Validation of a Crowdsourcing Methodology for Developing a Knowledge Base of Related Problem-Medication Pairs

A. B. McCoy^{1,2}; A. Wright^{3,4,5}; M. Krousel-Wood^{1,6,7}; E. J. Thomas^{8,9}; J. A. McCoy¹⁰; D. F. Sittig^{9,11}

¹Center for Applied Health Services Research, Ochsner Health System; ²Department of Biostatistics and Bioinformatics, Tulane University School of Public Health and Tropical Medicine, New Orleans, LA; ³Department of Medicine, Brigham and Women's Hospital; ⁴Department of Clinical Informatics Research and Development, Partners HealthCare; ⁵Harvard Medical School, Boston, MA; ⁶Department of Medicine, Tulane University School of Medicine; ⁷Department of Epidemiology, Tulane University School of Public Health and Tropical Medicine, New Orleans, LA; ⁸Department of Internal Medicine, The University of Texas Medical School at Houston; ⁹The University of Texas at Houston-Memorial Hermann Center for Healthcare Quality and Safety, Houston, TX; ¹⁰Department of Urology, Ochsner Health System, New Orleans, LA; ¹¹The University of Texas School of Biomedical Informatics at Houston, Houston, TX

Keywords

Crowdsourcing, electronic health records, knowledge bases, problem-oriented medical records, computer-assisted drug therapy, validation studies

Summary

Background: Clinical knowledge bases of problem-medication pairs are necessary for many informatics solutions that improve patient safety, such as clinical summarization. However, developing these knowledge bases can be challenging.

Objective: We sought to validate a previously developed crowdsourcing approach for generating a knowledge base of problem-medication pairs in a large, non-university health care system with a widely used, commercially available electronic health record.

Methods: We first retrieved medications and problems entered in the electronic health record by clinicians during routine care during a six month study period. Following the previously published approach, we calculated the link frequency and link ratio for each pair then identified a threshold cutoff for estimated problem-medication pair appropriateness through clinician review; problem-medication pairs meeting the threshold were included in the resulting knowledge base. We selected 50 medications and their gold standard indications to compare the resulting knowledge base to the pilot knowledge base developed previously and determine its recall and precision.

Results: The resulting knowledge base contained 26,912 pairs, had a recall of 62.3% and a precision of 87.5%, and outperformed the pilot knowledge base containing 11,167 pairs from the previous study, which had a recall of 46.9% and a precision of 83.3%.

Conclusions: We validated the crowdsourcing approach for generating a knowledge base of problem-medication pairs in a large non-university health care system with a widely used, commercially available electronic health record, indicating that the approach may be generalizable across health-care settings and clinical systems. Further research is necessary to better evaluate the knowledge, to compare crowdsourcing with other approaches, and to evaluate if incorporating the knowledge into electronic health records improves patient outcomes.

Correspondence to:

Tulane University School of Public Health and Tropical
Medicine

Department of Biostatistics and Bioinformatics

1440 Canal Street, Suite 2001

New Orleans, LA 70112

Email: amccoy1@tulane.edu

Appl Clin Inform 2015; 6: 334–344

<http://dx.doi.org/10.4338/ACI-2015-01-RA-0010>

received: January 21, 2015

accepted: April 5, 2015

published: May 20, 2015

Citation: McCoy AB, Wright A, Krousel-Wood M, Thomas EJ, McCoy JA, Sittig DF. Validation of a crowdsourcing methodology for developing a knowledge base of related problem-medication pairs. *Appl Clin Inf* 2015; 6: 334–344

<http://dx.doi.org/10.4338/ACI-2015-01-RA-0010>

1. Introduction

Clinical knowledge bases of problem-medication pairs are necessary for many electronic health record (EHR) based solutions that improve patient safety, such as clinical summarization [1]. One approach to developing these knowledge bases is crowdsourcing, which takes advantage of manually asserted links between medications and problems by clinicians during e-prescribing [2]. Our prior work has demonstrated that the crowdsourcing approach generated an accurate, up-to-date problem-medication knowledge base, but the study was limited in its generalizability, as it was completed at a university-based, academic ambulatory practice with one EHR and only validated internally. This paper reports the validation of the crowdsourcing methodology externally at a large, not-for-profit community-based academic medical center with a different, widely used, commercially available EHR.

2. Background

2.2. Electronic Health Record Summarization

EHRs have great potential to improve patient safety, provider efficiency, and costs across healthcare settings [3]. Given the potential benefits, implementation and “meaningful use” of EHRs is required by the Health Information Technology and Clinical Health (HITECH) Act, part of the American Recovery and Reinvestment Act of 2009 [4]. In most EHRs, patient data elements are organized by source or content type, such as medications, laboratory results, problems, allergies, notes, visits, and health maintenance items [5, 6]. However, some research has found that more optimal presentations of this information, such as displays organizing data by clinical conditions, may prevent medical errors and increase quality [7–10]. In particular, as EHRs are containing increasing amounts of data, these simplified, problem-oriented displays to summarize patient data may reduce frustration and inefficiency among care providers [1, 11–15]. In order for EHRs to summarize clinical data, the systems must contain knowledge bases about the relationships between the data types [1]. For example, a relationship between a medication and a problem is the “treats” relationship. Until recently, systems have not contained such knowledge bases, as they are difficult to develop, encode, and maintain.

2.2. Knowledge Base Development

Many approaches exist for developing knowledge bases of related problem-medication pairs, including standards-based ontologies, data mining, and literature mining, each having advantages and disadvantages. The use of standards-based ontologies may require little effort by investigators if the local data are already mapped to the standardized terminology, but these ontologies do not exist for all data types, they require considerable maintenance by the owners, they are often incomplete, and mapping to the terminology may be difficult [16,17]. Data mining does not require mapping to a standardized terminology, and it is one approach that can simply be repeated in multiple settings and may develop a knowledge base reflecting local practice trends [18–20]. However, it, along with literature mining [21–23], may result in a knowledge base that is biased toward more common links and will often miss more rare but often important links. Ensemble approaches can help overcome some of these limitations, but mapping to a common terminology is still required [24, 25].

2.3. Crowdsourcing

Crowdsourcing is the act of outsourcing a task to a group or community of people [26, 27]. Crowdsourcing has been used to create popular resources such as Wikipedia [28], a free Internet encyclopedia, and it has been used in various biomedical approaches, including drug discovery resources [29–31], document tagging [32], clinical decision support evaluation [33], and bioinformatics tasks [34].

Prior work utilized crowdsourcing to generate a pilot problem-medication knowledge base in a university-based, academic ambulatory practice [2]. The approach extracted links between clinical problems in a patient's problem list and medications that were entered by clinicians during e-prescribing, a task required in an increasing number of health centers. In this crowdsourcing scenario, clinician EHR users represented the community, and generating problem-medication pairs represented the outsourced task [35]. To evaluate the appropriateness of the entered links, investigators calculated two metrics for each problem-medication pair: the link frequency and the link ratio. The link frequency is defined as the number of patients for whom a clinician asserted a link between the medication and problem. For example, if the medication simvastatin was linked to hypercholesterolemia for 50 distinct patients, the link frequency for that problem-medication pair would be 50. The link ratio is defined as the number of patients for whom a link between a medication and problem has been manually asserted divided by the number of patients having been prescribed that medication and with that problem on their problem list, regardless of whether the two were linked. If 100 patients have simvastatin on their medication list and hypercholesterolemia on their problem list, but only 50 patients have a link between simvastatin and hypercholesterolemia, the link ratio for the pair is 0.5. To replicate the approach, an investigator would extract the linked problem-medication pairs, compute the link frequency and link ratio, and determine a threshold for inclusion in the final knowledge base.

Our prior work has demonstrated that crowdsourcing is an effective, inexpensive method for generating an accurate, up-to-date problem-medication knowledge base, which healthcare information systems can employ to generate problem-oriented summaries or infer missing problem list items to improve patient safety. However, the study was limited in its generalizability, as it was completed at a single study site with one EHR and only validated internally. This paper reports the validation of the crowdsourcing methodology externally at an additional clinical site with a different commercially available EHR.

3. Methods

We conducted this validation study at a large, not-for-profit community-based academic medical center that includes eight hospitals and over 38 outpatient health centers in urban and rural settings throughout southeast Louisiana. We first retrieved from the EHR reporting database all medications and problems entered in structured form in the EHR by all clinicians throughout the medical center during routine care (EpicCare 2010, Madison, WI) during a six-month time frame (July 1, 2013 through December 31, 2013). The EHR uses First DataBank (FDB) as the underlying terminology to populate the medication dictionary and ICD-9 to populate the problem dictionary. We followed the same approach previously described in the pilot work by calculating the link frequency and link ratio for each pair, dividing the pairs into 25 groups by range of link frequency and link ratio, and identifying a threshold cutoff for estimated problem-medication pair appropriateness after clinician review (i.e., the minimum acceptable estimated appropriateness) [2]. A clinician (JAM) reviewed 100 problem-medication pairs from each of the 25 threshold groups (2500 total pairs) to determine whether the pair was appropriate (i.e., the medication is used in the treatment of the linked problem). Problem-medication pairs in groups meeting the threshold cutoff were included in the resulting knowledge base.

To evaluate and compare the resulting knowledge base to the pilot knowledge base developed previously, we used the Lexi-Comp drug database (Wolters Kluwer, Hudson, Ohio, USA), a free-text resource, curated by pharmacists and pharmacology experts, as the gold standard for appropriate medication indications. We selected a sample of 50 medications that had been prescribed in at least one of the pilot or validation study settings; we initially retrieved 50 randomly selected medications from the union set, and when there were overlapping medications or classes, we discarded the duplicate and replaced it with another randomly selected medication. We then retrieved the corresponding labeled and unlabeled uses (i.e., problems) from Lexi-Comp. We manually reviewed the resulting knowledge bases from both the current study and the pilot crowdsourcing study to determine the recall (i.e., the proportion of gold standard pairs that are in the resulting knowledge base, or sensitivity) and precision (i.e., the proportion of pairs in the resulting knowledge base that are in the

gold standard set, or positive predictive value) with 95% confidence intervals (CIs), relevance-based measures traditionally used to evaluate information retrieval approaches, of each [36]. The study was approved by the Ochsner Health System Institutional Review Board.

4. Results

During the study period, clinicians ordered 4,611,939 medications, entered 4,864,892 problems, and linked 733,629 medications (15.9%) to problems (166,378 distinct problem-medication pairs). These pairs included 2,137 distinct medications (56.11% of ordered medications) and 22,636 distinct problems (35.89% of entered problems).

Evaluation of the appropriateness of the problem-medication pairs found results similar to those in the pilot evaluation. ► Table 1 depicts the percent of appropriate pairs as determined by clinician review and the total number of pairs in each group, and ► Figure 1 includes heat maps indicating the appropriateness of the reviewed pairs for both the pilot and the validation studies. For this study, we selected 92% as the threshold cutoff so that we did not include groups having a lower link ratio and link frequency range than any group that was excluded, as the estimated appropriateness did not linearly correlate with link frequency and link ratio in each group. For example, the estimated appropriateness for the threshold group with a link frequency of 3–4 and a link ratio of 0.2–0.9 was 97%, and the groups with the same link frequency and link ratios of 0.3–0.39 and ≥ 0.5 were 92% and 94%, respectively. Including only groups with an estimated appropriateness greater than or equal to 95%, as was done in the pilot study, would include the first group but not the second two groups, even though the link ratio was greater in the second two groups, but the threshold cutoff of 92% included all of these groups. One group (link frequency ≥ 10 and link ratio < 0.1) met the criteria in the pilot study but not in the current study. Our resulting knowledge base of pairs having an estimated appropriateness greater than or equal to 92% contained 26,912 problem-medication pairs; the pilot crowdsourcing knowledge base contained 11,167 problem-medication pairs.

The 50 randomly selected medications and corresponding uses from Lexi-Comp represented 194 problem-medication pairs. ► Table 2 lists the generic forms of the 50 included medications. Of the 194 gold standard pairs, 91 were included in the pilot crowdsourcing knowledge base, and 121 were included in the current study knowledge base, resulting in recalls of 46.9% (95% CI: 39.8%, 54.2%) and 62.3% (95% CI: 55.1%, 69.1%), respectively. The pilot crowdsourcing knowledge base contained 91 pairs not included in the gold standard (83.3% precision, 95% CI: 79.8%, 86.3%), and the current study knowledge base contained 294 pairs (87.5% precision, 95% CI: 86.1%, 88.8%).

5. Discussion

The crowdsourcing approach, when applied to a community-based, academic health system using a widely used, commercially-available EHR, had high recall and preprecision and provided external validation of a previously developed approach to developing a knowledge base of related problem-medication pairs. Our findings indicate that the crowdsourcing approach is generalizable to different settings, and, in fact, can be more accurate than previously described. These findings are important, as knowledge bases of problem-medication pairs can be used in a variety of different tasks to automate EHR use and improve clinical care; such tasks include patient summarization to overcome information overload, suggestion of indicated problems during e-prescribing to speed up order entry, and identification of undocumented patient problems to increase the accuracy and completeness of patient records [18]. Further, the crowdsourcing approach may be more efficient and accurate than existing approaches for generating such knowledge.

Although different methods may be used to identify a threshold for appropriateness to include linked problem-medication pairs in the resulting knowledge base, we repeated the previously described methods, including keeping the same threshold groups for link frequency and link ratio so that we could directly compare the results. The estimated appropriateness of each threshold group was similar between the pilot study and the current validation study, although we ultimately selected a different cutoff for inclusion in the final knowledge base. The pilot study included all links having

greater than or equal to 10 links regardless of link ratio, but in the current study we excluded those having a link ratio less than 0.1. These differences likely reflect variations in prescribing practices and patient populations across the two sites. At the pilot study site clinicians are required to link a problem to a prescribed medication, and at the current study site clinicians are not, suggesting that clinicians may be more likely to link an incorrect problem during e-prescribing to proceed in the workflow.

The differences in the appropriateness of linked pairs also highlights the need for investigators utilizing the crowdsourcing approach to execute methods for determining appropriateness for inclusion in the resulting knowledge base as necessary to meet the needs for their unique situation. For example, a knowledge base used for clinical summarization would benefit from having a higher sensitivity to make sure that all elements are included in the summary, so a lower threshold for inclusion could be selected; a knowledge based used to prompt clinicians about potentially undocumented problems would benefit from having a higher specificity to prevent delivery of false advice, so a higher threshold for inclusion would be selected.

After completing the approach in the validation setting, we found that the resulting knowledge base had a higher recall and precision than those of the pilot study, despite the lower threshold cutoff that we selected. One likely cause of this finding is the larger data set on which we completed the methods. Another potential cause is the differing levels of granularity in the underlying EHR data; in the pilot study, investigators utilized medications in prescribed form, including dosage and some brand names, while in the current study we utilized only generic forms of medications. Greater detail in the pilot study likely decreased the frequencies of given links compared to grouped medications (i.e., by generic form) in the current study. For example, previously, “Ciprofloxacin HCL 250 MG Oral Tablet”, “Ciprofloxacin HCL 500 MG Oral Tablet”, “Cipro 250MG/ML (5%) Oral Suspension Reconstituted”, and “Cipro 500MG/ML (10%) Oral Suspension Reconstituted” all resulted in distinct links to clinical problems, often to the same problem, such as a urinary tract infection. In the current study, these medications were all condensed to “Ciprofloxacin HCL” and fewer distinct problem-medication pairs but with a greater link frequency. As a result, we recommend using the most generic forms of medications and least specific problems (e.g., “Otitis Externa” vs. “Otitis Externa of the Right Ear”) while maintaining clinical relevance to complete the crowdsourcing approach. Finally, it is possible that differences in the documentation practices by clinicians at the study sites contributed to the higher recall and precision in the current study. These differences could be due to an easier-to-use EHR that facilitates linking, better training of clinicians at documentation, or more attentive clinicians in performing the task.

5.1. Limitations

This work has some limitations. First, the different granularity within the data sets (i.e., the use of medication dose and form in the initial study compared to the use of generic medications in the current study) may have accounted for differences in the accuracy of the resulting knowledge bases. Still, application of the approach with both data types resulted in accurate knowledge bases, thus increasing the generalizability of the approach. Our study is also limited in that the pairs were reviewed by a single clinician for appropriateness. However, similar reviews by the clinician have been validated in prior work, so we expect that minimal bias was introduced. Another limitation is that, because of the differing underlying terminologies used in the two studies (i.e., Medispan and FDB for medications, MEDCIN and ICD-9 for problems), we were not able to directly compare the resulting knowledge bases and review overlapping and non-overlapping pairs across the entire sets. However, our manual evaluation compared to the gold standard provided adequate assurance that the approach was still successful.

5.2 Future Work

In the future, we plan to improve the accuracy of our developed knowledge bases by adopting an ensemble approach and optimizing the advantages of the crowdsourcing and other approaches to generating this knowledge. Work is currently underway to adequately map the terminologies and determine the optimal approach to combining each source [24]. We also plan to repeat the methodology

at multiple sites, to determine whether site features (e.g., EHR vendor, terminology, culture), affect the resulting knowledge base. Finally, we plan to incorporate our knowledge bases into EHRs for summarization and clinical decision support, and we will assess whether there is any improvement on both patient and provider outcomes.

6. Conclusion

We validated a previously developed crowdsourcing approach for generating a knowledge base of problem-medication pairs in a large integrated healthcare network using a different commercially available EHR. The validation resulted in a knowledge base of problem-medication pairs that had a similar recall and precision compared to the previous study, supporting the generalizability of the crowdsourcing approach across EHRs and clinical settings. Further research is necessary to better evaluate the knowledge with different EHRs and clinical settings and to incorporate other approaches, in addition to implementing the knowledge into EHRs to improve outcomes.

Clinical Relevance Statement

Knowledge bases of problem-medication pairs can be used in a variety of different tasks to automate EHR use and improve clinical care, including patient summarization, suggestion of indicated problems during e-prescribing, and identification of undocumented patient problems. This study validates the previously described crowdsourcing approach to developing such knowledge bases.

Conflict of Interest

The authors declare that they have no conflicts of interest in the research.

Acknowledgments

This work was supported in part by NLM grant 1 K22 LM011430–01A1 and by the Ochsner Health System Center for Applied Health Services Research.

Fig. 1 Heat maps indicating the appropriateness of linked problem-medication pairs by link frequency and link ratio

		Link Frequency				
		1	2	3–4	5–9	≥ 10
Link Ratio	< 0.1	69% 43,585	72% 9,885	72% 6,368	81% 3,881	86% 2,455
	0.1 – 0.19	81% 14,304	91% 3,159	94% 2,371	96% 1,791	95% 1,954
	0.2 – 0.29	77% 10,737	94% 2,367	97% 1,408	95% 1,127	99% 1,560
	0.3 – 0.39	82% 8,171	95% 1,666	92% 1,532	97% 1,211	99% 1,807
	≥ 0.5	80% 36,921	98% 3,582	94% 1,928	95% 1,163	98% 1,445

Table 1

Appropriateness of linked problem-medication pairs by link frequency and link ratio.

Medications
Adalimumab, Alfuzosin, Amitriptyline, Aprepitant, Azithromycin, Brimonidine, Budesonide, Capsaicin, Chlorpheniramine/Pseudoephedrine, Chlorthalidone, Ciprofloxacin, Clotrimazole, Darunavir, Desonide, Dexamethasone, Digoxin, Docusate, Eptifibatide, Estazolam, Ethinyl estradiol/Norelgestromin, Famotidine, Fluocinonide, Griseofulvin, Hydrocodone/pseudoephedrine, Hydrocortisone, Indomethacin, Ketorolac, Lansoprazole, Leucovorin, Mesalamine, Moxifloxacin, Naproxen, Nedocromil, Nicardipine, Olmesartan/Hydrochlorothiazide, Orphenadrine, Oxycodone/Acetaminophen, Paricalcitol, Phenobarbital, Plerixafor, Polysaccharide-iron complex, Poly-ureaurethane, Rifampin, Salsalate, Selenium sulfide, Sitagliptin/Simvastatin, Sumatriptan, Travoprost, Valsartan, Vasopressin

Table 2

Randomly selected medications for gold standard evaluation.

References

1. Feblowitz JC, Wright A, Singh H, Samal L, Sittig DF. Summarization of clinical information: a conceptual model. *J Biomed Inform* 2011; 44(4): 688–699.
2. McCoy AB, Wright A, Laxmisan A, Ottosen MJ, McCoy JA, Bitten D, Sittig DF. Development and evaluation of a crowdsourcing methodology for knowledge base construction: identifying relationships between clinical problems and medications. *J Am Med Inform Assoc JAMIA* 2012; 19(5): 713–718.
3. Bates DW, Cullen DJ, Laird N, Petersen LA, Small SD, Servi D, Laffel G, Sweitzer BJ, Shea BF, Hallisey R. Incidence of adverse drug events and potential adverse drug events. Implications for prevention. ADE Prevention Study Group. *JAMA J Am Med Assoc* 1995; 274(1): 29–34.
4. Blumenthal D, Tavenner M. The “Meaningful Use” Regulation for Electronic Health Records. *N Engl J Med* 2010; 363(6): 501–504.
5. Laxmisan A, McCoy AB, Wright A, Sittig DF. Clinical Summarization Capabilities of Commercially-available and Internally-developed Electronic Health Records. *Appl Clin Inform* 2012; 3(1): 80–93.
6. Health IT and Patient Safety: Building Safer Systems for Better Care – Institute of Medicine. Available from: <http://iom.edu/Reports/2011/Health-IT-and-Patient-Safety-Building-Safer-Systems-for-Better-Care.aspx>
7. Han YY, Carcillo JA, Venkataraman ST, Clark RSB, Watson RS, Nguyen TC, Bayir H, Orr RA. Unexpected increased mortality after implementation of a commercially sold computerized physician order entry system. *Pediatrics* 2005; 116(6): 1506–1512.
8. Horsky J, Kuperman GJ, Patel VL. Comprehensive analysis of a medication dosing error related to CPOE. *J Am Med Inform Assoc JAMIA* 2005; 12(4): 377–382.
9. Koppel R, Metlay JP, Cohen A, Abaluck B, Localio AR, Kimmel SE, Strom BL. Role of computerized physician order entry systems in facilitating medication errors. *JAMA J Am Med Assoc* 2005; 293(10): 1197–1203.
10. McCoy AB, Waitman LR, Lewis JB, Wright JA, Choma DP, Miller RA, Peterson JF. A framework for evaluating the appropriateness of clinical decision support alerts and responses. *J Am Med Inform Assoc JAMIA* 2012; 19(3): 346–352.
11. Ash JS, Berg M, Coiera E. Some Unintended Consequences of Information Technology in Health Care: The Nature of Patient Care Information System-related Errors. *J Am Med Inform Assoc* 2004; 11(2): 104–112.
12. Ash JS, Sittig DF, Poon EG, Guappone K, Campbell E, Dykstra RH. The extent and importance of unintended consequences related to computerized provider order entry. *J Am Med Inform Assoc JAMIA* 2007; 14(4): 415–423.
13. Sittig DF, Singh H. Legal, ethical, and financial dilemmas in electronic health record adoption and use. *Pediatrics* 2011; 127(4): e1042–e1047.
14. Sittig DF, Teich JM, Osheroff JA, Singh H. Improving clinical quality indicators through electronic health records: it takes more than just a reminder. *Pediatrics* 2009; 124(1): 375–377.
15. Sittig DF, Singh H. Rights and responsibilities of users of electronic health records. *CMAJ Can Med Assoc J J Assoc Medicales Can.* 2012; 184(13): 1479–1483.
16. Carter JS, Brown SH, Erlbaum MS, Gregg W, Elkin PL, Speroff T, Tuttle MS. Initializing the VA medication reference terminology using UMLS metathesaurus co-occurrences. *Proc AMIA Symp* 2002; 116–120.
17. Elkin PL, Carter JS, Nabar M, Tuttle M, Lincoln M, Brown SH. Drug knowledge expressed as computable semantic triples. *Stud Health Technol Inform* 2011; 166: 38–47.
18. Wright A, Chen ES, Maloney FL. An automated technique for identifying associations between medications, laboratory results and problems. *J Biomed Inform* 2010; 43(6): 891–901.
19. Brown SH, Miller RA, Camp HN, Guise DA, Walker HK. Empirical derivation of an electronic clinically useful problem statement system. *Ann Intern Med* 1999; 131(2): 117–126.
20. Zeng Q, Cimino JJ, Zou KH. Providing concept-oriented views for clinical data using a knowledge-based system: an evaluation. *J Am Med Inform Assoc JAMIA* 2002; 9(3): 294–305.
21. Chen ES, Hripsak G, Xu H, Markatou M, Friedman C. Automated acquisition of disease drug knowledge from biomedical and clinical documents: an initial study. *J Am Med Inform Assoc JAMIA* 2008; 15(1): 87–98.
22. Kilicoglu H, Fiszman M, Rodriguez A, Shin D, Ripple A, Rindflesch TC. Semantic MEDLINE: A Web Application for Managing the Results of PubMed Searches 2008; 69–76.
23. Duke JD, Friedlin J. ADESSA: A Real-Time Decision Support Service for Delivery of Semantically Coded Adverse Drug Event Data. *AMIA Annu Symp Proc AMIA Symp AMIA Symp* 2010; 2010: 177–181.

24. Wu Y, Wright A, Xu H, McCoy AB, Sittig DF. Development of a Unified Computable Problem-Medication Knowledge Base. AMIA Annu Symp Proc AMIA Symp AMIA Symp 2014; 2014.
25. Berk RA. An introduction to ensemble methods for data analysis. *Sociol Methods Res* 2006; 34(3): 263–295.
26. Tapscott D. *Wikinomics : how mass collaboration changes everything*. New York: Portfolio; 2006.
27. Howe J. The rise of crowdsourcing. *Wired Mag* 2006; 14(6): 1–4.
28. Giles J. Internet encyclopaedias go head to head. *Nature* 2005; 438(7070): 900–901.
29. Ekins S, Williams AJ. Reaching out to collaborators: crowdsourcing for pharmaceutical research. *Pharm Res* 2010; 27(3): 393–395.
30. Hughes S, Cohen D. Can Online Consumers Contribute to Drug Knowledge? A Mixed-Methods Comparison of Consumer-Generated and Professionally Controlled Psychotropic Medication Information on the Internet. *J Med Internet Res* 2011; 13(3).
31. Brownstein CA, Brownstein JS, Williams DS 3rd, Wicks P, Heywood JA. The power of social networking in medicine. *Nat Biotechnol* 2009; 27(10): 888–890.
32. Parry DT, Tsung-Chun Tsai. Crowdsourcing techniques to create a fuzzy subset of SNOMED CT for semantic tagging of medical documents. 2010 IEEE International Conference on Fuzzy Systems (FUZZ). IEEE 2010; 1–8.
33. Waghlikar KB, MacLaughlin KL, Kastner TM, Casey PM, Henry M, Greenes RA, Liu H, Chaudhry R. Formative evaluation of the accuracy of a clinical decision support system for cervical cancer screening. *J Am Med Inform Assoc* 2013; Apr 5; amiajnl – 2013–001613.
34. Good BM, Su AI. Crowdsourcing for. *Bioinformatics* 2013; btt333.
35. Sweidan M, Williamson M, Reeve JF, Harvey K, O'Neill JA, Schattner P, Snowdon T. Evaluation of features to support safety and quality in general practice clinical software. *BMC Med Inform Decis Mak* 2011; 11(1): 27.
36. Hersh W. Evaluation of biomedical text-mining systems: lessons learned from information retrieval. *Brief Bioinform* 2005; 6(4): 344–356.