Extracting Rx Information from Clinical Narrative

James G. Mork[1], MSc

jmork@mail.nih.gov

Olivier Bodenreider[1], M.D., PhD

obodenreider@mail.nih.gov

Dina Demner-Fushman[1], M.D., PhD

ddemner@mail.nih.gov

Rezarta Islamaj Doğan[2], PhD

islamaj@ncbi.nlm.nih.gov

François-Michel Lang[1], MSE

flang@mail.nih.gov

Zhiyong Lu[2], PhD

luzh@ncbi.nlm.nih.gov

Aurélie Névéol[2], PhD

neveola@ncbi.nlm.nih.gov

Lee Peters[1], MSc

lpeters@mail.nih.gov

Sonya E. Shooshan[1], MLS

sshooshan@mail.nih.gov

Alan R. Aronson[1], PhD

alan@nlm.nih.gov

[1]Lister Hill National Center for Biomedical Communications (LHNCBC),

[2]National Center for Biotechnology Information (NCBI),

U.S. National Library of Medicine, National Institutes of Health, Bethesda, MD 20894

Contact Information:
Alan R. Aronson, PhD
National Library of Medicine
Building 38A, Room 9N-905
8600 Rockville Pike, MSC-3826
Bethesda, MD 20894
Phone: 301.435.3162
Fax: 301.496.0673
alan@nlm.nih.gov

Abstract

OBJECTIVE: The i2b2 Medication Extraction Challenge provided an opportunity to evaluate our entity extraction methods, contribute to the generation of a publicly available collection of annotated clinical notes and start developing methods for ontology-based reasoning using structured information generated from the unstructured clinical narrative. DESIGN: We addressed the task of extracting salient features of medication orders from the text of de-identified hospital discharge summaries with a knowledge-based approach using simple rules and lookup lists. We combined our entity recognition tool, MetaMap, with dose, frequency and duration modules specifically developed for the Challenge as well as a prototype module for reason identification. MEASUREMENTS: Evaluation metrics and corresponding results were provided by the Challenge organizers. RESULTS: Our results indicate that robust rule-based tools achieve satisfactory results in extraction of simple elements of medication orders, but more sophisticated methods are needed for identification of reasons for the orders and durations. LIMITATIONS: Due to the time constraints and nature of the Challenge, some obvious follow-on analysis has not been completed yet. CONCLUSIONS: We plan to integrate the new modules with MetaMap to enhance its accuracy. This integration effort will provide guidance in retargeting our existing tools for better processing of clinical text.

I. INTRODUCTION

Extraction of the elements of medication orders from clinical narrative is a preliminary step in many important applications of medical informatics. These applications include but are not limited to: support of quality assurance through reconciliation of patient's medication lists and clinical notes [1, 2]; detection of adverse reactions to drugs [3] and medication non-compliance [4]; study of a population's response to a drug [5]; support of care plan development [6]; and identification of inactive medications [7].

Whereas evaluation of the individual efforts in extraction of medication names from biomedical literature could use "found data", such as Medical Subject Headings (MeSH[®]) assigned to MEDLINE[®] abstracts in the manual indexing process [8], until recently, no annotated resources for evaluation of extraction of medication orders from clinical narrative were publicly available. The opportunity to evaluate our named entity extraction methods and to contribute to development of an annotated publicly available large collection of clinical notes presented itself with the third i2b2 (Informatics for Integrating Biology and the Bedside) Medical Extraction Challenge [9].

To date, most algorithms and systems for extraction of drug order elements are knowledge-based. In fact, the absence of any large annotated collection makes it difficult to employ supervised machine learning. In contrast the availability of nomenclatures such as RxNorm [10] (which contains drug names, ingredients, strengths, and forms) encourages the use of rule-based systems. For example, Evans et al. developed a set of about 50 rules encoded as regular expressions to identify drug dosage objects and their attributes [11]. A Natural Language Processing (NLP) system augmented with the above rules and two lexicons (one containing drug names extracted from the Unified Medical Language System[®] (UMLS[®]) [12] and another one containing unusual words and abbreviations found in drug dosage phrases) identified about 80% of drug dosage expressions. Gold et al. expanded Evans' definition of drug dosage and implemented a system (the MERKI parser) that uses an RxNorm-based lexicon to extract known drug names and contextual clues to extract out-of-vocabulary drug names. Xu et al. developed an approach that attempts to extract a formal medication model (consisting of the drug name, signature modifiers and temporal modifiers) from clinical text using a chart parser and a

semantic grammar, and backs off to regular expressions if the chart parser fails [13].

The U.S. National Library of Medicine (NLM) tool (referred to as NLM's i2b2 Challenge Tool or simply, the Tool) developed to extract all fields originally defined in the i2b2 medication extraction guidelines is also knowledge-based and relies on lexical-semantic processing and pattern matching similar to the above systems. Our approach differs from the previously explored ones in that we 1) expanded a large number of term lists obtained for each element of drug phrases generating potential spelling variants and mining the UMLS for related terms as well as using corpus-based expansion, 2) developed a module for identification of negated drug mentions, 3) applied a UMLS-based approach to identification of reasons for medication orders, and 4) developed a module for validating drug and reason combinations.

II. METHODS

Early in the planning phase for this Challenge, the decision was made to use simple rules and lookup lists of various entities due to the time constraints of the Challenge. Our processing of the discharge summaries for this Challenge was relatively straightforward and is depicted in Figure 1. This section follows the course of our processing efforts beginning with a description of the lookup lists developed for the Challenge. Complete details can be found in our Appendix - Full-length Paper online (http://jamia.org).

[Figure 1 about here]

II.0. **Developed lookup lists required for the Challenge.**

The discovery of coverage gaps in our terminology resources (e.g., short forms of drug names such as *aspart* are not always covered in the UMLS, although the long form, *insulin aspart*, maps to two concepts) led to the decision to augment our initial resources with lookup lists. The lists that we developed used existing, publicly available resources with some minor manual curation based on processing the training set and reviewing what was missed by the Tool described here. Although many of the resources have items in common, each of the resources was added for specific reasons. Figure 2 graphically depicts the data sources with arrows connecting the entities and the lists where they made contributions.

[Figure 2 about here]

The drug identification list was created using DailyMed [14] for a list of common prescription drug names. We then added display names from RxTerms [15], Ingredients and Brand Names from RxNorm, and a list of drugs, drug classes, dosages, modes, frequencies, and durations from MERKI. In an attempt to complement the list of drugs we already had, we started looking at pharmacologic classes (e.g., diuretics), as opposed to drug names and added about 5,000 names from 1,360 UMLS concepts. RxHub [16], which is derived from drug names obtained from de-identified patient medication records, provided us with a list of common drug name misspellings. The U.S. Food and Drug Administration (FDA) Structured Product Labeling web site [17] provided us with extensive lists of Dosage Forms (dosages) and Routes of Administration (modes). Finally, manual curation was done to extend all of the lists based on reviews of the Tool results for the training set.

II.1 & 2. **The discharge summaries were read into the program and tokenized.**
Each line was tokenized using white-space as the token boundary. List boundaries were simply identified by which sections corresponded to the Challenge list of valid "list" sections. Sentence boundaries were identified using the simple rule of finding a "period" followed by spacing as long as the previous character wasn't a number. Sentence boundaries helped to define the extent of both drugs and reasons. Section identification was most crucial to this Challenge for several reasons: it 1) allowed us to decide if we wanted to process specific sections or ignore them, 2) assisted in limiting the scope of drugs and reasons, 3) was instrumental in determining whether a drug was in a "list" or "narrative", and 4) helped eliminate some ambiguity (e.g., not identifying drugs within Allergy sections). Candidate section names were defined as all strings occurring at the beginning of a line, consisting of uppercase letters only (a mixed case review was attempted, but found to be too noisy), and followed by a period, a colon, or the end of the line. We identified 10,454 such potential section names, 937 of them unique. The list of unique names was then manually reviewed, scrubbed, and some mixed case section names were manually added to the list e.g., "Attending". We consequently created a list of twenty-one triggers (see Table 1) that denoted sections we could ignore. We ended with 632 section names extracted

from the training set.

lab	laboratory	laboratories	allergies	allergy
attending	fam hx	family history	family hsitory	discharge date
service	labs	escription document	dictated by	entered by
vital sign	vitals	signs	vital signs	diet

Table 1: List of Trigger Phrases for Sections to be Ignored

II.3. **Text was reformatted into a single text line.**

Early testing showed that by simply processing the summaries line by line, we ended up missing some drugs and reasons because the text was broken across lines. So, once the sections were identified, we combined all of the text to be processed into a single line. A mapping between the reformatted and original text was maintained.

II.4. **Reasons were identified using MetaMap and exact matches from the Gopher list.**

We used both MetaMap [18] and a list derived from the Gopher [19] project to identify reasons. In this Challenge, the discharge summaries sometimes had misspellings, acronyms/abbreviations, and different ways of stating a medical reason for prescribing a drug. While MetaMap was able to identify most of the spelling variations and any text inversions, it was limited to the contents of the UMLS Metathesaurus. The Gopher lookup list was introduced to expand our coverage and to assist with less well-behaved occurrences. In the end, the two approaches seemed to complement each other fairly well. We also maintained a "bad reason" list to eliminate as many false positives as possible (see section *II.6*).

II.5. **Reasons were then reconciled with the original text and tagged using the mapping information from the single free-text line back to the original discharge summary.**

We used exact text matches to the lookup lists to tag drugs, modes, dosages, durations, and frequencies. Drug boundaries were also identified by noting the first position of each drug so we could know when we came to the end of the current drug during filtering. Drug boundaries expanded left and right depending on where the components were identified with the final drug boundary encompassing the drug name and any of its associated components.

II.6. **Filtering was performed to add, remove, and extend tagged items.**

Filtering involved simple rules, a "bad reason" trigger list (e.g., "*ruled out for*"), and a "bad drugs" list for what should be removed (e.g., *insulin* within *insulin-dependent diabetes*). We developed rules for limiting the scope of a drug to try and eliminate the crossover of components, and we also tried to identify non-active medications (e.g., *should not take aspirin*) and allergy-specific drugs to remove false positives. Simple rules for expanding components by looking at the tokens to the left and right of the component were developed as needed.

II.7. **Drug/reason pairings identified.**

Once drugs and reasons had been initially identified, we attempted to match each drug name with a nearby reason. Initially we had a very simple rule to use the closest reason if there were two possibilities. This was refined to ensure that reason assignment did not violate a drug, list, or section boundary. We also created a small set of trigger phrases to use in combining certain nearby reasons and drugs (see Figure 3). In some cases, we allowed multiple reasons for a drug if they were next to each other and connected with a comma, "*and*", or "*or*".

[Figure 3 about here]

II.8. **Drug/reason pairings validation.**

Once drug/reason pairings were identified, we attempted to validate the pairings via knowledge contained in the UMLS. The validation of the drug/reason pairings was accomplished via a constrained traversal of the UMLS relations involving two main steps as described below.

Drugs and reasons were first mapped to UMLS concepts, using exact and normalized matches, and further restricting mappings to the semantic group *Chemicals & Drugs* and *Disorders*, respectively. All successful mappings were considered, including several pairs of UMLS concepts generated by one original drug/reason pairing.

Selected UMLS relations were then used to identify plausible relations between drugs and reasons. The key relations were provided by the NDF-RT source vocabulary where ingredients are associated with diseases through *may_treat* and *may_prevent* relationships.

The algorithm did not explore all paths, but rather stopped at the first path reached between the drug and the reason. For example, "*albuterol / asthma*" was identified through a direct link between ingredient and disease. 9,415 possible drug/reason pairings were found with 2,785 of these having at least one path through the UMLS tying them together.

III. RESULTS

We finished fourth overall out of 20 teams that participated in the Challenge. Since two of the three teams who scored best had pre-existing systems that were modified for the Challenge, we were pleased that a system developed expressly for the Challenge performed so well. The lessons learned during this effort are being evaluated for inclusion in our NLP tool suite. Results are shown in Table 3 and Table 7 in the i2b2 JAMIA overview paper [20]. It is clear from Table 7 that all teams had significant problems with identifying both Durations and Reasons.

IV. DISCUSSION

In general we are satisfied with our vocabulary and rule-based identification of drug names, doses, modes and frequencies. The lack of significant difference between our exact and inexact scores confirms this view since it shows that we either found the entire element or missed it completely. Our dose and duration results are satisfactory, considering they are based on very simple heuristics. However, the approach is brittle in the presence of pattern changes in the middle of an enumeration of drugs. Deeper understanding of the context is needed to overcome this weakness.

Low scores for durations and reasons, on the other hand, show that our methods are clearly insufficient for those drug elements. In the absence of creating a full-fledged natural language understanding system, some improvement might be achieved using corpus-based methods. Any corpus-based methods would need to be judiciously applied given their known weaknesses: they are noisy if not supervised, and they are ambiguous even when supervised. For example, using our corpus-based expansion we identified *HCT* as an abbreviation of *hydrochlorothiazide* (more commonly abbreviated as *HCTZ*); however, *HCT* is also common shorthand for *hematocrit*.

Finally, we intend to incorporate some of our Tool's features into the MetaMap algorithm. Specifically, we will include the overall identification of drug mentions with the expectation that it will reduce ambiguity because of the coordination of a drug's elements. In addition, augmenting MetaMap's negation algorithm with the drug-specific negation detection developed for the Challenge should be useful in applying it to clinical text.

V. LIMITATIONS

Many of the limitations of this research occurred because we are reporting on the development of an NLP application in the context of a time-sensitive Challenge rather than fundamental research. In-depth analysis that we would normally have done will be done in the future. Examples of such analysis include determining the relative contributions to our results from the many knowledge sources we used, a similar analysis of the contributions of the filtering rules, and a study to determine an optimal balance between the knowledge sources and the rules. In addition, the relations identified between drugs and diseases from selected UMLS relations are not intended to be used as a reference set of relations reflecting therapeutic intent. Rather, we use constraints on the UMLS graph of relations in order to identify plausible drug-reason relations for the purpose of validating drug-reason pairings. Despite the presence of many false positives and false negatives, our algorithm proved useful in the context of this Challenge.

Acknowledgements

References

1. Gold S, Elhadad N, Zhu X, Cimino JJ, Hripcsak G. Extracting structured medication event information from discharge summaries. *AMIA Annu Symp Proc.* 2008 Nov 6:237-41.

2. Cimino JJ, Bright TJ, Li J. Medication reconciliation using natural language processing and controlled terminologies. *Stud Health Technol Inform.* 2007;129(Pt 1):679-83.

3. Carol Friedman. Discovering Novel Adverse Drug Events Using Natural Language Processing and Mining of the Electronic Health Record. *Artificial Intelligence in Medicine, 12th Conference on Artificial Intelligence in Medicine, AIME 2009, Verona, Italy, July 18-22, 2009*. Proceedings 2009.

4. Turchin A, Wheeler HI, Labreche M, Chu JT, Pendergrass ML, Einbinder JS. Identification of documented medication non-adherence in physician notes. *AMIA Annu Symp Proc.* 2008 Nov 6:732-6.

5. Sirohi E, Peissig P. Study of Effect of Drug Lexicons on Medication Extraction from Electronic Medical Records. *Pac Symp Biocomput.* 2005:308-18.

6. Demner-Fushman D, Seckman C, Fisher C, Hauser SE, Clayton J, Thoma GR. A Prototype System to Support Evidence-based Practice. *AMIA Annu Symp Proc.* 2008 Nov 6:151-5.

7. Breydo EM, Chu JT, Turchin A. Identification of inactive medications in narrative medical text. *AMIA Annu Symp Proc.* 2008 Nov 6:66-70.

8. Aronson AR, Mork JG, Névéol A, Shooshan SE, Demner-Fushman D. Methodology for Creating UMLS Content Views Appropriate for Biomedical Natural Language Processing. *AMIA Annu Symp Proc.* 2008 Nov 6:21-5.

9. Informatics for Integrating Biology and the Bedside, i2b2, a National Center for Biomedical Computing, Third Shared-Task Challenge in Natural Language Processing for Clinical Data Medication Extraction Challenge: https://www.i2b2.org/NLP/Medication.

10. RxNorm: http://www.nlm.nih.gov/research/umls/rxnorm.

11. Evans DA, Brownlow ND, Hersh WR, Campbell EM. Automating concept identification in the electronic medical record: an experiment in extracting dosage information. *Proc AMIA Annu Fall Symp.* 1996:388–92.

12. UMLS Knowledge Sources: http://www.nlm.nih.gov/research/umls.

13. Xu H, Stenner SP, Doan S, Johnson KB, Waitman LR, Denny JC. MedEx: a medication information extraction system for clinical narratives. *J Am Med Inform Assoc., 2010 Jan–Feb;17(1):19-24.*

14. DailyMed: http://dailymed.nlm.nih.gov.

15. RxTerm: http://wwwcf.nlm.nih.gov/umlslicense/rxtermApp/rxTerm.cfm.

16. Fung KW, Applied Medical Terminology Research. A Report to the Board of Scientific Counselors. The RxHub Project – a sneak preview. April 2009. Page 32. http://www.lhncbc.nlm.nih.gov/lhc/docs/reports/2009/tr2009001.pdf#page=33.

17. FDA Structured Product Labeling Resources: http://www.fda.gov/ForIndustry/DataStandards/StructuredProductLabeling

18. Aronson AR, Lang FM. The Evolution of MetaMap, a Concept Search Program for Biomedical Text. *AMIA Annu Symp Proc.* 2009 Nov:22.

19. McDonald CJ, Tierney WM: The medical gopher-A microcomputer system to help find, organize and decide about patient data. *West J Med* 1986 Dec; 145(6):823-9.

20. Ozlem Uzuner, Imre Solti, Eithon Cadag. Extracting Medication Information from Clinical Text. *J Am Med Inform Assoc., 2010 (to appear).*

Figure 1 Processing Flow Diagram

Figure 2: Lookup Lists and Their Sources

Figure 3: Simple Reason Grouping Rules