Published in final edited form as:

Clin Pharmacol Ther. 2012 August ; 92(2): 228–234. doi:10.1038/clpt.2012.54.

Detection of Pharmacovigilance-Related adverse Events Using Electronic Health Records and automated Methods

K Haerian¹, D Varn¹, S Vaidya¹, L Ena¹, HS Chase¹, and C Friedman¹

¹Department of Biomedical Informatics, Columbia University, New York, New York, USA

Abstract

Electronic health records (EHRs) are an important source of data for detection of adverse drug reactions (ADRs). However, adverse events are frequently due not to medications but to the patients' underlying conditions. Mining to detect ADRs from EHR data must account for confounders. We developed an automated method using natural-language processing (NLP) and a knowledge source to differentiate cases in which the patient's disease is responsible for the event rather than a drug. Our method was applied to 199,920 hospitalization records, concentrating on two serious ADRs: rhabdomyolysis ($n = 687$) and agranulocytosis ($n = 772$). Our method automatically identified 75% of the cases, those with disease etiology. The sensitivity and specificity were 93.8% (confidence interval: 88.9–96.7%) and 91.8% (confidence interval: 84.0–96.2%), respectively. The method resulted in considerable saving of time: for every 1 h spent in development, there was a saving of at least 20 h in manual review. The review of the remaining 25% of the cases therefore became more feasible, allowing us to identify the medications that had caused the ADRs.

Pharmacovigilance is an essential component of pharmaceutical safety.¹ Traditional sources of information used for detecting new, rare, and serious adverse drug reactions (ADRs) are clinical trials, pharmaceutical industry reports, and adverse-event spontaneous reporting databases.^{2–4} Electronic health records (EHRs) contain occurrences of ADRs, yet few current techniques exist to directly extract these potential pharmacovigilance signals. Even with multiple layers of research and regulation, cases of medications requiring withdrawal from the market secondary to serious adverse reactions continue to emerge, confirming the need for new, complementary methods of detection.

The mainstay for pharmacovigilance has been spontaneous reporting systems (SRSs), such as the US Food and Drug Administration's Adverse Event Reporting System and VigiBase, the World Health Organization's global Individual Case Safety Reports database.^{5,6} Limited information is available in standardized spontaneous reports.⁷ In addition, SRSs require the use of algorithms to estimate statistical measures of reporting frequency.⁸ Only a fraction of

© 2012 American Society for Clinical Pharmacology and Therapeutics

Correspondence: K Haerian (krh7003@dm.columbia.edu).

SUPPLEMENTARY MATERIAL is linked to the online version of the paper at <http://www.nature.com/cpt>

AUTHOR CONTRIBUTIONS

K.H. wrote the manuscript, designed research, performed research, and analyzed data. D.V. contributed new reagents/analytical tools. S.V. performed research and analyzed data. L.E. designed research, performed research, analyzed data, and contributed new reagents/analytical tools. H.S.C. wrote the manuscript, designed research, performed research, analyzed data, and contributed new reagents/analytical tools. C.F. wrote the manuscript, designed research, performed research, analyzed data, and contributed new reagents/analytical tools.

CONFLICT OF INTEREST

The authors declared no conflict of interest.

adverse drug events are identified and reported.⁹ The inherent limitations in SRSs, including underreporting, biased reporting rates, incomplete patient information, and indeterminate population exposure, creates a need for complementary data sources and methods.¹⁰

Others have advocated developing methods for pharmacovigilance signal detection from longitudinal observational databases such as EHRs.^{11,12} An advantage of EHR data over SRSs is the availability of more comprehensive medical information obtained during the usual course of care.¹³ EHRs typically contain elements of interest for pharmacovigilance such as the timing of medication administration, symptom development, and a detailed clinical history. Events are captured as part of standard of care, and it is therefore unnecessary to estimate a reporting frequency; this may result in a better estimation of the prevalence of ADRs. Current approaches to pharmacovigilance are beginning to recognize the advantage of utilizing supplementary sources of information such as EHRs; at the same time, the adoption of EHRs is increasing throughout the United States, potentially providing more data.^{14,15} EHRs need to be explored in a systematic way to augment and complement the information in spontaneous reporting databases. Extracting signals from EHR data for pharmacovigilance requires methods to assist with the unique issues that arise from working with non-case report data.¹⁶

A paper by Ramirez *et al.* in *Clinical Pharmacology & Therapeutics* illustrated the preponderance of serious adverse drug reactions (SADRs) that can be identified in electronic medical records.¹⁷ The authors utilized abnormal laboratory signals (ALS), such as those consistent with a diagnosis of agranulocytosis, as a method of extracting data relating to patients with potential SADRs. However, because the majority of abnormal lab values resulted from the underlying disease condition in the patients rather than from their medications, intensive manual chart review was required to identify the occurrences of SADRs. Expert manual review is expensive in terms of manpower and time.¹⁸ The work by Ramirez *et al.* is an important demonstration of the use of EHRs for pharmacovigilance purposes. However, the method could be substantially improved if the abnormal lab values that are likely to have been caused by disease- or patient-related conditions rather than medication-related conditions are filtered out. For example, musculoskeletal trauma, convulsions, or arterial ischemia may cause rhabdomyolysis with no drug being involved in these events.¹⁹ In utilizing EHR data for signal detection, if an abnormal finding such as rhabdomyolysis is the basis for extracting data on potential patients, the SADR cases must be separated from those arising from other causes. Ramirez and colleagues' study illustrated the preponderance of non-SADRs in the EHRs; they found that, in 70% of patients, an alternative cause (non-SADR) was responsible for the ALS.¹⁷ In terms of pharmacovigilance, this means that 70% of the cases that a knowledge expert spends time reviewing can clearly be disregarded as being non-SADRs.

One approach for extracting comprehensive data from unstructured clinical notes, like those in EHRs, is to use natural-language processing (NLP). NLP systems can extract medication information, patient problem lists, and comprehensive clinical information.²⁰⁻²² We used the Medical Language Extraction and Encoding System (MedLEE), a clinical NLP system developed at Columbia University.²³ It identifies semantic structures and concepts, as well as modifiers such as time and certainty. MedLEE has been used previously for such applications as extracting adverse events from discharge summaries, assessing quality of care in cardiovascular disease, and performing automated knowledge acquisition.²⁴⁻²⁶ Our research was aimed at developing a method to substantially reduce the need for manual review and thereby increase the efficiency of the manual review process. The approach involves an automated method that combines NLP with an expert-generated knowledge source that we are creating, called the related disease identifier (RDI). This will enable researchers to quickly filter out, either prospectively or retrospectively, recurrent abnormal

signals that are disease-related so that identification efforts can focus on finding potential ADRs.

In the present study, we focused on patients with an ALS consistent with rhabdomyolysis or agranulocytosis. We then applied our automated method and assessed the sensitivity and specificity of the method by comparing the classification results to a gold standard generated by expert manual review by two physicians. In addition, we applied the method retrospectively to a 5-year time period. We present the generalized findings, including identification of several medications and combinations of medications that resulted in either of the two ADRs: rhabdomyolysis or agranulocytosis.

RESULTS

Part 1: automated method for identification of suitable cases

ALSs and corresponding ADR cases identified—We used the EHRs of patients at New York Presbyterian Hospital after obtaining institutional review board approval. The number of hospitalizations during our study period from 2004 to 2009 was 199,920. The cases of rhabdomyolysis and agranulocytosis were detected from laboratory values. The diagnosis of rhabdomyolysis was based on a cutoff plasma creatine kinase (CK) level 5 times the upper limit of the normal reference range (i.e., CK level of ~1,000 U/l, five times the upper limit of normal).²⁷ The data contained a total of 8,172 CK test results that showed values higher than 5 times the upper limit of normal. These tests corresponded to 3,213 unique individuals (patients often undergo more than one test). The existence of a positive laboratory signal triggered a query about the existence of a corresponding patient discharge summary within 7 days after the abnormal laboratory test value. For many patients, there was no discharge summary because they had been examined in the outpatient clinic, were admitted for fewer than 3 days, or were emergency room patients only; patients with no discharge summary were excluded from the study. There were 687 patients with a CK >5 times the upper limit of normal for whom discharge summaries were available. Agranulocytosis was defined as neutrophils <500 mm³, hemoglobin >10 g/dl, and platelet count >100,000 × 10⁹/l.^{28,29} There were 3,886 laboratory test results that met the criteria for agranulocytosis. Of these, we found 772 unique individuals within our chosen time frame.

A total of 1,459 patients with an ALS were included for further ADR workup (rhabdomyolysis, $n = 687$; agranulocytosis, $n = 772$). Application of our RDI filtered out the data for 75% of these patients with a disease-related ALS, representing as many as 1,095 cases (522 rhabdomyolysis; 573 agranulocytosis).

Evaluation of the method: NLP and RDI—Two physicians independently reviewed the EHR data, which consisted of 275 randomly selected cases (see Methods) in order to assess the etiology of the ALS signal. Interreviewer agreement for classification of disease etiology of ALS was calculated using Cohen's κ . The Cohen's κ was 0.70, indicating substantial agreement between reviewers.³⁰ A comparison of the RDI classification of patients vs. the expert manual reviewers' classification is shown in Table 1. The overall sensitivity of the NLP-plus-RDI system was 93.8%, and the specificity was 91.8%. The positive predictive value was 95.4%; the negative predictive value was 89.1%.

Part 2: manual review for ADR detection

After part 1 was complete, we manually reviewed the remaining patients to identify ADRs. There were 364 ALS cases (25%) for which the NLP-processed discharge summary text was not disease-related per our RDI knowledge base. The process used for identifying cases is illustrated in Figure 1. These 364 cases were considered potential ADRs; each of the

corresponding discharge summaries was then manually reviewed. The findings are described below.

Rhabdomyolysis—After application of the NLP-plus-RDI approach and thereby excluding non-SADR cases based on information on diseases known to cause elevated CK levels, there were 165 cases remaining for further ADR workup. A manual review of these cases uncovered 34 with drug-induced rhabdomyolysis. In patients with elevated CK levels, 4.9% were the result of an ADR. The number of rhabdomyolysis ADR cases attributable to each drug, including delineation in cases in which there was a co-suspected or interacting drug, is shown in Figure 2.

Agranulocytosis—After application of the NLP-plus-RDI approach to exclude non-SADR cases based on information on diseases known to cause agranulocytosis, 199 cases remained for further ADR workup. A manual review of these cases revealed the following: 125 were related to chemotherapy, 31 were related to transplant immunosuppression, and 12 were associated with neither chemotherapy nor immunosuppression but were drug-induced. Medications were responsible for 21.8% of the agranulocytosis cases, of which 1.6% were not chemotherapy or immunosuppression medications. The drugs responsible for the ADR agranulocytosis are shown in Figure 3. Our findings included a case of clopidogrel, which is infrequently associated with agranulocytosis.^{31,32} The following drugs were co-suspected or interacting combinations responsible for the ADR: sulfasalazine and gabapentin, azathioprine and valganciclovir (Valcyte), Bactrim (trimethoprim/sulfamethoxazole) and Keflex (cephalexin), carbamazepine and clonazepam, and Dilantin and Ceftriaxone.

Cohort characteristics—Patient demographic data are shown in Table 2. It is known that baseline CK levels are higher in men than in women and that these levels are higher in blacks than in whites, Hispanics, and Asians.³³ Our study findings showed a higher proportion of black men in our patient population with rhabdomyolysis.

Race-related differences in prevalence of agranulocytosis are shown in the table. Certain ethnic and racial groups, such as those with African ancestry, have lower baseline neutrophil counts.³⁴ We did not observe a disproportionate increase in agranulocytosis in racial groups known to have these lower baselines.

DISCUSSION

Our study focused on two serious ADRs: rhabdomyolysis and agranulocytosis, both of which are known to be potentially life-threatening.^{19,35-38} We applied an automated method of ADR signal detection to demonstrate a more efficient way to extract pharmacovigilance data from EHRs. Using the NLP-plus-RDI approach, 1,096 cases (75%) in which the underlying disease was responsible for the ADR were identified without the need for manual review. The sensitivity and specificity of the automated method were 93.8 and 91.8%, respectively. The automated method resulted in considerable time savings: for every 1 h spent in development, there was a saving of at least 20 h in manual review.

Comparison with other EHR methods

Most computerized methods were developed for SRSs; in these databases, the original reporter has already screened the data for confounders or covariates. Current efforts to use EHR data, however, will need to address the confounding problem. One current approach primarily uses EHRs to strengthen signals first detected in SRSs.³⁹ By contrast, our approach involves finding ADRs directly in the EHR. We anticipate that this may allow us to detect ADRs that might not be reported to an SRS, given factors such as reporting bias.

Others have used simulated EHR data to study the performance of statistical methods.⁴⁰ Our method was developed and verified using real EHR data, including narrative records, rather than simulated data, because the latter may lack the multilevel confounding present in real EHRs.

An advantage of our method is that it does not rely on a pre-determined hypothesis for identifying ADRs. The potential drug candidates are not pre-identified and then studied; instead, they are isolated as the last attributable cause left standing. For example, we need not prospectively identify the drugs that are to be investigated, such as “patient on medication X AND patient’s serum lab > Y,”⁴¹ which would require the predetermined variables for X to be explicitly coded into the automated method. An important advantage of our approach is that not only will known drugs be identified but also the potential exists for novel (e.g., those not detected through clinical trials or other methods) and previously unidentified ADRs to be discovered. Another advantage of our method is that it is not dependent on an abnormal lab signal but is also applicable to adverse events noted elsewhere in the EHR, including in the textual notes. The EU-ADR (Exploring and Understanding Adverse Drug Reactions) project identified 23 key adverse events to monitor in EHRs, of which only 10 can be identified using an ALS.⁴² The other 13, which include conditions such as bullous skin eruptions, could be identified using NLP or ICD-9 codes (although some conditions lack ICD-9 codes (e.g., cardiac valve fibrosis)). Here the RDI system could be used to extract cases suspected of being ADRs.

Another approach to pharmacovigilance is to study the application of disproportionality methods to structured health databases such as claims data.⁴³ The use of these data has inherent disadvantages because they are very limited and only certain elements are coded for billing purposes. The NLP aspect of our method has the benefit of using comprehensive data occurring in narrative notes. Our automated method was successful as the next step involving the identification of ADRs in real, unstructured EHR data. Ultimately, we believe that a variety of methods will be necessary to monitor and detect the wide variety of drugs and ADRs that exist.

Limitations

One limitation of our study is its reliance on complex knowledge. Certain confounding diseases or conditions do not consistently result in the same findings but instead depend on timing and degree. Severe trauma can result in muscular skeletal damage and release of CK from the muscle cells, but a small cut on the finger would not elevate CK to a significant degree. The choice to include certain concepts in the RDI system is not designed to address this degree of granularity. This limitation materialized in our manual review. In this review, there were eight cases in which the NLP-plus-RDI method classified a case as having disease etiology for the ALS, but the gold standard did not. The use of our method would therefore have excluded these cases from further review, potentially resulting in the loss of actual ADR cases. In our study, none of these eight cases contained evidence of an ADR, and all of them lacked enough information to enable definitive determination of the etiology of the ALS. This was a particular problem with respect to substance abuse—it was unclear when the patient had last used illicit substances such as cocaine, which can cause rhabdomyolysis or neutropenia if adulterated with levamisole.^{44,45} Of note, if the automated method failed to classify a case as “disease etiology for ALS,” indicating that a disease-related cause was present (11 cases in our study), then the case would be included in the next step of our process, namely, manual review. This type of error would not result in the loss of ADR information, as the cause would be determined by manual review, albeit with additional effort and time.

A second limitation was the selection of patients for the study. We chose patients who had been discharged within 7 days after an abnormal laboratory result. This discharge cutoff was based on the decision of two clinical experts who determined that this time period was clinically reasonable for our specific ADRs. Also, the longer the time between the abnormal result and the discharge, the lower the likelihood that an ADR occurred. For example, some discharge summaries referred to a time point years after the abnormal laboratory value had been recorded and were therefore obviously unrelated. Our goal was to find acute, serious, transient ADRs resulting in hospitalization. However, our window of interest may have excluded some patients with longer hospital stays. It also excluded those admitted for fewer than 3 days and those who were examined only in outpatient clinics or the emergency room. The longitudinal relationship may need to be modified when other ADRs are studied. Future areas of study could be the use of different types of notes (e.g., outpatient) or adjustments to the longitudinal window of study.

A third limitation, and a point we wish to emphasize, is that the use of this method removes the possibility of studying disease-exacerbated adverse drug events. We are not able to evaluate how a drug worsens the effect of a disease. As a case in point, chemotherapeutics further lower the neutrophil counts that are already reduced in certain cancers. Our method excludes confounders completely, and therefore a researcher who wanted to study how a drug may exacerbate an adverse effect of a disease would need to use a different method. However, our method does ensure that a clinical expert reviews the record for each finding and allows for identification of single medications or combinations of medications.

Summary

By automatically assigning the etiology of a laboratory signal to either the disease itself or to adverse events, the extent of chart review that must be done by the expert reviewer is minimized. Our method could augment a laboratory-signal pharmacovigilance program, such as the one implemented by Ramirez et al.¹⁷ It would also be useful in removing confounding from statistical data-mining methods applied to EHRs. The RDI alone could be embedded in a computer decision-support system to alert physicians to the likelihood of an ADR when no matching comorbidities are present in a patient's current medical history. Although our method required some expert time to choose the diseases to be incorporated into the RDI, it was a one-time investment and the automated method resulted in considerable time savings overall. Our method achieved a high level of sensitivity and specificity in its ability to correctly identify the records of cases in which a disease was responsible for the abnormal laboratory signal.

In our study, we found that almost half of the ADRs were attributable to a combination of coadministered drugs with the same known side effects. These medications may have had their effects in one of the following ways: (i) as co-suspects, with only one actually causing the ADR; (ii) both contributing to a higher cumulative effect; or (iii) a drug metabolism interaction, such as is seen with drugs that utilize the same cytochrome P450 enzyme for metabolism. Other groups have focused on the importance of drug-drug interactions and biclustering to identify novel adverse drug events.^{39,46}

In conclusion, we present a novel method for reducing the time and amount of work required to detect drug-related serious ADRs from EHR data. Although it is challenging to correct for all the covariates that are present in EHR data, our method provides a way to prespecify these and exclude cases that would otherwise need to be corrected for when statistical modeling was applied. We compared our method to the gold standard of manual review and showed that the method has high sensitivity and specificity. We then applied our method to 5 years of inpatient EHR data to extract information on rhabdomyolysis and agranulocytosis ADRs. We found cases of well-established drugs corresponding to the

ADR, drugs that had been implicated only in case reports, and drug interactions causing the ADRs. Multiple drug exposure is a known independent risk factor for serious ADRs.⁴⁷ Health-care providers should not only be mindful of the number of medications that they prescribe to an individual but also be cautious when prescribing medications with overlapping side-effect profiles.

METHODS

Study setting

The study was conducted at the Columbia University Medical Center/New York Presbyterian Hospital, an 888-bed hospital in New York City. The institutional review board approved the protocol for this study. Institutional-level, retrospective data were obtained from the clinical data warehouse, and we studied laboratory data, medication lists, and information in discharge summary notes from 2004 to 2009.

Methods framework

Step 1: laboratory signal detection. The rhabdomyolysis cases were selected based on serum CK level of >5 times normal, and the agranulocytosis cases were selected based on neutrophil count <500 mm³, hemoglobin level >10 g/dl, and platelet count >100,000 × 10⁹/l. The laboratory values were obtained from a structured laboratory table in the New York Presbyterian Hospital Clinical Data Warehouse. The existence of a positive laboratory signal triggered a query for the existence of a discharge summary after the date of the abnormal laboratory value but within 7 days. If a discharge summary was available, then that patient was included in the study.

Step 2: NLP of discharge summaries. To capture the data from the unstructured narrative discharge summaries with a view to filtering concepts, we used the MedLEE System to produce a structured, coded output, which was then stored in a database.⁴⁸ MedLEE identified the semantic structures and concepts that we used to determine diseases, symptoms, and medications noted in each report. MedLEE also produced modifiers such as time and certainty, which we used to include events of high certainty and to exclude events that were negated or had occurred in the past. The concepts map to standardized Unified Medical Language System (UMLS) concept unique identifiers.⁴⁹ The processing of the data in each discharge summary would result in a list of diagnoses and symptom concept identifiers, which would then be compared to the extensive list of concept identifiers in the RDI.

Step 3: development of the RDI. Two clinical experts generated a list of known medical conditions that commonly result in the laboratory signal of interest. Elevation in the levels of CK is known to occur in medical conditions such as myocardial infarction, stroke, Duchenne muscular dystrophy, and muscle compression or overuse. Low neutrophil counts occur in patients with Kostmann syndrome, glycogen-storage disease type 1b, and in several malignancies and other diseases. As the next step, a list of diseases and conditions present in patients with the adverse event was extracted and used by a clinical expert to identify different granularities in the terms used to document the medical conditions of interest. (Supplementary Materials and Methods online contains further details.) This method allowed us to capture a more comprehensive list of relevant concepts, such as ST segment elevation in myocardial infarction and acute myocardial infarction.

All the identified non-ADR causes were then coded using the UMLS. The standardized list of UMLS concept unique identifiers was used as the knowledge-based filter that we refer to as the RDI, to automatically remove cases in which the abnormal laboratory value could clearly be attributed to a patient's disease state. The remaining cases were those in which the

abnormal laboratory value was caused by a medication, an unknown etiology, or a filtering error. More information on the RDI can be found in Supplementary Materials and Methods online.

Step 4: evaluation of design (analysis of NLP-plus-RDI categorization of patients vs. a gold standard set of cases). A block random set of cases was selected to evaluate the method. This involved selecting cases (discharge summaries), generating a gold standard using physician review, using the automated NLP-plus-RDI method, and comparing and analyzing the results.

The gold-standard set of cases was created by two clinical experts (physicians) who read and annotated 150 discharge summaries for rhabdomyolysis and 125 discharge summaries for agranulocytosis. The sample size was associated with a 95% confidence interval with a margin of error of $\pm 5.3\%$ per the Krejcie and Morgan formula.⁵⁰ The patients were block-randomized and chosen from cases in which an abnormal laboratory value was present. Interrater agreement was calculated using Cohen's κ . Differences were resolved by consensus. For each discharge summary, the presence of a disease etiology (not ADR) was documented if present. The same group of discharge summaries was categorized using the NLP-plus-RDI system. The results of the NLP-plus-RDI automated system were then compared with the categorization determined by our gold standard, namely, the independent clinical expert manual review. These results were used to calculate the sensitivity and specificity.

Step 5: apply the NLP/RDI to the full data set. After measuring the performance of the NLP/RDI automated system, we applied it to the full set of discharge summaries established in step 1. All the patients for whom the data contained a concept unique identifier that matched a concept in the RDI were classified as having a disease-related ALS; these were excluded from further review. The other cases were then manually reviewed to determine the cause of the ALS. When an ADR was found, the suspect drug or combination of drugs was documented.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank Martin Lindquist for statistical review. We acknowledge the receipt of the following: NLM grants R01:LM010016, R01:LM010016-0s1, R01:LM010016-0s2, R01:LM008635, and R01:LM06910. K.H. received NLM training grant t15:LM007079.

References

1. Härmak L, van Grootheest AC. Pharmacovigilance: methods, recent developments and future perspectives. *Eur. J. Clin. Pharmacol.* 2008; 64:743–752. [PubMed: 18523760]
2. Atuah KN, Hughes D, Pirmohamed M. Clinical pharmacology: special safety considerations in drug development and pharmacovigilance. *Drug Saf.* 2004; 27:535–554. [PubMed: 15154826]
3. Hartford CG, et al. Pharmacovigilance during the pre-approval phases: an evolving pharmaceutical industry model in response to ICH E2E, CIOMS VI, FDA and EMEA/CHMP risk-management guidelines. *Drug Saf.* 2006; 29:657–673. [PubMed: 16872240]
4. Rodriguez EM, Staffa JA, Graham DJ. The role of databases in drug postmarketing surveillance. *Pharmacoepidemiol. Drug Saf.* 2001; 10:407–410. [PubMed: 11802586]

5. Lindquist M, Edwards IR. The WHO Programme for International Drug Monitoring, its database, and the technical support of the Uppsala Monitoring Center. *J. Rheumatol.* 2001; 28:1180–1187. [PubMed: 11361210]
6. M, Lindquist, Vigibase, the WHO Global ICSR Database system: Basic Facts. *Drug Info. J.* 2008; 42:409–419.
7. S, Ahmad, Spontaneous reporting in the United states. In: Strom, BL., editor. *Pharmacoepidemiology.* Wiley and sons; West sussex, UK: 2005.
8. Hauben M, Bate A. Decision support methods for the detection of adverse events in post-marketing data. *Drug Discov.* 2009; 14:343–357.
9. Jha AK, Laguet J, Seger A, Bates DW. Can surveillance systems identify and avert adverse drug events? A prospective evaluation of a commercial application. *J. Am. Med. Inform. Assoc.* 2008; 15:647–653. [PubMed: 18579834]
10. Goldman SA. Limitations and strengths of spontaneous reports data. *Clin. Ther.* 1998; 20(suppl. C):C40–C44. [PubMed: 9915089]
11. Wilson AM, Thabane L, Holbrook A. Application of data mining techniques in pharmacovigilance. *Br. J. Clin. Pharmacol.* 2004; 57:127–134. [PubMed: 14748811]
12. Madigan D, Ryan P. What can we really learn from observational studies?: the need for empirical assessment of methodology for active drug safety surveillance and comparative effectiveness research. *Epidemiology.* 2011; 22:629–631. [PubMed: 21811110]
13. Weaver J, Willy M, Avigan M. Informatic tools and approaches in postmarketing pharmacovigilance used by FDA. *AAPS J.* 2008; 10:35–41. [PubMed: 18446503]
14. Shea S, Hripcsak G. Accelerating the use of electronic health records in physician practices. *N. Engl. J. Med.* 2010; 362:192–195. [PubMed: 20089969]
15. Wilke RA, et al. The emerging role of electronic medical records in pharmacogenomics. *Clin. Pharmacol. Ther.* 2011; 89:379–386. [PubMed: 21248726]
16. Wang X, Hripcsak G, Markatou M, Friedman C. Active computerized pharmacovigilance using natural language processing, statistics, and electronic health records: a feasibility study. *J. Am. Med. Inform. Assoc.* 2009; 16:328–337. [PubMed: 19261932]
17. Ramirez E, et al. A pharmacovigilance program from laboratory signals for the detection and reporting of serious adverse drug reactions in hospitalized patients. *Clin. Pharmacol. Ther.* 2010; 87:74–86. [PubMed: 19890254]
18. Phansalkar S, Hoffman JM, Hurdle JF, Patel VL. Understanding pharmacist decision making for adverse drug event (ADE) detection. *J. Eval. Clin. Pract.* 2009; 15:266–275. [PubMed: 19335483]
19. Huerta-Alardín AL, Varon J, Marik PE. Bench-to-bedside review: Rhabdomyolysis – an overview for clinicians. *Crit. Care.* 2005; 9:158–169. [PubMed: 15774072]
20. 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; 17:19–24. [PubMed: 20064797]
21. Meystre SM, Haug PJ. Randomized controlled trial of an automated problem list with improved sensitivity. *Int. J. Med. Inform.* 2008; 77:602–612. [PubMed: 18280787]
22. Friedman C, Hripcsak G, Shagina L, Liu H. Representing information in patient reports using natural language processing and the extensible markup language. *J. Am. Med. Inform. Assoc.* 1999; 6:76–87. [PubMed: 9925230]
23. Hripcsak G, Friedman C, Alderson PO, DuMouchel W, Johnson SB, Clayton PD. Unlocking clinical data from narrative reports: a study of natural language processing. *Ann. Intern. Med.* 1995; 122:681–688. [PubMed: 7702231]
24. Melton GB, Hripcsak G. Automated detection of adverse events using natural language processing of discharge summaries. *J. Am. Med. Inform. Assoc.* 2005; 12:448–457. [PubMed: 15802475]
25. Chiang JH, Lin JW, Yang CW. Automated evaluation of electronic discharge notes to assess quality of care for cardiovascular diseases using Medical Language Extraction and Encoding system (MedLEE). *J. Am. Med. Inform. Assoc.* 2010; 17:245–252. [PubMed: 20442141]
26. Wang X, Chused A, Elhadad N, Friedman C, Markatou M. Automated knowledge acquisition from clinical narrative reports. *AMIA Annu. Symp. Proc.* 2008:783–787. [PubMed: 18999156]

27. Cervellin G, Comelli I, Lippi G. Rhabdomyolysis: historical background, clinical, diagnostic and therapeutic features. *Clin. Chem. Lab. Med.* 2010; 48:749–756. [PubMed: 20298139]
28. Standardization of definitions and criteria of causality assessment of adverse drug reactions. Drug-induced cytopenia. *Int. J. Clin. Pharmacol. Ther. Toxicol.* 1991; 29:75–81. [PubMed: 2026468]
29. Tavassoli N, et al. Detection and incidence of drug-induced agranulocytosis in hospital: a prospective analysis from laboratory signals. *Eur. J. Clin. Pharmacol.* 2007; 63:221–228. [PubMed: 17225990]
30. Landis JR, Koch GG. The measurement of observer agreement for categorical data. *Biometrics.* 1977; 33:159–174. [PubMed: 843571]
31. CAPRIE steering Committee. A randomised, blinded, trial of clopidogrel versus aspirin in patients at risk of ischaemic events (CAPRIE). *Lancet.* 1996; 348:1329–1339. [PubMed: 8918275]
32. Dieter RS. Risk of neutropenia with clopidogrel. *J. Am. Coll. Cardiol.* 2000; 36:1436–1437. [PubMed: 11028512]
33. Neal RC, Ferdinand KC, Ycas J, Miller E. Relationship of ethnic origin, gender, and age to blood creatine kinase levels. *Am. J. Med.* 2009; 122:73–78. [PubMed: 19114174]
34. Hsieh MM, Everhart JE, Byrd-Holt DD, Tisdale JF, Rodgers GP. Prevalence of neutropenia in the U.S. population: age, sex, smoking status, and ethnic differences. *Ann. Intern. Med.* 2007; 146:486–492. [PubMed: 17404350]
35. Melli G, Chaudhry V, Cornblath DR. Rhabdomyolysis: an evaluation of 475 hospitalized patients. *Medicine (Baltimore).* 2005; 84:377–385. [PubMed: 16267412]
36. Saloustros E, Tryfonidis K, Georgoulas V. Prophylactic and therapeutic strategies in chemotherapy-induced neutropenia. *Expert Opin. Pharmacother.* 2011; 12:851–863. [PubMed: 21254862]
37. Andersohn F, Konzen C, Garbe E. Systematic review: agranulocytosis induced by nonchemotherapy drugs. *Ann. Intern. Med.* 2007; 146:657–665. [PubMed: 17470834]
38. Garbe E. Non-chemotherapy drug-induced agranulocytosis. *Expert Opin. Drug Saf.* 2007; 6:323–335. [PubMed: 17480181]
39. Tatonetti NP, et al. Detecting drug interactions from adverse-event reports: interaction between paroxetine and pravastatin increases blood glucose levels. *Clin. Pharmacol. Ther.* 2011; 90:133–142. [PubMed: 21613990]
40. Schuemie MJ. Methods for drug safety signal detection in longitudinal observational databases: LGPS and LEOPARD. *Pharmacoepidemiol. Drug Saf.* 2011; 20:292–299. [PubMed: 20945505]
41. Kuperman GJ, Teich JM, Bates DW, McLatchey J, Hoff TG. Representing hospital events as complex conditionals. *Proc. Annu. Symp. Comput. Appl. Med. Care.* 1995:137–141. [PubMed: 8563253]
42. Trifirò G, et al. Data mining on electronic health record databases for signal detection in pharmacovigilance: which events to monitor? *Pharmacoepidemiol. Drug Saf.* 2009; 18:1176–1184. [PubMed: 19757412]
43. Zorych I, Madigan D, Ryan P, Bate A. Disproportionality methods for pharmacovigilance in longitudinal observational databases. *Stat. Methods Med. Res.* 2011 e-pub ahead of print 30 August 2011.
44. Roth D, Alarcón FJ, Fernandez JA, Preston RA, Bourgoignie JJ. Acute rhabdomyolysis associated with cocaine intoxication. *N. Engl. J. Med.* 1988; 319:673–677. [PubMed: 3412385]
45. Chung C, et al. Characteristic purpura of the ears, vasculitis, and neutropenia—a potential public health epidemic associated with levamisole-adulterated cocaine. *J. Am. Acad. Dermatol.* 2011; 65:722–725. [PubMed: 21658797]
46. Harpaz R, Perez H, Chase HS, Rabadan R, Hripcsak G, Friedman C. Biclustering of adverse drug events in the FDA's spontaneous reporting system. *Clin. Pharmacol. Ther.* 2011; 89:243–250. [PubMed: 21191383]
47. Macedo AF, Alves C, Craveiro N, Marques FB. Multiple drug exposure as a risk factor for the seriousness of adverse drug reactions. *J. Nurs. Manag.* 2011; 19:395–399. [PubMed: 21507111]
48. Friedman C, shagina L, Lussier Y, Hripcsak G. Automated encoding of clinical documents based on natural language processing. *J. Am. Med. Inform. Assoc.* 2004; 11:392–402. [PubMed: 15187068]

49. Browne AC, Divita G, Aronson AR, McCray AT. UMLS language and vocabulary tools. AMIA Annu. Symp. Proc. 2003; 2003:798. [PubMed: 14728303]
50. Krejcie RV, Morgan DW. Determining sample size for research activities. Educ. Psychol. Meas. 1970; 30:607–610.

Figure 1. Flowchart of automated method for facilitating pharmacovigilance event detection. ADR, adverse drug reaction; NLP, natural-language processing; RDI, related disease identifier; SADR, serious adverse drug reaction.

Figure 2. Identified single, co-suspect, and interacting drugs causing cases of rhabdomyolysis adverse reaction.

Figure 3. Identified single, co-suspect, and interacting drugs causing cases of agranulocytosis adverse reaction.

Table 1
Sensitivity and specificity of the NLP-plus-RDI method vs. the gold standard (expert manual reviewers' classification)

Condition	Number of cases evaluated	Sensitivity (95% CI)	Specificity (95% CI)	Positive predictive value (95% CI)
Rhabdomyolysis	150	96.75% (91.38–98.95)	81.58% (61.25–92.97)	95.97% (90.37–98.51)
Agranulocytosis	125	87.04% (78.48–94.20)	95.77% (87.33–98.90)	94.00% (82.46–98.44)
Overall	275	93.79% (88.87–96.70)	91.84% (84.08–96.15)	95.40% (90.83–97.85)

CI, confidence interval; NLP, natural-language processing; RDI, related disease identifier.

Table 2
Demographics of the patient population

Variable	Total hospitalizations	Unique patients	Rhabdomyolysis	Agranulocytosis
<i>N</i>	199,920	110,625	3,213	1,942
Mean age ($\pm$ SD)	52.0 ($\pm$ 24.9)	51 ($\pm$ 24.9)	49.8 ($\pm$ 25.7)	37.8 ($\pm$ 26.6)
Sex (% female)	55.9%	55.8%	31.9%	46.6%
Race (% of group)				
White	32.1%	31.8%	26.7%	37.4%
Hispanic	30.7%	22.0%	30.3%	23.6%
Black	17.2%	15.0%	24.6%	15.8%
Asian	2.0%	1.8%	2.1%	3.3%
Other/undocumented	18.0%	29.4%	23.5%	19.9%