Comparison of Two Knowledge Bases on the Detection of Drug-Drug Interactions

Guilherme Del Fiol, M.D., M.S.[1,2], Beatriz H.S.C. Rocha, M.D., Ph.D.[3], Gilad J. Kuperman, M.D., Ph.D.[4], David W. Bates, M.D., M.Sc.[4], Percy Nohama, Ph.D.[1]
[1] Pontificia Universidade Catolica do Parana, Curitiba, Brazil
[2] Assessoria de Informatica, Universidade Federal do Parana, Curitiba, Brazil
[3] Department of Internal Medicine, Universidade Federal do Parana, Curitiba, Brazil
[4] Center for Applied Medical Information Systems Research, Partners HealthCare System and Division of General Medicine, Brigham and Women's Hospital, Boston, MA.

This paper describes a drug ordering decision support system that helps with the prevention of adverse drug events by detecting drug-drug interactions in drug orders. The architecture of the system was devised in order to facilitate its use attached to physician order entry systems. The described model focuses in issues related to knowledge base maintenance and integration with external systems. Finally, a retrospective study was performed. Two knowledge bases, developed by different academic centers, were used to detect drug-drug interactions in a dataset with 37,237 drug prescriptions. The study concludes that the proposed knowledge base architecture enables content from other knowledge sources to be easily transferred and adapted to its structure. The study also suggests a method that can be used on the evaluation and refinement of the content of drug knowledge bases.

INTRODUCTION

Adverse drug events (ADEs) are known to be a major health problem worldwide. In the US, it is estimated that more than 7,000 people die from medication-related errors annually [1] and that the annual costs related to ADEs are approximately 136 billions dollars [2]. Additionally, medication errors are responsible for one out of 131 outpatient deaths [1]. In hospitals, the occurrence of ADEs increases patient length of stay from 2 to 5 days, and hospitalization costs from US$2,500 to 4,700[3,4]. Notably, it has also been demonstrated that about 28% of the ADEs resulted from errors that could have been prevented [3].
Physician order entry (POE) systems, coupled with decision support tools, have shown to be effective in reducing ADE rates and health care costs. The Brigham Women's Hospital, in a recent study, reported a reduction of $5-10 million dollars annually in hospital costs after the introduction of internally developed POE [5]. The same system, enhanced with decision support facilities that help with the drug ordering process, enabled a reduction of 80% in the incidence of medication errors [6].
Despite the evidences that reveal the ability of decision support systems in reducing medication errors, the use of these systems is still largely restricted to the academic medical centers where they were initially developed. Issues related to the portability and integration with other hospital information systems, knowledge base development and maintenance, and the establishment of formal evaluation methodologies must be addressed before a widespread use of drug ordering decision support systems can be achieved [7].
Although some drug ordering decision support systems have been developed using commercial drug knowledge bases, these knowledge sources often produce too many insignificant alerts, decreasing the confidence of users and increasing the likelihood that an important warning might be missed [8].
Drug-drug interactions (DDIs) are one of the causes of ADEs in hospitalized patients. Thousands of DDIs have been reported, but only a few are worth of attention. In addition, each site has its own preferences and needs. Thus, the same set of DDIs that suits to one site might not be entirely appropriate to others. Consequently, the elaboration of a DDI knowledge base for drug ordering decision support systems is not a trivial task. Besides the definition of a minimal set of DDI rules, it is also desirable that these rules can be adjusted for local needs and continually evaluated and improved.
This paper describes the evaluation of the performance of a drug ordering decision support system using two drug knowledge bases (DKB) developed in different academic medical centers. The objectives were to verify if the architecture proposed by the system allows the use of external knowledge bases, evaluate knowledge bases (KBs) performance,

1067—5027/00/$5.00

and observe how this kind of evaluation could help to improve the content of drug knowledge bases.

METHODS

The study was performed in three main phases. Initially, a drug ordering decision support system (DODSS) prototype was developed and a DDI rule knowledge base was created. In the second phase, a retrospective study using a dataset with inpatient drug prescriptions was performed to evaluate the system performance. In the third phase, rules from an external KB were transferred to our drug ordering decision support system and the retrospective study described in phase two was performed again with this new KB.

The Drug Ordering Decision Support System

The DODSS is a rule-based system that detects drug-drug interactions in drug orders. The system development had two main requirements: open architecture, with an interface module for integration with Hospital Information Systems; KB structured to make its maintenance a straightforward process.

The system has four main modules: a knowledge module, an inference module, an integration module, and a maintenance module. The knowledge module is a production rule knowledge base that represents drug-drug interactions. Three components constitute the knowledge module: a drug vocabulary, drug interaction rules and a drug category tree.

The vocabulary is a repository of codes and terms that represent the domain of drugs. The rules represent drug interactions using elements from different levels of specificity, i.e., from specific drugs to categories from the drug category tree. The drug category tree is a hierarchical configuration with multiple levels where drugs are classified according to common characteristics. The main attribute of the tree is that descendants inherit the characteristics of the parent categories. The representation of drug interactions by rules that use drug categories makes these rules as generic as possible. For example, the rule *first generation cephalosporins interact with aminoglycosides* is composed by two drug categories. Since there are dozens of drugs in each of these categories, a huge number of rules would be necessary to represent all possible interactions among them one-by-one. The whole KB was implemented in a relational database. This approach allowed rules to be easily included and modified by domain specialists, without the need of programming.

The inference module receives a pair of drugs from the integration module and, with facts and rules from the knowledge module, checks if there is an interaction between these drugs. If a drug interaction is found, its description is returned back to the integration module.

The integration module has an application programming interface that allows external systems, such as drug ordering systems, to access the DODSS. The integration module also performs mappings among external systems vocabularies and the knowledge module internal vocabulary.

The maintenance module is an application with a graphical user interface that allows rules, drug categories, and vocabulary items to be inspected, added and modified as needed.

The knowledge module rules were obtained with help from domain specialists from Hospital de Clinicas da Universidade Federal do Parana (HC-UFPR), a tertiary university hospital situated in the south of Brazil, and from the literature. In order to decrease the KB development time, its scope was reduced to moderate and severe DDIs involving cardiovascular agents and antibiotics, since these categories of drugs were considered by specialists to be responsible for most of the DDIs in our environment. In this phase, 207 rules were added to the knowledge base. All rules were included with support from the maintenance module.

Evaluation using the local KB

The DODSS was evaluated by a retrospective study. A dataset with 37,237 inpatient drugs prescriptions, ordered from January to March of 1999, was obtained from the HC-UFPR POE system, through which physicians make drug orders since 1994. The dataset was interfaced to the DODSS using the interface module and the POE vocabulary was mapped to the knowledge base vocabulary.

Evaluation using the BWH KB

The second KB was obtained from a set of DDI rules used by the Brigham Women's Hospital POE system and provided to us in an unstructured format. The Brigham Women's Hospital (BWH) rules were also included in the DODSS knowledge base with support from the maintenance module (326 rules were included). Since BWH rules had compatible levels of granularity with the DODSS vocabulary, the inclusion of each rule in the KB consisted on seeking the similar drug or drug category within the DODSS vocabulary items. A new item was included in the vocabulary if a drug or drug category was not found. After the inclusion process, all rules were checked for correctness.

At last, the BWH KB was evaluated by the same method applied to the local KB, using the same dataset.

RESULTS

The Brigham Women's Hospital DDI rules were easily adapted to the DODSS structure, requiring only the addition of 12 new drug categories in the drug category tree and 5 new drugs. The whole process took 13 man-hours, 6 hours for the inclusion of rules in the database and 7 hours for rule checking (Table 1).

Table 1 – Time required adapting BWH KB to the drug ordering decision support system.

Process	Time
Rule inclusion	6 hours
Rule checking	7 hours

Content of the knowledge bases

The scope of the KBs was quite different. The local KB focused on cardiovascular agents (89.9% of the rules) and antibiotics (19.3%), while BWH KB had a more general focus (Table 2).

Table 2 – Rules of both KBs classified according to the category of the drugs involved.

Drug Category	Local KB	BWH KB
Antibiotics	19.3%	19.0%
Antiviral drugs	-	20.9%
Cardiovascular agents	89.9%	15.0%
Oral anticoagulants	4.8%	8.0%

Results of the evaluation study

In inpatients, the decision support system, using the local KB, detected 16,880 drug-drug interactions among 37,237 orders, and 10,044 (27.0%) orders had at least one DDI. The mean duration of patient exposition to DDIs was 6.6 days. With BWH KB, the decision support system detected 4,863 DDIs and 4,283 (11.5%) orders had at least one DDI. The mean duration of patient exposure to DDIs was 5.3 days (Table 3).

Table 3 – Main outcomes.

Outcome	Local KB	BWH KB
Orders analyzed	37,237	37,237
DDIs detected	16,880	4,863
Orders with DDIs	10,044 (27.0%)	4,279 (11.5%)
DDIs Mean duration	6.6 days	5.3 days

The DDIs detected by both knowledge bases were classified according to their severity level (Table 4). The mean time patients were exposed to DDIs from each severity level was also obtained (Table 5).

Table 4 – Classification of detected DDIs according to severity level.

Severity level	Local KB	BWH KB
Mild	1,588 (9.4%)	-
Require monitoring	5,651 (33.5%)	1,139 (23.4%)
Moderate	8,553 (50.6%)	3,284 (67.5%)
Severe	1,088 (6.4%)	440 (9.0%)
Total	16,880	4,863

Table 5 – Mean duration of DDIs by severity level.

Severity level	Local KB	BWH KB
Mild	6.0	-
Require monitoring	6.7	5.4
Moderate	7.1	3.9
Severe	5.0	8.7

The DDIs were also classified according to the category of the drugs. The DDIs detected by the local KB were represented mainly by anti-hypertensive agents (51.2%) and antibiotics (33%) (Table 6). On the other side, DDIs detected by the BWH KB were most frequently anti-hypertensive agents (35.6%), antibiotics (21.9%), anticonvulsants (16.4%), and H2 blockers (14.6%) (Table 7).

Table 6 – Main drug categories of the DDIs detected by the local KB.

Drug category	Number of DDIs
Anti-hypertensives	8,649 (51.2%)
Antibiotics	5,563 (33.0%)
Cardiac glycosides	2,539 (15.0%)
Adrenocorticosteroids	2,400 (14.2%)
Diuretics	1,956 (11.7%)

Table 7 – Main drug categories of the DDIs detected by the BWH KB.

Drug category	Number of DDIs
Antibiotics	1,731 (35.6%)
Anti-hypertensives	1,066 (21.9%)
Anticonvulsants	797 (16.4%)
H2 antagonists	712 (14.6%)
Cardiac glycosides	372 (7.6%)

Some of the detected DDIs need special attention because they are potentially fatal. The BWH KB, for example, detected 53 (1.2%) interactions between azole antifungals and cisapride. The local KB did not detect these interactions. In addition, patients were exposed to such interaction for 10 days in average, almost twice as the overall exposition time.

Regarding the percentage of rules used in each KB,

28% of the local KB rules were used at least once and 10.3% of the rules were responsible for detecting more than 90% of the interactions. When analyzing the BWH KB, 16.3% of the rules were used in one or more occasions and 5.5% of the rules were responsible for more than 90% of the interactions (Table 8).

Table 8 – Percentage of rules used in the experiment by each knowledge base.

	Local KB	BWH KB
Rules used	58 (28%)	53 (16.3%)
Rules responsible for 90% of the interactions	23(10.3%)	18 (5.5%)

DISCUSSION

The DODSS architecture has proved to be quite open, allowing the inclusion of DDI rules from an external source easily and in a very short period of time. The feasibility of representing drug-drug interactions in a generic fashion, using drug categories, reduced the number of rules and, as a result, facilitated the inclusion of new rules in the knowledge base. Therefore, generic rules, using drug categories hierarchically structured have shown to be a suitable method for representing drug-drug interactions, as also suggested by others [9]. It is expected that this structure is also appropriate for other sorts of drug interactions, such as drug-allergy, drug-lab and drug-food. We plan to use and test the suitability of the developed structure to cover these domains in future studies.

In both of the knowledge bases evaluated, a small number of rules was responsible for the detection of most of the DDIs. These findings suggest that DKBs for the detection of drug-drug interactions need to, and possibly should, represent only a small subset of interactions reported in the literature to be effective. This fact demonstrates the feasibility of developing a minimal set of rules, representing the most important DDIs, that can be used as a starting point for any drug ordering decision support system.

The differences found between the performance of the two KBs can be attributed to a series of factors. First, the scope of the knowledge bases was different. While the local KB concentrated efforts on antibiotics and cardiovascular agents, the BWH KB had a wider and more generic scope. Second, the BWH KB has been in use for a longer time and in a real-time environment. So, its content has certainly been refined in order to satisfy user demands. On the other side, the local KB is merely a prototype and was tested exclusively against retrospective data sets, without feed-back from users that could further refine its content. Third, BWH rules were developed in order to attend local needs that certainly are not entirely applicable to Hospital de Clinicas needs. The drugs used in each hospital are different, medical practice is different and so are the patients. Supports this hypothesis the fact that only 16% of the BWH rules were actually used during the experiment. On the other hand, even the local KB had a low rate of rule utilization (28% of the rules were used). This data suggests that even a small but carefully selected set of rules would be able to detect a large amount of drug-drug interactions.

The incidence of drug-drug interactions obtained by both KBs is comparable to incidences reported by others [10,11,12].

Even though the BWH rules have shown to be quite different from our rules, it doesn't mean that an external and mature KB cannot be used as a "seed" to other KBs, as long as the original KB is "translated" to meet local needs. This process of refinement requires the definition of such local needs by in-house specialists. The methodology applied in our study seems to be suitable and useful for the refinement of DDI KBs. Retrospective datasets provide a faster feedback of the characteristics of local practices. Data gathered from the retrospective evaluations can be used to determine which classes of drugs are being considered by the KB and which classes are not. In that manner, new rules can be included covering drug categories that are not being contemplated properly by the KB. In addition, rules that represent DDIs with high incidences tend to produce too many alerts to the users. If such interactions are common practice in the environment where the system is in use and cost-effectiveness is being evaluated, the presence of these rules should be reconsidered or refined.

Due to the large amount of drug-drug interactions detected by both knowledge bases we expect to integrate, as soon as possible, the DODSS to the Physician Order Entry system of the HC-UFPR. In this context, the results of this study will be of great value for the refinement of our local KB.

CONCLUSION

Although there is a consensus indicating the need for the use of decision support systems to prevent adverse drug events, the use of these systems is still generally restricted to the academic centers where they were initially developed. One of the factors that hinder a widespread use of those systems is related to knowledge base content creation and maintenance.

In this study, we present a DODSS with a knowledge base architecture that facilitates the maintenance of its content. The proposed model has also proved of

being capable to fit content from an external knowledge source. Concluding, the study suggests a methodology for comparing drug knowledge bases that can also be used on the refinement of their content.

Acknowledgments

The authors are grateful to Drs. David Bates and Gilad Kuperman, from the Brigham Women's Hospital, for giving us access and rights to use the BWH drug-drug interaction rules in our experiment.

References

1. KOHN LT, CORRIGAN JM, DONALDSON MS, editors. To Err is Human: Building a Safer Health System. Washington, D.C.: National Academy Press; 1999.
2. JOHNSTON JA, BOOTMAN JL. Drug-related Morbidity and Mortality: a Cost of Illness Model. Arch Intern Med 1995;155:1949-1956.
3. BATES DW, CULLEN DJ, LAIRD N, et al. Incidence of Adverse Drug Events and Potential Adverse Drug Events. JAMA 1995;274:29-34.
4. CLASSEN DC, PESTOTNIK SL, EVANS RS, LLOYD JF, BURKE JP. Adverse Drug Events in Hospitalized Patients. JAMA 1997;277:301-6.
5. GLASER J, KUPERMAN GJ. Impact of Information Events on Medical Care. Proc of the Health Inf Manag Syst Soc Conf 1996;2:2-9.
6. BATES DW, TEICH, JM, LEE J, et al. The Impact of Computerized Physician Order Entry on Medication Error Prevention. J Am Med Inform Assoc 1999;6:313-21.
7. MILLER RA. Medical Diagnostic Decision Support Systems – Past, Present, and Future. J Am Med Inform Assoc 1994;1:8-27.
8. POIKONEN J, LEVENTHAL JM. Medication-Management Issues at the Point of Care. JHealth Inf Manag 1999;13(2):43-51.
9. PATERNO MD, TEICH JM, SEGER DL, BATES DW. A Practical Method for Presenting Drug Interactions to Clinicians. In: Proc AMIA Ann Fall Symp 1996:872.
10. COSTE B, SENE B, MILSTEIN C, BOUEE S, VENOT A. Indicators for the Automated Analysis of Drug Prescribing Quality. Methods of Information in Medicine 1998;37:38-44.
11. BRAVERMAN SE, HOWARD RS, BRYANT PR, BELANDRES PV. Potential Drug Interaction in a Physical Medicine and Rehabilitation Clinic. Am J Phys Med Rehabil 1996;75:44-49.
12. DAMBRO MR, KALLGREN MA. Drug Interactions In a Clinic Using Costar. Comput Biol Med 1988;18(1):31-38.