Identification of Surgery Indicators by Mining Hospital Data: A Preliminary Study

Marie Persson and Niklas Lavesson

Abstract—The management of patient referrals is an interesting issue when it comes to predicting future patient demand to increase hospital productivity. In general, a patient is referred from the general practitioner to hospital care. A patient referral contains information that indicates the need for hospital care and this information is differently structured for different medical needs. In practice, these needs can be viewed as the forthcoming patient demand at the hospital, analogous to a volume of orders. Today, the structure of the referrals is very much up to the general practitioner who is referring the patient. This implies that the data provided to the hospital can vary extensively between cases. We suggest that, by enforcing a certain structure on the referral data, it may be possible to make early predictions about the patient demand. Such predictions could then be used as a basis for managing resources more efficiently to increase hospital productivity. This paper investigates the possibility of using data mining techniques to automatically generate prediction models by extracting conclusive information from patient records combined with surgical suite statistics, e.g., surgery preparations and anesthesia type, that are of significance for estimating patient demand in a surgery department, e.g., probability of surgery, surgery duration and recovery. We hypothesize that the generated models may provide new knowledge about, and a basis for, how to structure a patient referral. In addition, these models may also be used for the actual prediction of patient demand.

Index Terms— Learning systems, Medical diagnosis, Medical services, Prediction methods

I. INTRODUCTION

SIMILARLY to many European countries today, Sweden is struggling with an increasing healthcare cost. The political pressure to reduce this cost together with a growing patient demand, provide strong incentives to enhance productivity and cost-efficiency. One of the most expensive areas in healthcare is surgery, which necessitates many expensive resources in terms of staff, equipment, and medical resources. Generally, these resources have to be managed and divided between several surgical departments, e.g., orthopedics, gynecology or the general surgery department, within the hospital in order to meet the total surgery demand

(depicted in Fig. 1). The question is how and when the surgery demand should be assessed. Commonly, the patient queue, i.e., the waiting list for surgery, is viewed as the surgery demand. Let us review a quite common sequence of steps for processing a patient within the Swedish healthcare system: typically, the patient begins by contacting the general practitioner. If deemed necessary, the patient is then referred from the general practitioner to the hospital care. The referral is assessed by an expert, i.e., a surgeon specialist, and the patient is then put on wait for an appointment to meet a surgeon with the appropriate subspecialty. Next, the patient meets the surgeon at the outpatient clinic, i.e., the hospital care, and together they decide upon treatment, e.g., surgery. If surgery is decided, the patient is added to the patient queue for surgery. At this point, when the patient has been added to the waiting list, the accumulated surgery demand can be estimated and hence provide information to facilitate and improve surgery management.

Fig. 1. The surgery demand. Several departments are managed by the surgical suite, which is responsible for operating room scheduling and resource allocation.

A. Aims and Objectives

In this paper, we investigate whether it is possible to make (early) predictions about the surgery demand by using data mining based models that are automatically generated from patient records in combination with surgery related statistics and the associated known outcomes. We argue that, if these predictions can be made at the referral stage, the resource management and productivity related to surgery can be improved. The paper reports on a preliminary study, which serves to lay the foundation for a larger project, which involves collaborations with domain experts at a Swedish hospital. The work will be continued by obtaining actual data from the hospital in order to further develop the theory and to perform validations by using empirical experiments.

B. Outline

The remainder of this paper is organized as follows: first, we describe the problem more thoroughly from a healthcare

Manuscript received March 24, 2009. This work was supported in part by Blekinge County Council, Sweden.

M. Persson is with Blekinge Institute of Technology, School of Computing, Box 520, SE-372 25, Ronneby, Sweden (phone: +46-457-385838; fax: +46-455-385057; e-mail: Marie.Persson@bth.se).

N. Lavesson is with Blekinge Institute of Technology, School of Computing, Box 520, SE-372 25, Ronneby, Sweden (e-mail: Niklas.Lavesson@bth.se).

point of view. We then review related work. In Section III, our approach to address the problem is described and this is followed by a theoretical analysis. Finally, we present conclusions and some pointers to future work.

II. BACKGROUND

In Sweden, there is an intense debate about, and work conducted on, the systematization and standardization of patient records, as well as of patient referral procedures from the general practitioner to hospital care. Consequently, the discussions about this referral standardization are of great interest for both the general practitioners and the hospital care.

The hospital care wants the referred patients to be more extensively and suitably examined on by the general practitioner. In turn, the general practitioners require guidelines about which examinations should be conducted before referral. In theory, this problem may seem to have a trivial solution. However, considering the fact that the causes for, and data provided in, the patient referral can be quite dissimilar between different cases; this would imply that the problem is not easily solved in practice. Additionally, a recent study conducted in Sweden reveals that patients with identical diagnoses from two geographically close Swedish counties were managed using significantly different measures by the county hospitals [2]. The authors present a variable, denoted surgery indication, which can be calculated given the appropriate information for different types of surgeries. However, they note that most referrals or patient records do not include the necessary information.

A. Referral Contents

The main question is what kind of information should be included in the patient referral. This is indeed an interesting and difficult question. Despite the fact that a patient referral is strongly related to the symptoms and/or disorder/disease of the patient, the possibility of relying on diagnosis codes may be dismissed for two primary reasons: firstly, the patient might not have been given a diagnosis because the general practitioner has not been able to identify the disorder/disease with any degree of certainty. In many cases that is after all partly the motivation for having the patient referred to the hospital care. Secondly, the diagnosis of a patient may not in itself provide the necessary information on how to treat a patient. In other words, a diagnosis is a description of what is wrong rather than a description of how to treat a patient.

There are often several other variables of interest that need to be known in order to provide effective treatment. Consequently, such variables would also be of interest when trying to predict future healthcare demand. The diagnosis, if available, is of course one of these variables. However, the question is which additional variables are important for different types of cases and how information about these variables may be extracted.

Let us review an example case, where the described problem occurs. Consider a patient, referred from a general practitioner to an orthopedic department at the hospital (as visualized in Fig. 2).

Fig. 2. A conceptual view of the studied healthcare process. The referrals from general practitioners are submitted to the hospital care.

One of the main interests of the orthopedic department is that of knowing or being able to quickly finding out what type of subspecialty is required, e.g., hip, knee, back, and so forth. Additionally, it is crucial to determine whether the patient needs surgery or not. In order to determine this, prerequisite medical examinations have to be conducted. Many of these medical examinations can be performed by the general practitioner. This may also be beneficial in terms of cost since hospital care is more expensive. Additionally, the patients will be prevented from burdening the surgery waiting list if the results of the medical examinations contraindicate surgery or if the prerequisite medical examination is incomplete and therefore has to be complemented at the hospital. Moreover, it seems that the number of patients having surgery in comparison to the patients referred to the hospital care is lower than preferred. Interviews with domain experts suggest that only 30-40% of the referred patients are having surgery and that this number should instead be around 80% in order to achieve cost efficiency.

What type of information is crucial for the orthopedic department or the surgeon in order to accurately predict the probability of surgery depends on the particular case. Each subspecialty requires a different set of data. However, we point out that all subspecialties have a collective need for several types of data, such as: age, sex, and laboratory results. If it would be possible to determine the necessary information for a set of generic problems under which more specific problems can be sorted, it seems plausible to assume that the prediction can be made more accurately. If sufficiently good predictions can be made, the upcoming surgery demand for different types of surgery can be estimated. This means that a more efficient resource planning can be performed. If one type of surgery seems to increase in demand, certain actions can be taken to prevent waiting list congestions, e.g., operating room planning, scheduling of surgeons and/or more long-term actions like directions for education or employments to meet the approaching surgery demand.

Accordingly, we want to find out which variables are of interest for an orthopedic department in order to perform predictions about the future patient demand. It should be observed that these predictions are by no means meant to serve as a substitution for the visit to the outpatient clinic. A physical visit at the outpatient clinic before surgery, to meet the surgeon, is in most cases a prerequisite for surgery. Indeed it is also vital for patient trust.

We have identified three target variables that are of great importance when making patient demand predictions: a) probability of surgery, b) surgery duration, and c) surgery recovery or patient length of stay. Consequently, we strive to identify which types of patient data have the potential to be

good indicators for these three variables. Obviously, if the probability of surgery is determined to be low, the two remaining variables are not of interest for the orthopedic department.

B. Related Work

Quite a few studies have been conducted within the relatively new scientific field of healthcare management. There has been a contrast between the rather haphazard approach of healthcare management and the more scientific approach related to the medical profession [7]. Consequently, the field is now headed towards, what can be referred to as, evidence-based healthcare management. One of the main areas of research is healthcare informatics, which is about using methods from, e.g., computer science, data mining, and statistics on different types of data to develop decision support systems that can improve healthcare management [9].

Data mining methods in general have been studied extensively in this area; see, for example, [8] and [6]. However, studies about data mining based approaches to surgeon and operating room scheduling are scarcer. One of the studies related to this topic [1] compares the performance of four different machine learning techniques on the problem of operating room scheduling. Many types of input data were collected for this study. However, potentially important variables such as age, sex, and morbidity, were left unused.

It is important to stress that most studies about operating room scheduling are focusing on the patients that are already up for surgery. On the contrary, our paper investigates the possibility to make predictions about the patient demand for surgery much earlier; namely at the referral stage. To our knowledge, the problem of predicting the probability or occurrence of surgery based on referral data and patient records has not been studied previously.

III. APPROACH

We now suggest a novel approach for solving some of the main problems discussed in the previous section. The motivation for this approach is primarily based on the fact that large amounts of potentially useful data about patients, procedures, and consumed resources are stored in the hospital databases. These data can possibly be associated with data about, e.g., outcomes of treatments and decisions about whether to perform surgery or not. In context of a major restructuring project related to surgery management at a hospital in Sweden, this approach was introduced and received great interest.

Let us review the case under study: Blekinge hospital, which is a medium sized hospital, stores several types of structured data about the managed patients. At the surgical suite, additional important stored data include: the type of surgery, as well as the duration of the surgery among several other data related to surgery, e.g., anesthesia methods and surgery preparation. Moreover, a host of variables are stored in the patient records, e.g., diagnoses, drug treatments, whether surgery was decided upon for a certain problem, length of stay, and so forth.

We hypothesize that patterns may exist in the patient record data that can be used together with data from the surgical suite

to predict the occurrence, type and, duration of surgery and recovery. Such predictions could then be used by the surgeon specialist (here, orthopedic specialist) and surgery managers to schedule resources in terms of, e.g., staff, operating rooms, and beds more efficiently.

If the described patterns exist, it would be possible to generate models that indicate which variables or attributes are important to include in the referral for a particular problem. The generated models could also be used to predict the probability of surgery at the referral level. We therefore conclude that a data mining based approach seems to be quite suitable for this purpose. In particular, there is one type of technique that would be useful to employ.

C. Supervised Learning

If our assumptions are correct, it would be possible to apply supervised learning algorithms for the task of generating classifiers, i.e., by generalizing from patient record data and the associated occurrence, type, duration of surgery and recovery for known cases. A classifier is a generated model that can categorize a certain type of data into a discrete number of classes. Many types of classifiers are also able to output probabilities for the complete set of classes [5]. These probabilities may be used by the domain experts, e.g., the surgeon specialists, to reason about which decision to make when presented with several possible options. For example, it would most certainly be beneficial for a surgeon specialist to know that the predicted probability of the need for surgery is 0.48 (i.e., a borderline case) as opposed to just knowing that the predicted action is to not perform surgery.

Generally, a supervised learning algorithm is given a data set that features inputs and the associated output. The objective is then to generalize from these known observations of input and output. The resulting generalization is represented either by a classifier, if the output is discrete, e.g., the need of surgery (*true* or *false*), or a regression function, if it is continuous, e.g., the duration of recovery.

The input variables, or attributes, can be of different data types out of which two of the most frequently used types are the nominal (discrete) and the numeric (continuous) type.

In order to evaluate the generalization performance of the generated models it is crucial to divide the known data set into subsets of data that can be used for training and testing the models, respectively. Naturally, these subsets should be disjoint. Otherwise, the performance estimate may be overly optimistic [10].

D. Theoretical Modeling

It would of course be interesting to generate a data set using existing referral data as the input and, e.g., the occurrence of surgery as the output. We could then generate models that predict the occurrence or probability of surgery given the referral data. However, this only works in theory today since the currently used referrals are still represented by unstructured text, meaning that one would first need to extract pieces of information from these documents to transform them to structured representations, i.e., databases that could be mined. Of course, one could always employ text mining techniques on the unstructured data in order to extract useful information [3]. However, the lack of structure is not the only

problem, as we have already discussed: the main problem associated with the referrals is that they do not necessarily include the type of information requested by the specialist.

Instead, we would like to approach this problem from another angle by trying to find out what kind of information is indeed required by the hospital care for different types of referrals. From a surgeon specialist, we may obtain a set of generic referral types, $R = \{r_1, \dots, r_n\}$. For each type, we need to gather data from a large set of patients. Subsequently, for each patient, i , we need to obtain data about the target variable that is to be predicted at the referral stage. For example, we may obtain data from the hospital patient records about whether or not a surgery was carried out, s_i . Once we have gathered this data, it may even be possible to perform clustering to generate additional generic referral types; referral type definitions. A similar approach has been used to generate patient type definitions [4].

The difficulty lies in deciding on which input variables to include from the set of patient record variables, V , and this is indeed what we want to investigate, i.e., which variables, V_r , are good indicators of s for a certain referral type, r . Whether or not V represents the complete set of patient record variables or a subset, we would like use V as the set of input attributes, and s as the target attribute when applying the supervised learning algorithms. Our assumption is then that V_r can be extracted from the generated models on the condition that they are accurate in their prediction of s .

E. Models for Structured Referrals

If we can generate accurate models for the prediction of s , the variables from V_r can be used as a basis for structuring the referral of type r . Consequently, the general practitioner would know which pieces of information need to be included in each different type of referral. In fact, it would even be possible to design a software-based referral system, using domain expert knowledge. In this system, the general practitioner would specify the type of referral and in return the software application would describe what kind of information the general practitioner needs to provide and which tests and examinations to perform. Additionally, this type of interactive referral system may help the general practitioner in obtaining a better understanding of which possible actions e.g., surgery, that are needed to treat the patient and in taking decisions on whether it is useful or not to refer the patient.

F. Prediction Models for the Surgical Suite

Predicting the probability or occurrence of, surgery at the referral stage would simplify resource allocation and staff management at the surgical suite and the orthopedic department since it would allow the surgeon specialist to rule out patients with a low predicted probability of surgery in order to decrease the patient queue. However, the patients that are likely to need surgery would still pose a difficult problem of optimization and resource allocation. We suggest that the patient surgery demand can be better and earlier estimated than what is managed presently in the healthcare process by applying supervised learning algorithms to generate surgery prediction models from patient record data associated with surgery duration or recovery as target variables.

G. Algorithm Selection

It is difficult at this preliminary stage to determine which algorithms to employ in order to best solve the problems described in sections E and F. The main reason behind this is that we first need to gain more knowledge about the actual databases in terms of the amount, complexity, and type of available data. However, it is still possible to discuss some general aspects that need to be considered when selecting algorithms.

When presented with the task of constructing models for structured referrals, we first need to generate accurate surgery prediction models from data available in the various hospital databases. As suggested earlier, we may then extract suitable surgery indicators for each individual referral type from these prediction models. The set of identified indicators for each referral type (along with any known or uncovered relationships between them) would serve as a basis for constructing the structured referral models. The first step in the data mining process would be to generate a data set that includes all possible variables that may be correlated with the target variable from the existing hospital databases. Feature selection can then be used in order to systematically reduce the effect of the curse of dimensionality and to generate robust prediction models.

With regard to the task of generating surgery prediction models, we have already indicated the possible benefit of using classifiers that are able to output probabilities in addition to the actual classification since this additional information may be used to make a more informed decision, e.g., about whether to perform surgery or not. Coincidentally, both of the tasks require the generation of surgery prediction models. However, the intended uses of the prediction models differ between the tasks. For the second task, the goal is the prediction model itself. The intended use of this model is to serve as decision support for the surgeon specialist when allocating resources and scheduling personnel. Considering this intended use, we hypothesize that a human understandable (interpretable) model might be more valuable than one that is overly complex or opaque (assuming that both models are equally accurate in their predictions).

For example, neural network-based models or ensemble models, although accurate and robust for many domains, are generally considered as opaque and incomprehensible. However, there are some possible approaches available to remedy this issue, e.g., by extracting transparent and comprehensive rule sets from the generated (opaque) models [11]. A more common approach is to make use of learning algorithms that are able to generate interpretable models, e.g., decision tree inducers or rule-based learners, although it is important to recognize the fact that the models generated by these learning algorithms are not always interpretable. For example, we may generally regard tree-based models as an intuitive means for us to understand the decision process but if the tree structure is too complex (e.g., because of the high number of nodes, the existence of reoccurring sub trees, and so on) it might be time consuming or even impossible to understand the decision process.

For the first task, the model is merely used as a means to identify potential indicators that can be used as a basis for

constructing a new referral system. Thus, we do not need to restrict ourselves to using learning algorithms that generate interpretable models. It is important, however, that we are able to understand (and perhaps rank) the importance of different input variables in predicting the target variable.

IV. CONCLUSION

This paper investigates issues related to the management of patient referrals from general practitioners to hospital care. In short, we hypothesize that hospital productivity could be improved if the surgical suite and the orthopedic department would be able to predict the patient surgery demand at an earlier stage than what is possible today. We therefore suggest supervised learning as an approach to find patterns in patient and surgical suite records that could be used to predict whether a certain patient would need surgery and, if so, what would be the duration of the surgery and the subsequent recovery. From the models generated by the learning algorithms, we further assume that it would be possible to extract information about which variables, or factors, are good indicators, i.e., predictors for a particular type of referral. Our conclusion is that such a list of factors, given for each referral type, would provide a good basis for developing a software-based patient referral system in which the general practitioner may specify the referral type and in return receive advice about which information to include in the referral.

For future work, we intend to further establish our theoretical models of the studied problem and validate these models empirically by performing experiments on actual data obtained from a medium sized hospital in Sweden.

ACKNOWLEDGMENT

We would like to thank Professor Johan Berglund at Blekinge Institute of Technology for his valuable comments and suggestions.

REFERENCES

- [1] S. I. Davies, "Machine learning at the operating room of the future: A comparison of machine learning techniques applied to operating room scheduling," Master's Thesis, Massachusetts Institute of Technology, Cambridge, 2004.
- [2] T. Ekblom, J. Bonnerstig, S. Vähärautio, T. Vikerfors, and L. Rytterberg, "Olika indikation för knäplastik vid olika ortopedkliniker," in *Läkartidningen*, vol. 106, no. 13, 2009.
- [3] R. Feldman and J. Sanger, *The text mining handbook Advanced Approaches in Analyzing Unstructured Data*, location: Cambridge University Press, 2006.
- [4] M. W. Isken and B. Rajagopalan, "Data mining to support simulation modeling of patient flow in hospitals," in *J. Medical Systems*, vol. 26, no. 2, 2002.
- [5] N. Lavesson and P. Davidsson, "Evaluating learning algorithms and classifiers," in *Intelligent Information & Database Systems*, vol. 1, no. 1, pp. 37-52, 2007.
- [6] M. K. Obenshain, "Application of data mining techniques to healthcare data," in *Statistics for Hospital Epidemiology*, vol. 25, no. 8, 2004.
- [7] Y. A. Ozcan and P. Smith, "Towards a science of the management of health care," in *Health Care Management Science*, vol. 1, 1998.
- [8] M. Silver, H.-C. Su, and S. B. Dolins, "Case study: How to apply data mining techniques in a healthcare data warehouse," in *J. Healthcare Information Management*, vol. 15, no. 2, 2001.
- [9] T. T. H. Wan, "Healthcare informatics research: From data to evidence-based management," in *J Medical Systems*, vol. 30, no. 1, 2006.

- [10] I. H. Witten and E. Frank, *Data Mining: Practical Machine Learning Tools and Techniques*, San Francisco, CA: Morgan Kaufmann Publishers, 2005.
- [11] U. Johansson, *Obtaining Accurate and Comprehensible Data Mining Models: An Evolutionary Approach*, Doctoral thesis, Linköping University, Sweden, 2007.