benefit healthcare by applying expertise gained in other domains to address challenges that are still new to the health informatics discipline. The IS discipline can benefit as well, both by testing its theories and methods in healthcare settings and through cross-pollination with health informatics expertise.” (Wilson and Lankton 2004). While existence of common and specific treatment processes are exciting from the IS perspective, clinicians may not share the enthusiasm, burdened as they are with the responsibility for human lives. The treatment decision support described in this paper may bring closer emergency department information systems where “The personal, moral, and legal responsibility for timely care no longer rests solely with the doctor or nurse” (Rucker 2003), but such systems can only be developed with the support and sanction of healthcare professionals.

Acknowledgement

Thanks to the referees who's helpful and constructive comments helped form an improved final version of this paper and provided insights valuable for future works.

References

Acute Health Division (2001) *Emergency Demand Management: A new approach* Melbourne, Victorian Government Department of Human Resources: Feb

- Averill R F, Muldoon J M, Vertrees J C, Goldfield N I, Mullin R L, Fineran E C, Zhang M Z, Steinbeck B and Grant T (1998) *The evolution of casemix measurement using Diagnosis Related Groups (DRGs)* Wallingford, CT, 3M Health Information Systems
- Becker J, Kugeler M and Rosemann M (2003) *Process management: a guide for the design of business processes* Berlin ; Heidelberg ; New York, Springer
- Bond M, Baggoley C, Erwich-Nijhout M and Phillips D (1998) "Urgency, disposition and age groups: A casemix model for Emergency Departments" *Emergency Medicine* **10** (2): 103-110
- Cameron J, Baraff L and Sekhon R (1990) "Case-mix classification for emergency departments" *Medical Care* **28**: 146-158
- Ceglowski A, Churilov L and Wassertheil J (2004a) "Process Focused Patient Clustering for a Hospital Emergency Department" *Casemix 2004* Sydney
- Ceglowski A, Churilov L and Wassertheil J (2005) "Data rich, socially responsible OR: Process modelling in hospital emergency departments" *IFORS Triennial Conference* Hawaii, INFORMS
- Ceglowski A, Churilov L and Wasserthiel J (2004b) "Data driven process modelling for a hospital emergency department" *Computer Supported Activity Co-ordination Workshop at ICEIS 2004* Porto, Portugal
- Ceglowski A, Churilov L and Wasserthiel J (2004c) "Process mining informed industrial engineering in hospital emergency departments" *APIEMS 2004* Gold Coast, Australia
- Coleridge J T, Cameron P A, White J B and Epstein J (1993) "Profile of emergency department attendances" *Emergency Medicine* **5** (1): 18-25
- Cooke M W, Wilson S and Pearson S (2002) "The effect of a separate stream for minor injuries on accident and emergency department waiting times" *Emerg Med J* **19** (1): 28-30
- De Vries G, Bertrand J W M and Vissers J M H (1999) "Design requirements for health care production control systems" *Production Planning and Control* **10** (6): 559-569
- Department of Health (2003) *Improving the flow of emergency admissions* London, NHS Modernisation Agency
- Djorhan V and Churilov L (2003) "Business interactions in an acute emergency department in Australia: A clinical process modeling perspective" *The 6th Pacific Asia Conference on Information Systems (PACIS 2002)* Tokyo, The Japan Society for Management Information (JASMIN)
- Gunning K and Rowan K (1999) "ABC of intensive care: outcome data and scoring systems." *BMJ [Clinical Research Ed.] [NLM - MEDLINE]* **319** (7204): 241
- Jelinek G A (1995a) Doctor of medicine *Casemix classification of patients attending hospital emergency departments in Perth, Western Australia: Development and evaluation of an urgency-based casemix system* Perth, University of Western Australia
- Jelinek G A (1995b) *A Casemix information system for Australian Hospital Emergency departments* Perth, A Report to the Commissioner of Health, Western Australia
- Johnston S J (2002) *ERP: Payoffs and Pitfalls*
<http://hbsworkingknowledge.hbs.edu/item.jhtml?id=3141&t=technology> Harvard Business School Working Knowledge Weekly Newsletter Accessed: 5 Mar 2003
- Kennedy R, Lee Y, Van Roy B, Reed C and Lippmann R (1998) *Solving data mining problems through pattern recognition*, Prentice Hall
- Kohonen T (1995) *Self-organizing maps* Berlin, Springer
- Lane D C, Monefeldt C and Rosenhead J (2000) "Looking in the wrong place for healthcare improvements: A system dynamics study of an accident and emergency department" *Journal of the Operational Research Society* **51**: 518-531
- Neuroshell 2 3.0 (1996) Frederick, MD Ward Systems Group

- Rosemann M and van der Aalst W M P (2003) *A Configurable Reference Modelling Language* Brisbane, Queensland University of Technology: May
- Rucker D (2003) *Finally, A Tool to Re-Engineer Health Care: The Workflow Engine* San Francisco, Montgomery Research, Inc.: 7 Jan
- Salvendy G (2001) *Handbook of industrial engineering : technology and operations management* New York, Wiley
- San Pedro J, Burstein F, Wassertheil J, Arora N, Churilov L and Zaslavsky A c (2005) "On Development and Evaluation of Prototype Mobile Decision Support for Hospital Triage" *The 38th Hawaii International Conference on System Sciences*, IEEE
- Sánchez A (2004) " A Chronic Wound Healing Information Technology System: Design, Testing, and Evaluation in Clinic" *Electronic Journal of Information Systems Evaluation* **7** (1): 57-
- Shim J P, Warkentin M, Courtney J F, Power D J, Sharda R and Carlsson C (2002) "Past, present, and future of decision support technology" *Decision Support Systems* **33**: 111-126
- The American College of Emergency Physicians (2003) *Emergency Department Technology* <http://www.acep.org/> The American College of Emergency Physicians Accessed: 31 May 2004
- van der Aalst W M P, ter Hofstede A H M and Weske M (2003). "Business Process Management: A Survey" *Business Process Management (Proceedings of BPM2003)*. Berlin, Springer-Verlag: 1-12
- Viscovery SOMine Standard Edition 3.0 (1999) Wien Eudaptics Software Gmbh
- Vissers J M H (2002). "A framework for production control in health care organisations" *Evidence based management in health care : the role of decision support systems* Boldy D, Braithwaite J and Forbes I Eds. Kensington, N.S.W., School of Public Health & Community Medicine University of New South Wales: viii, 146
- Walley P (2003) "Designing the accident and emergency system: lessons from manufacturing" *Emerg Med J* **20** (2): 126-130
- Walley P, Watt A, Davies C, Huang A and Ma K (2001) *A study of demand for emergency access health services in two UK health regions* Warwick, Warwick Business School
- Ward J (1963) "Hierarchical grouping to optimise an objective function" *Journal of the American Statistical Association* **58**: 236-244
- Wilson E V and Lankton N K (2004) "Interdisciplinary Research And Publication Opportunities In Information Systems And Healthcare" *Communications of the Association for Information Systems* **14**: 332-343