Towards a Collaborative Filtering Approach to Medication Reconciliation

Sharique Hasan, MS, George T. Duncan, PhD, Daniel B. Neill, PhD, Rema Padman, PhD
The Heinz School, Carnegie Mellon University, Pittsburgh, PA

Abstract

A physician's prescribing decisions depend on knowledge of the patient's medication list. This knowledge is often incomplete, and errors or omissions could result in adverse outcomes. To address this problem, the Joint Commission recommends medication reconciliation for creating a more accurate list of a patient's medications. In this paper, we develop techniques for automatic detection of omissions in medication lists, identifying drugs that the patient may be taking but are not on the patient's medication list. Our key insight is that this problem is analogous to the collaborative filtering framework increasingly used by online retailers to recommend relevant products to customers. The collaborative filtering approach enables a variety of solution techniques, including nearest neighbor and co-occurrence approaches. We evaluate the effectiveness of these approaches using medication data from a long-term care center in the Eastern US. Preliminary results suggest that this framework may become a valuable tool for medication reconciliation.

Key words: *Patient Safety, Data Quality, Medication Reconciliation, Collaborative Filtering, Machine Learning*

Introduction

Well-informed and safe medication prescribing decisions depend on a variety of inputs. These include a patient's demographic characteristics, diagnoses, allergies, and their current and past list of prescription and non-prescription medications. The health care provider incorporates this information with clinical knowledge to arrive at a decision about which drugs will best address the patient's ailments. The failure to access relevant information as well as the cognitive limitations of decision-makers can negatively affect the final prescribing decision, resulting in adverse drug events (ADEs).

A 2006 Institute of Medicine report estimates that errors in prescription and dispensation of medications cause 1.5 million preventable ADEs each year, with an estimated 800,000 in long-term care facilities alone [1]. Studies also suggest that 50 percent of medication errors in hospitals result from failure to reconcile medications [2]. Consequently, the Joint Commission on the Accreditation of Healthcare Organizations recommends medication reconciliation, the process of creating an accurate list of all medications a patient is taking, for reducing errors.

Many studies indicate that there are significant discrepancies between clinic-derived medication histories, admissions orders, patient self-reports and even claims data [3]. A recent study comparing self-reported drug consumption against medical records data found that 80.4% of patients had discrepancies, with nearly three discrepancies per patient [4]. Omissions of drugs from a patient's list constitute the majority of discrepancies, followed by commissions and other inconsistencies [5]. The medical consequences of discrepancies are not trivial. Related research suggests that discomfort, clinical deterioration, or death can occur in patients if these discrepancies are not adequately resolved [6].

Several strategies have been proposed for addressing the problem of discrepancies in medication lists [7]. Most of these strategies focus on improving processes through assignment of responsibilities, improving inter-organizational communication, and increasing access to information. The most common strategy involves using reconciliation forms at different points in the process and ensuring they are completed and verified. Form-based interventions do improve list accuracy, but often hindered by consistency of application and a failure to maintain the process [8].

Information technology also plays a role in medication reconciliation. Electronic medical records, prescribing systems, and computerized physician order entry are beneficial on several fronts. Above all, they provide a means to store medication data in a structured and easily accessible format. Many of these systems also incorporate decision-support components with pre-programmed rules that alert prescribers about potentially harmful interactions. However, the usefulness of these alerts depends on the accuracy of the stored patient information. In turn, the accuracy of the stored list depends on the robustness of the reconciliation process to various factors including a patient's accurate recall of their current drug regimen [9].

In many ways, the electronic medication record and form-based reconciliation are natural complements. The reconciliation form allows patients to report use of medications not currently stored in the electronic record. This information can be used to create a more accurate and accessible list for that

patient. The relationship in the other direction is less obvious. The electronic record is a repository of medication lists for hundreds or thousands of patients. At the most basic level, the electronic medication list for a patient consists of drugs recorded for that patient in a given clinical setting. Additional information in these records includes patient demographics, diagnoses, allergies and other pertinent health information such as laboratory test results. Even the most basic information, if processed in a sophisticated manner, can provide insights about potential discrepancies in a patient's medication list.

Additionally, electronic records provide us with the capacity to use medication information from a large population of patients in order to increase the accuracy of an individual patient's list. As a simple example, if every patient in the database who was prescribed drug A was also prescribed drug B, the occurrence of drug A on the patient's list suggests that drug B is also very likely to be present.

Collaborative filtering (CF) is a set of methods for processing information about users in order to make inferences or predictions about the information of other users. Many online retailers use CF to make predictions about products that an individual may enjoy based on aggregated information from users with similar observed tastes [10]. Successful applications of CF include the movie recommendations used by Netflix and product recommendations on Amazon.com.

In this paper, we outline an approach for using collaborative filtering as a methodology for the problem of medication reconciliation – specifically detecting omissions of medications from a patient's list. We use a variety of CF methods to answer the following question: *if a patient's medication list is incomplete, what drugs are most likely to be missing?* The methods produce an ordered list of drugs considered to have the highest likelihood of being omitted from a patient's record. In practical terms, the ordered list of potentially omitted drugs can be used to develop individualized memory aids that can improve recall and strengthen reconciliation efforts [11]. Patients can be queried as to whether they are taking drugs that are likely to have been omitted and/or may cause serious adverse drug events.

We organize the remainder of the paper into four sections. Section 2 presents our formulation of medication reconciliation as a CF problem. In Section 3, we describe three computational and statistical methods for CF. Section 4 evaluates these methods using medication data from a long-term care center in the eastern U.S. Finally, in Section 5 we provide a discussion of our results, limitations, and directions for future work.

Problem Formulation

At the center of the reconciliation problem is a patient's list of medications. This medication list is a set of entities, where each entity represents a drug. The most granular view of a drug entity is a *brand name drug with a certain dose and route* (e.g. Tylenol Oral Tablet 325 MG). This same entity can also be viewed in more general terms, as a *brand name drug* (e.g. Tylenol), as a *generic* chemical name (e.g. Acetaminophen), or more generally as a member of a *therapeutic class* (e.g. Non-Narcotic Analgesics). Our original medication data contained only the *branded drug with a certain dose and route*. Using the Center for Disease Control's Ambulatory Care Drug Database System (www2.cdc.gov/drugs) we classified each drug-dose-route entity into its respective branded drug, generic and therapeutic class

Regardless of the granularity of the drug entities, we can represent the complete and accurate medication list of all patients in a population as a matrix $M = \{m_{ij}\}$, for patients $i = 1, \cdots, I$ and drugs $j = 1, \cdots, J$, and where:

$$m_{ij} = \left\{ \begin{array}{l} 1 \; drug - entity \; j \; occurred \; in \; medication \; list \; of \; patient \; i \\ 0 \; otherwise \end{array} \right.$$

An analogous representation of the medication list is set of lists l_i, where l_i constitutes the set of drug entities e_j for a given patient i and $e_j \in l_i$ if and only if $m_{ij} = 1$ [12].

Because knowledge of a patient's true medication list l_i is often incomplete, prescribers observe a partial list of drugs for a patient, denoted by $\tilde{l}_i$. The observed partial list may be incomplete for several reasons. These include the failure to record a previous prescribing decision or the unintentional or intentional omission of a drug during patient self-report. The actual probability of omitting a given drug from a patient's list depends on a variety of factors. For instance, over-the-counter drugs or herbal supplements may have a higher probability of being missing than those prescribed by the current provider. On the other hand, omission of drugs may occur with no discernable pattern, that is to say each entity has an equal chance of being omitted [13].

Regardless of the distribution of omissions, discrepancies often result in a variety of negative outcomes for the patient. These include duplication of medication as well as prescribing drugs that negatively interact with those not recorded in the observed list.

Collaborative Filtering Methods

In most applications, the goal of collaborative filtering is to make predictions about products an individual may enjoy based on the aggregate tastes of similar individuals. In our case, we predict whether specific drugs have been omitted from an individual's medication list based on the known medications of similar individuals and the observed list of medications for that patient. Many computational and statistical methods for CF exist, each with its own advantages. Our preliminary experiments use three methods for ranking the drug-entities not observed in the partial list [12]. In each case, the algorithm assigns a score p_j for each drug not observed in the partial list. We then sort entities in decreasing order based on this score. We assume that the entity with the highest score is the one with the highest chance of being missing from the partial list, and so forth.

Popular

The "*popular*" algorithm considers each drug-entity e_j not observed in the partial list $\tilde{l}_i$, counts the number of lists l_{-i} in the training set which contain e_j, and chooses the most commonly occurring entities. The score p_j for each entity e_j is assigned according to the following equation, where $I(x)$ is the indicator function, returning 1 if x is true and 0 otherwise:

$$p_j = score(e_j) = \begin{cases} \sum_{l_{-i}} I(e_j \in l_{-i}) & e_j \notin \tilde{l}_i \\ 0 & e_j \in \tilde{l}_i \end{cases}$$

The *popular* algorithm can be expected to perform well if there are relatively few drugs that occur in many lists.

Co-occurrence counting

The "*co-occurrence counting*" algorithm scores each entity e_j not present in the observed partial list $\tilde{l}_i$ according to the number of times it has co-occurred with drug-entities $e_k \in \tilde{l}_i$ that are observed in the partial list. We calculate the score for each of the relevant entities in the following way, where e_k is the k^{th} drug entity in list $\tilde{l}_i$.

$$p_j = score(e_j) = \begin{cases} \sum_{e_k} \sum_{l_{-i}} I(e_k \in l_{-i}) \times I(e_j \in l_{-i}) & e_j \notin \tilde{l}_i \\ 0 & e_j \in \tilde{l}_i \end{cases}$$

The *co-occurrence counting* algorithm tends to do well when there is a strong pair wise structure in the prescribing patterns.

K-Nearest Neighbors

K-Nearest neighbors (KNN), a standard memory-based machine learning approach, operates in a straightforward manner [14]. Given an observed partial list, we find the K training lists $l_1 \cdots l_K$ that are closest to it according to some distance metric. Scores for the missing entities are assigned using majority vote of the K nearest neighbors. For our purpose, we use the Ochiai Similarity Measure, the binary form of cosine similarity to compare the observed partial list $\tilde{l}_i$ with each of the lists l_{-i} in the training set. We define a as the number of drug-entities that are present is both lists, b as the number of drug-entities present in $\tilde{l}_i$ but not in l_{-i} and c as the number of drug-entities present in l_{-i} but not in $\tilde{l}_i$. The Ochiai similarity measure is:

$$Dist(\tilde{l}_i, l_{-i}) = \sqrt{\left(\frac{a}{a+b}\right)\left(\frac{a}{a+c}\right)}$$

The nearest neighbors approach tends to do well when there are patients who are on similar drug regimens. Because our data is relatively sparse, we use a smoothed nearest neighbors approach, which is a weighted average of the base rates and the votes of the nearest neighbors. We specify the smoothed nearest neighbor vote as:

$$p_j = \frac{\alpha_j + v_j^K}{\alpha_j + \beta_j + K}$$

where $\alpha_j = r_j \times s$ and $\beta_j = s - \alpha_j$, are the relevant parameters for the base rates in the smoothed nearest neighbor method. The parameter s is the strength of the base rate information and r_j is the base rate for drug j. The smaller the value of s, the less emphasis placed on the base rates. The term v_j^K is the number of occurrences of drug j in the K nearest neighbors of list $\tilde{l}_i$. We use $K=3$ and $s=1$ in our evaluation discussed below.

Random

To establish a baseline for comparison of these methods, we use the *random* algorithm, which uses no information. This algorithm is the simplest possible approach, meant to provide a baseline for ordering the list. For each drug-entity not observed in the partial list, the algorithm assigns a score $p_j \in [0,1]$ uniformly at random.

Medication Data and Experiment

To evaluate the collaborative filtering methods on the medication reconciliation task, we used medication data from an online pharmacy that provides medications to a long-term care center in the Eastern United States. This data set contains 182 patients and a total of 177 unique branded drugs (excluding dose and route information) and 64 therapeutic classes. In our data, 49% of the patients are female and the median age is 81 with a 1st and 3rd quartiles of 63 and 88, respectively. The median number of medications

a patient has in our data is 14, with a 1^{st} and 3^{rd} quartiles of 7 and 21 medications, respectively.

For our analysis, we removed drugs that occurred no more than twice, since they would be difficult to predict based on data driven methods alone. The median number of drug occurrences is 9, and 1^{st} and 3^{rd} quartiles of 6 and 15 occurrences, respectively.

Cross-validation

We use a leave-one-out cross-validation approach to test how well each of the collaborative filtering methods described earlier perform in predicting omitted drug-entities. For each patient i in our data we randomly remove one drug from their list l_i to construct our observed list $\tilde{l}_i$. We then use information l_{-i} about all other patients, excluding patient i, as our training data to estimate our models for collaborative filtering and use these models to rank the drug list for patient i.

Results

Our experiments compared the four algorithms described in the previous section. For each algorithm, we attempted to predict the correct branded drug that was missing from a patient's record (without dose and route information). Table 1 summarizes these results. Columns 1, 10, 25, 50 and 100 give the proportion of patients whose missing drug is ranked at or below 1, 10, 25, 50, and 100 in the ordered list of drugs generated by each of the algorithms. An asterisk (*) next to the proportion indicates that the result was significantly better than the popular algorithm at $\alpha = .05$. The *median* and *mean* columns indicate median and mean rank of the omitted drugs in the ordered list generated by the algorithms.

Algorithm	1	10	25	50	100	Mean	Med.	Max
Random	.01	.07	.12	.28	.58	84.3	85	169
Popular	.13	.35	.53	.68	.84	44.3	19	174
KNN	.22*	.47*	.62	.76	.88	33.6	14	160
Co-Occ	.23*	.47*	.64*	.75	.90	33.7	12	162
*Indicates the algorithm performed significantly better than popular at $\alpha = .05$								

Table 1: Results for Missing Branded Drug

As expected, the baseline *random* method performed poorly. The *popular* algorithm, which used only the base rates, was able to guess the missing drug on the first try 13% of the time and required a median of 19 and a mean of 44.3 guesses to correctly guess the omitted drug. The modified *nearest neighbor* (with *k=3* and *s=1*) performed better than the popular algorithm, significantly improving the percentage correct on the first guess to 22% and improving the median and mean guesses to 14 and 33.6, respectively. The *co-occurrence* algorithm also

performed significantly better than the *popular* algorithm on the first guess as well, and was able to reduce the median and mean from 19 to 12 and 44.3 to 33.7, respectively.

We also used these algorithms to predict therapeutic class of the missing drug entity, by first predicting the drug and then choosing the corresponding therapeutic class. For instance, if the prediction was "Allegra", then we generalize the result to the therapeutic class "Antihistamines."

Algorithm	1	10	25	50	100	Mean	Med.	Max
Random	.03	.26	.51	.92	1.0	24.8	25	60
Popular	.14	.47	.77	.96	1.0	15.6	12	63
KNN	.23*	.53	.84	.98	1.0	12.9	9	61
Co-Occ	.24*	.59*	.82	.98	1.0	12.8	7.5	63
*Indicates the algorithm performed significantly better than popular at $\alpha = .05$								

Table 2: Results for Missing Therapeutic Class

Table 2 presents the results when we generalized the algorithms' predictions to the therapeutic class. Our results improved for two reasons: (1) the number of therapeutic classes was 64 vs. 177 brand names, and (2) one drug from a therapeutic class is often substituted for another. The therapeutic class is therefore easier to predict than the brand name. We were able guess the missing drug class on the first try 24% of the time with the Co-occurrence counting algorithm and the missing drug class was in the list of top 10 classes constructed by this algorithm for nearly 60% of the patients. The co-occurrence counting and nearest neighbor algorithms still performed considerably better than random guessing or using the popular algorithm.

Discussion and Conclusions

Our preliminary results suggest that simple collaborative filtering approaches that use only medication information can do a relatively good job at "recommending" missing medications. We expect that additional information about the patient, such as demographics and diagnoses, will further improve these results.

In practice, we can use the relevant top k drugs (e.g. top 10) as a decision aid for the "What other drugs are you taking?" question on intake forms. The ordered list can also be used during the prescribing process to account for potential drug interactions when an important drug may be missing from a patient's record.

Medical data is increasingly being structured, stored, and linked across organizational boundaries. This phenomenon will undoubtedly improve efforts to reduce medication errors. Nevertheless, better access to information alone cannot fully address the problem. There will always be occasions where important information is not

stored in any available database. Using a collaborative filtering approach, we are able to look beyond what is recorded using information from many other patients' records to predict omissions and improve the accuracy of each individual patient's medication list.

One limitation of our current research is that it assumes a fixed set of drugs, drawn only from the long-term care center itself. A second limitation is the implicit assumptions about the accuracy of the training data. In real life situations, the training data may also be spotty and imperfect. In future work, we plan to evaluate how robust the predictions are to relaxation of these assumptions.

As next steps, we will extend our collaborative filtering framework to include other information pertinent to predicting missing drugs, including information on patient demographics, diagnoses, prescribing physicians, and allergies.

We note that the current approach only assigns a score to the potentially missing drugs when ordering the list of candidates. This allows us to evaluate how well our algorithms work in identifying the missing information. In practice, we would need to take a decision theoretic approach, by looking at both the probability as well as the consequence of a missing drug.

Our current experiments suggest that the collaborative filtering approach to medication reconciliation holds promise. We anticipate improvements as additional information is used and we evaluate additional clinical settings. We do not envision a one-size-fits-all solution to this problem since clinical settings are heterogeneous and different collaborative filtering approaches may work better in different scenarios. We also hypothesize that a collaborative filtering approach may be beneficial in dealing with other types of discrepancies in medical data, such as laboratory tests, diagnoses and allergies.

Acknowledgements

References

1. Aspden, P., et al., Preventing Medication Errors. Quality Chasm Series, Washington, DC: National Academies Press. 2006.

2. Institute for Healthcare Innovation, Percent of Unreconciled Medications. 2006. http://www.ihi.org/IHI/Topics/PatientSafety/Medicat ionSystems/Measures/PercentofUnreconciledMedica tions.htm

3. Glintborg, B., S.E. Andersen, and K. Dalhoff, Insufficient communication about medication use at the interface between hospital and primary care. British Medical Journal, 2007. 16(1): p. 34-39.

4. Lindberg, M., Medication discrepancy: A concordance problem between dialysis patients and caregivers. Scandinavian Journal of Urology and Nephrology, 2007(1): p. 1-7.

5. Seaton, T.L., et al., Concordance Between Medication Histories and Outpatient Electronic Prescription Claims in Patients Hospitalized With Heart Failure. AMIA Annual Symposium Proceedings, 2005. 2005: p. 1109.

6. Midlöv, P., et al., Medication report reduces number of medication errors when elderly patients are discharged from hospital. Pharmacy World & Science, 2007: p. 1-7.

7. Santell, J.P., Reconciliation Failures Lead to Medication Errors. Joint Commission Journal on Quality and Patient Safety, 2006. 32(4): p. 225-229.

8. Nicol, N., Case study: An interdisciplinary approach to medication error reduction. Am J Health System Pharmacy, 2007. 64(14): p. 17.

9. Crichton, E.F., Patient recall of medication information. The Annals of Pharmacotherapy, 1978. 12(10): p. 591-599.

10. Breese, J.S., D. Heckerman, and C. Kadie, Empirical Analysis of Predictive Algorithms for Collaborative Filtering. Learning, 1992. 9: p. 309-347.

11. Mitchell, A.A., L.B. Cottler, and S. Shapiro, Effect of Questionnaire Design on Recall of Drug Exposure in Pregnancy. American Journal of Epidemiology, 2004. 123(4): p. 670-676.

12. Goldenberg, A., et al., A comparison of statistical and machine learning algorithms on the task of link completion. Proceedings of the ACM SIGKDD Workshop on Link Analysis for Detecting Complex Behavior, 2003.

13. Bedell, S.E., et al., Discrepancies in the Use of Medications Their Extent and Predictors in an Outpatient Practice. 2000, Am Med Assoc. p. 2129-2134.

14. Hastie, T., R. Tibshirani, and J. Friedman, The Elements of Statistical Learning: Data Mining, Inference, and Prediction. 2001: Springer.