Proceedings of the 3rd INFORMS Workshop on Data Mining and Health Informatics (DM-HI 2008)
J. Li, D. Aleman, R. Sikora, eds.

A NURSING CARE PLAN RECOMMENDER SYSTEM USING A DATA MINING APPROACH

Lian Duan
Management Sciences Department
University of Iowa
Iowa City, IA
lian-duan@uiowa.edu

W. Nick Street
Management Sciences Department
University of Iowa
Iowa City, IA

Der-Fa Lu
College of Nursing
University of Iowa
Iowa City, IA

Abstract

Recommender systems have been extensively studied to present items such as movies, music, and books that are likely of interest to the user. We propose to use correlations among nursing diagnoses, outcomes, and interventions to create a recommender system for constructing nursing care plans. Nursing care plan recommender systems can provide clinical decision support, nursing education, clinical quality control, and serve as a complement to existing practice guidelines. In the current study, we used nursing diagnosis data to develop the methodology. Our system utilizes a prefix-tree structure common in itemset mining to construct a ranked list of suggested care plan items based on previously-entered items. Unlike common commercial systems, our system makes sequential recommendations based on user interaction, modifying a ranked list of suggested items at each step in care plan construction. We rank items based on traditional association-rule measures such as support and confidence, as well as a novel measure that anticipates which selections might improve the quality of future rankings. Since the multi-step nature of our recommendations presents problems for traditional evaluation measures, we also present a new evaluation method based on average ranking position and use it to test the effectiveness of different recommendation strategies.

Keywords: nursing care plan, recommender system, data mining.

1. Introduction and Related Work

According to a report published in 2000 by the Institute of Medicine, at least 44,000 and perhaps as many as 98,000 patients die in the hospital each year as a result of medical errors alone [7]. These data point to adverse healthcare events as the leading cause of deaths in the US. Total national costs are estimated to be between $37.6 billion and $50 billion for adverse events and between $17 billion and $29 billion for preventable adverse events [7]. To contribute to the effectiveness, safety and efficiency of nursing care, we propose a nursing care plan recommender system which can facilitate clinical decision support nursing education, clinical quality control, and serve as a complement to existing practice guidelines.

Clinical nurses use care plans to describe complex patient care phenomena. Nursing care plans in the U.S. commonly include three standardized nursing terminologies: nursing diagnoses, encoded using NANDA [11]; interventions, using NIC [4]; and outcomes, using NOC [10]. We propose a system to interactively provide a ranked list of suggested diagnoses, interventions and outcomes based on those the nurse has already entered. The list is ordered based on historical data from the care facility, and is updated dynamically as new items are entered.

The study of recommender systems is an important and problem-rich research area. It provides an abundance of practical applications that help users deal with information overload by providing personalized recommendations, content, and services. Examples of such applications include recommending books, CDs, and other products at Amazon.com, movies by MovieLens [9], and news at VERSIFI Technologies [2]. These systems focus on commercial activities which have some differences from clinical activities. First, commercial recommender systems base recommendations on the user's history because users' preferences are relatively consistent, while a clinical system makes less use of the user's historical information since the same patient may have different problems for each visit. Second, for commercial recommender systems, the users' behavior is more flexible. Purchases depend on whether or not they want an item, while nursing care plans depend on the needs of the patients. Third, for most e-commerce recommender systems, the recommendation problem is reduced to the problem of estimating ratings for items the user has not selected. We frame the recommendation problem for clinical systems as minimizing the ranking position of required selections in the presented list. Lin et al. [8] proposed a recommender system using association rules; however, it uses only support and confidence which aim at optimizing the current selection without considering future selections.

While we believe that the application of recommender technology to clinical nursing practice is new, there are several examples in the literature of nursing expert systems [6, 12]. However, clinical expert systems are constructed based on the knowledge of experienced nurses, creating a development bottleneck. Further, as patterns change across time, the encoded rules need to be updated manually. By using data mining methods, we can extract rules from historical data automatically instead of relying on expert knowledge, and handle changes to practice standards by extracting patterns within sliding windows. Furthermore, data mining methods can find unique patterns for each individual hospital which is more accurate for clinical quality control.

2. Methodology

To facilitate electronic health record input, we provide a list of all the possible selections. At each step, the nurse selects one required item from the list. Ideally, the item at the top of the list is required, and in general, we wish to rank-order the list such that the required items are as near the top as possible. After each selection, the selected item is removed from the ranking list, and the list is re-ordered. Here we use the commonly used measurements for association rules, such as support and confidence [5], to construct the ranking list. In addition, due to the step-by-step process, we use a novel measure that anticipates which selections might improve the quality of future rankings. Throughout this paper, we use N to denote the total number of care plans. The notation N(S) is used to denote the number of care plans which contain the itemset S.

The first measurement is *support*, the percentage of records in which the item appears. We use support to measure popularity and recommend the most popular selection to the user first. The support of a given item A is calculated as N(A)/N.

The second measurement is *confidence*, the probability of the item being chosen conditioned on the previous set of selected items. The confidence of a given item A, given the set S that has already been chosen, is calculated as N($S \cap A$) / N(S).

We also introduce a new measure termed *information value* or simply IV. To measure IV(A) we consider how "orderly" the list of conditional probabilities would be if A is chosen, and for that we use a variation of the entropy equation from information theory. Here, p_i is used to denote the confidence of the ith remaining selection after if A has been selected. The entropy for item A is calculated as $\frac{\sum_{i=1}^{k}(p_i * \log_2(p_i) + (1-p_i) * \log_2(1-p_i))}{k}$. Ideally, any p_i should be either 1 or 0, leading to an entropy of 0. In this case, we would be able to identify exactly the set of selections that must be chosen, given the current set of selections plus A. Conversely, the most undesirable case is a p_i of 0.5. In this case, we have no information about future selections, and the future ranking list would be chaotic. We promote the selection that has both the high probability to be chosen and low entropy to predict future selections. With this measurement, we strike a balance between the gain of the current selection and that of future selections. The information value of the possible selection A is calculated as confidence(A) * (1 - entropy(A)).

Regardless of the measurement used, the fundamental element of this system is to easily obtain the occurrence of any selection set. Getting the occurrence of a set relies on a top-down search in the subset lattice of the items. Here we use a prefix tree structure to quickly retrieve the occurrence of any selection set with less memory use. More details can be seen in [3].

3. Experiments

The dataset was extracted from a community hospital in the Midwest for the year 1996. Our experiments used 10,000 care plans as a training set and 5,000 care plans as a testing set.

We use two different types of evaluation mechanisms, called *random selection* and *greedy selection*. For random selection, we randomly select one item from the remaining items in the care plan and evaluate its ranking in the ordered list. For greedy selection, we always select the remaining care-plan item with the highest ranking in the list. Both of these can be seen as simulating human behavior. When all required items are near the top of the list, human selection behaves like greedy selection. If all the required items are low in the list, people are not patient enough to go through the list and would instead select the needed item in an alphabetic list. In this case human selection behaves more like random selection. Actual human selection is likely between the results of these two methods.

We compute the average ranking of selected items and report the results, averaged over five separate runs, in Table 1. In addition to support, confidence and IV, the lift measurement [5] was also tested, but gave poor results which are not shown.

Table 1: Average selection ranking

		1	2	3	4	5	Mean	Variance
Random Selection	Support	5.396	5.338	5.439	5.434	5.341	5.390	0.049
	Confidence	5.152	5.132	5.214	5.199	5.093	5.158	0.050
	Information Value	5.133	5.126	5.220	5.202	5.101	5.157	0.052
Greedy Selection	Support	4.320	4.292	4.397	4.382	4.287	4.336	0.051
	Confidence	3.905	3.909	3.990	3.998	3.897	3.940	0.050
	Information Value	3.895	3.898	3.986	3.988	3.880	3.929	0.053

The simple measure of support is used as the baseline for comparison, and both confidence and IV are better than support under both selection strategies. The comparison between confidence and information value is less obvious. Intuitively, confidence is the best measurement under the random selection strategy since the current selection does not affect future selections, and confidence focuses only on minimizing the current selection ranking. However, in the experiment the performance of information value is almost the same as that of confidence under random selection. In the greedy selection strategy, information value always does slightly better than confidence. The improvement is small but consistent. All differences are diluted by the existence of two disproportionately probable diagnoses that occur in nearly every care plan.

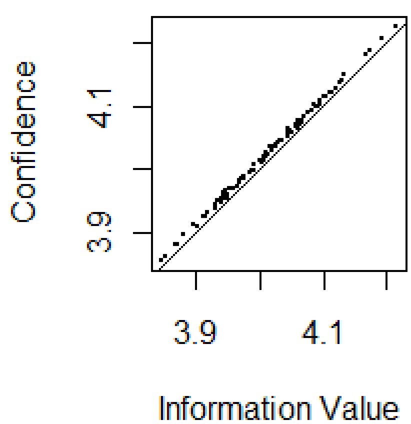

Figure 1: Information Value vs. Confidence

In order to examine the difference between confidence and information value in the greedy selection strategy, we repeat the experiment 100 times and compare the average ranking position of information value with that of confidence in the same experiment. In Figure 1, each point represents an experiment, the x-axis is the average ranking of information value, and the y-axis is the average ranking of confidence. Points above the line are experiments in which information value has a smaller average ranking than confidence. All the points in Figure 1 are above the line, i.e., information value outperforms confidence in each experiment. Moreover, information value has statistically significantly better performance (p = 1.70313E-60, using a pairwise t-test).

To examine what happens inside each method, we compute the average ranking of the selections in each iterative step of the selection process. In Figures 2 and 3, the x-axis represents the *i*-th step and the y-axis represents the average ranking value of choices made at that step. Under both greedy (Figure 2) and random (Figure 3) selection, both confidence and information value are consistently better than support. Since the performance difference between confidence and IV is difficult to see, we calculated the difference between them in each step, as shown in Figures 4 and 5. Under greedy selection, the performance of information value is constantly better than that of confidence, increasing through the 8th selection. After that, the improvement decreases but is still positive. However, no such pattern is evident under random selection, and overall there is no difference between the two values. Figures 4 and 5 support the conclusion that the performance of information value is almost the same as that of confidence in the random selection strategy consistently better than confidence under greedy selection.

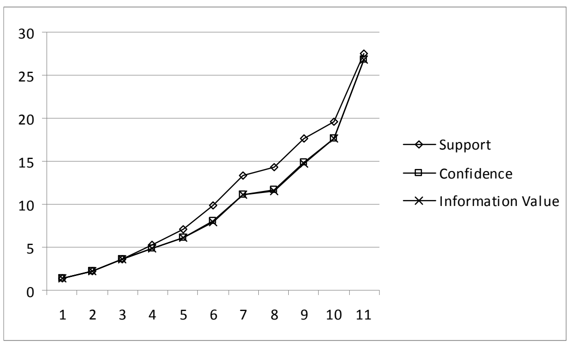

Figure 2. The average i-th step result of greedy selection

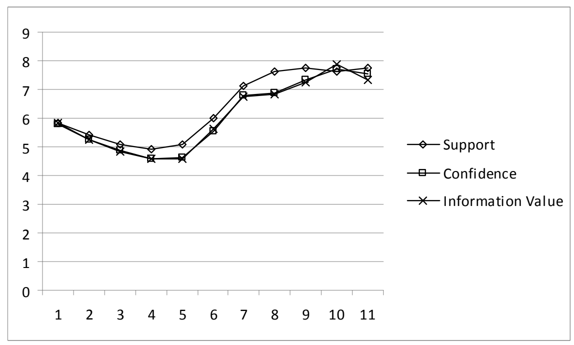

Figure 3. The average i-th step result of random selection

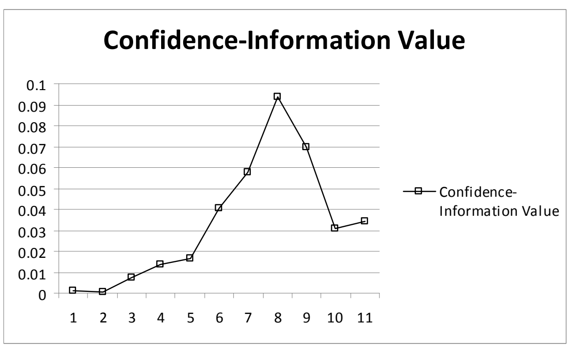

Figure 4. The difference in the i-th step result of greedy selection

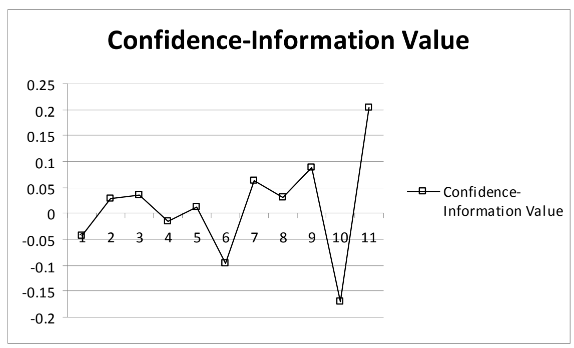

Figure 5. The difference in the i-th step result of random selection

4. Conclusion

We have described a new recommendation technique based on several measurements. In addition to traditional measurements support and confidence, we also test the effectiveness of a novel measurement - information value - which balances the gain between current and future selections. Its performance surpasses that of confidence, and it is still computable in real time. Such a system is a complement to expert systems and traditional practice guidelines, and can be very useful for nursing education and clinical quality control.

To date we have only experimented with relationships among diagnoses. In future work we will also examine the relationships among diagnoses, outcomes, and interventions. The effectiveness difference between expert systems and such a recommender is also interesting. Rules from experts' knowledge could be more accurate but they are not easily updated and specialized for different hospitals. Can we combine these two kinds of systems to achieve better results? Another promising direction is to incorporate contextual information into the recommendation process and make recommendations based on multiple dimensions, patient profiles, and other information [1]. Finally, we will examine the trade-off between confidence (immediate probability) and entropy (future probability) in the information value measurement, and adjust it to perform better on specific problems.

References

1. Adomavicius, G., Sankaranarayanan, R., Sen, S., and Tuzhilin, A., 2005. "Incorporating contextual information in recommender systems using a multidimensional approach," *ACM Trans. Information Systems* 23(1), 103-145.
2. Billsus, D., Brunk, C. A., Evans, C., Gladish, B., and Pazzani, M., 2002. "Adaptive interfaces for ubiquitous web access," *Commun. ACM* 45(5), 34-38.
3. Borgelt, C., 2003. "Efficient implementations of Apriori and Elcat," *Proc. Workshop on Frequent Item Set Mining Implementations*, Melbourne, FL, USA.
4. Dochterman, J., Bulechek, G.: Nursing Interventions Classification (NIC), 4th ed. St. Louis, MO: Mosby, 2004.
5. Han, J., and Kamber, M., 2006. "Data Mining: Concepts and Techniques," 2nd ed, Morgan Kaufmann.
6. Keenan, G. M., Yakel, E., and Marriott, D., 2006. "HANDS: A revitalized technology supported care planning method to improve nursing handoffs," *Proc. 9th Int. Congress on Nursing Informatics*, Seoul, South Korea, June 11-14, 2006.
7. Kohn, L., Corrigan, J., and Donaldson, M., 2000, "To error is human: Building a safer health system," Washington, DC: National Academy Press.
8. Lin, W., Alvarez, S. A., and Ruiz, C., 2002. "Efficient adaptive-support association rule mining for recommender systems," *Data Mining and Knowledge Discovery* 6(1), 83-105.
9. Miller, B. N., Albert, I., Lam, S. K., Konstan, J. A., and Riedl, J., 2003. "MovieLens unplugged: Experiences with an occasionally connected recommender system," In *Proc. 8th International Conference on Intelligent User Interfaces*, Miami, Florida, USA, January 12 - 15, 2003. ACM, New York, NY, 263-266.
10. Moorhead S, Johnson M, Maas M: Nursing Outcomes Classification (NOC), 3rd ed. St. Louis, MO: Mosby, 2004.
11. NANDA: Nursing Diagnoses: Definition & Classification. Philadelphia, PA: NANDA International, 2005.
12. Ryan, S. A., 1985. "An expert system for nursing practice," Journal of Medical Systems 9(1-2), 29-41.