Progress Report

Close-Out Documentation

Title: Mining complex clinical data for patient safety research

Principal Investigator and Senior Team Members:

George Hripcsak, MD, MS (PI)

Carol Friedman, PhD

Vimla Patel, PhD

Suzanne Bakken, RN, DNSc

Henry Spotnitz, MD

Alan Moskowitz, MD

Organization: The Trustees of Columbia University in the City of New York
New York, NY

Project Period: 9/01/2001 to 8/31/2004

Program Official: Stanley Edinger, PhD, Senior Science Advisor

Acknowledgment of Agency Support: This work was supported by the Agency for Healthcare Research and Quality.

Grant Number: R18 HS11806

Structured Abstract

Purpose: To use an electronic health record and natural language processing (NLP) to automate and improve medical event detection.

Scope: The project was carried out in a large, urban academic medical center with a repository of 15 years of data on 2.4 million patients seen in inpatient and outpatient areas. Electronic data included registration data, laboratory data, narrative ancillary reports (radiology, etc.), and notes written by providers (discharge summaries, resident signout notes, visit notes, etc.).

Methods: Narrative data were structured and coded using the MedLEE NLP system and merged with other coded clinical data. Queries were written by experts or derived from machine learning to detect a broad range of medical events. The accuracy of event detection was estimated (sensitivity, specificity, predictive value). Cognitive analyses and case-based reasoning were also applied.

Results: Narrative reports, such as discharge summaries and resident signout notes, were found to contain useful information for uncovering medical events. NLP successfully identified which of 45 NYPORTS events occurred in 57,452 discharge summaries, achieving a PPV of .44 and sensitivity of .27. NLP was found to complement and improve upon error detection based on coded data. Providers were found to explicitly report medical errors in the electronic health record at rates similar to voluntary error reporting (.3 to 1.9%). Cognitive studies revealed complex causation for errors and flawed decision-making processes. Case-based reasoning on an errors database was accurate. An errors terminology was developed and incorporated into the LOINC national standard.

Key Words: patient safety; medical errors; natural language processing; data mining; cognitive analysis; terminology; case-based reasoning; electronic health records

Purpose

Medical errors hurt patients, cost money, and undermine the healthcare system. The first step to reducing errors is detecting them, for what cannot be detected cannot be managed. A number of approaches have been applied to medical error detection, including mandatory event reporting, voluntary near-miss reporting, chart review, and automated surveillance using information systems. Automated surveillance promises large-scale detection, minimal labor, and, potentially, detection in real time to prevent or recover from errors. Unfortunately, large amounts of important clinical information lie locked in narrative reports, unavailable to automated decision support systems. A number of tools have emerged from medical informatics and computer science—natural language processing, visualization tools, and machine learning—as well as methods for understanding cognitive processes. We hypothesized that the electronic health record contains information useful for detecting errors and that natural language processing and other tools would allow us to retrieve the information.

Scope

We carried out our studies at Columbia University Medical Center, a large, urban academic medical center. Most of the studies were retrospective and potentially included all patients seen at the medical center since 1989, when a significant amount of clinical data began to be collected. This included approximately 2.5 million patients who had been seen in the inpatient or outpatient areas. Data from the clinical repository were prepared for analysis (using a number of tools, including natural language processing), and queries were applied to the data set to identify cases with significant medical events, including medical errors. A broad range of events was studied (including the 45 events defined in the New York Patient Occurrence Reporting and Tracking System). Incidence for most individual event types was 1% or less per admission or visit.

Methods

Overview of event detection

We developed a framework for automated event discovery from the electronic health record [10] (see also related publications [2,13–15,21,22]). Our framework consisted of the following assumptions. There exists a repository of rich clinical information that we hypothesize contains useful data about patient safety. The data are obscured, however, due to the way they are recorded. Much is in narrative form and therefore not amenable to traditional statistical analysis; even when coded, the data are stored in complex, nested structures that may be difficult to use. A set of informatics tools exists—natural language processing, machine learning, etc.—that is capable of extracting the useful patient safety information from the repository in an automated or partially automated fashion. Manual review of the paper and electronic medical record (and, in a smaller sample, interviews of the relevant care providers) can be used to create a reference standard to judge the accuracy of the automated system. Errors have structures that can be described using a systems approach to errors and using a cognitive approach. We may then use these descriptions to learn how to improve the automated system.

Within this framework, the following basic approach was reused for a series of experiments:

1. Target events—Pick the target events of interest (either an actual list of known errors or a conceptual type of error to look for).
2. Repository—Begin with the full clinical repository or a purposely defined subset.
3. Natural language processing—Use natural language processing to parse the narrative data and create a fully coded repository.
4. Queries—Generate queries that detect and classify errors. They may be generated manually or automatically.
5. Verification—Verify the accuracy of the detection and classification by manual review, thus calculating performance and adding to the database of known errors.
6. Error description—Use a systems approach or a cognitive approach to describe the newly detected errors.
7. Feedback—Based on the mistakes uncovered in step 5 and the information learned in step 6, improve the natural language processor (step 3) and the queries (step 4), and possibly steer the next selection of target events (step 1).

Target events

Because one of the primary purposes of the medical record is to document the patient's state, we expected that adverse outcomes would be recorded commonly in narrative reports, such as discharge summaries, and, for certain events, operative reports and outpatient notes. We expected that, although it might be possible to infer errors from evidence in the medical record, it would be uncommon to see errors documented explicitly as errors. We expected that near-miss events in particular would rarely if ever be documented explicitly as near misses in the medical record, although some near misses might be inferred from evidence in the record. We also expected that it would often be unclear whether an error actually occurred. The exact frequency of occurrence of such reports in the narrative part of the electronic record was unknown.

We used several approaches to detect events of interest:

a) Explicit voluntary reporting in the medical record

We looked for events that were explicitly referred to as being errors or adverse outcomes. That is, we looked for cases in which the provider stated not only that there was a medical condition but also that the condition represented an error or adverse outcome. Certain phrases were considered indicative of such events: “untoward,” “nosocomial,” “inadvertent,” “error,” “adverse,” “unexpected,” etc. With this approach, we hoped to document the rate of providers reporting errors voluntarily within the medical record. If voluntary reporting occurred in the medical record, then this approach might reveal errors that would not have been included in the other approaches below.

b) Conflicts in the record

The most common approach currently used in automated screening is to look for medical evidence of errors and adverse outcomes. Such events are usually inferred from conflicting evidence. For example, the occurrence of a myocardial infarction in a non-cardiac admission demonstrates an adverse outcome and may point to an error. There are several overlapping types of conflicts, including various kinds of mismatches of diagnoses and treatments. We characterized those conflicts as follows:

- Mention of a diagnosis (e.g., myocardial infarction, pneumothorax, aspiration pneumonia, nosocomial infection) in a case when the diagnosis would not normally be expected.
- Evidence of different providers (or the same provider over time) assigning different, competing diagnoses to the case, especially when the treatments differ or even conflict.
- Autopsy reports (a subset of pathology) that uncover diagnoses that were not mentioned before death, especially if a needed treatment was missing.
- A mismatch between treatments and diagnoses. This can be a simple lack of treatment for an important condition or a deviation from a more complex clinical guideline.
- A mismatch between treatments. Examples include drug interaction and invasive interventions in the setting of over-anticoagulation.
- An increase in the intensity of care (e.g., intensive care unit stay) in a case when it would not normally be expected.
- Diagnoses indicative of complications after a procedure.
- Repeat admission that would not normally be expected for original admitting diagnosis.
- The ordering of diagnostic tests not normally expected for the original admitting diagnosis (e.g., ventilation perfusion scan done post-surgically to rule out pulmonary embolism, or a radiology report to rule out a fracture after a fall).

In our study, we focused on a particular conflict, a mismatch between admission diagnosis and discharge diagnosis. We chose this conflict because we wanted to assess how well purely administrative data would perform. We choose the following discharge diagnoses—aspiration pneumonia, acute myocardial infarction, catheter-related infection, pulmonary embolism, and stroke—and looked for cases in which one of the discharge diagnosis codes matched the diagnosis but the admission diagnosis did not include it.

To create the reference standard, we manually reviewed electronic discharge summaries, electronic resident care transfer notes (“signout” notes), and paper charts. We took the opportunity to assess the potential for electronic notes to detect errors. We compared manual review of the electronic notes to manual review of the paper charts to derive a possible upper limit of performance that one could expect for natural language processing of electronic notes. High performance would imply that adequate information is present in the electronic notes to detect errors effectively.

c) Specific event detection—NYPORTS

We used the 45 events defined in New York State's mandatory occurrence reporting system, known as the New York Patient Occurrence Reporting and Tracking System (NYPORTS), as specific target events. Underreporting is a limitation of many mandatory reporting systems, including NYPORTS. In an effort to assist in the detection and capture of NYPORTS events, the electronic medical record and a natural language processor (MedLEE) were used to detect the occurrences recorded in the NYPORTS includes/excludes list. The goal was to assess the performance of natural language processing. We carried out two studies. In one, we choose three events and attempted detection using both coded and narrative data. In the other, we attempted to detect all 45 NYPORTS events based on parsing discharge summaries.

Clinical repository

The clinical repository comprises 15 years of data on 2.5 million patients, including both coded and narrative data. It has about five million diagnostic narrative reports and more than one million notes authored by clinicians, including discharge summaries, signout notes, visit notes, nursing notes, etc. Only a small proportion of admission notes and daily inpatient progress notes are in the system at this point in time, however.

Natural language processing

The Medical Language Extraction and Encoding System (MedLEE) [1] was developed at Columbia University (and Queens College, CUNY, NY) during the past 15 years. It initially was developed for the domain of radiological reports of the chest. It was designed using a modular approach to facilitate extension to new domains and applications. There are five programming components. The first is the preprocessor, which reads in the report, segments it into sections and sentences, and performs lexical lookup. This component uses a lexicon to categorize words and phrases and specify their target forms. For example, "abdominal" is a lexical entry that has body location as a category and abdomen as a target form. The preprocessor uses three other knowledge bases: full forms for abbreviations (e.g., CHF for congestive heart failure); sections of reports (e.g., impression, clinical information); and contextual rules used to resolve ambiguous terms (e.g., discharge from ER vs. discharge from hospital). The second component is the parser, which interprets the relations among the terms of a sentence and generates an intermediate structured output form. It uses a grammar containing rules specifying sequences of well-formed categories in which each sequence is associated with an interpretation. For example, "no relief of pain" is associated with the sequence of categories NEG CHANGE PREP FINDING. This is a well-formed sequence interpreted so that the primary information is the FINDING pain, which is modified by CHANGE (i.e., relief), which in turn is modified by NEG (i.e., no). The error recovery component is used if a parse of the complete sentence cannot be obtained. In that case, the sentence is segmented at certain points and a parse is attempted of each segment; this component is used to increase sensitivity, although specificity may be reduced somewhat. The next component uses compositional maps to compose phrases that have been separated in the sentence. For example, this component would obtain the phrase "swollen extremities" from the sentence "extremities appeared to be very swollen." The last module performs the encoding to map the findings to a controlled clinical vocabulary. To do this, it uses a table associating target output terms to codes. The vocabulary can be the UMLS, SNOMED, or a local vocabulary. For example at

Columbia, the output is transformed to MED codes. Finally, the intermediate structured output form is mapped to XML form. This form is ideal, because it facilitates retrieval of the structured output form and maps the original report to the structured output, which facilitates highlighting applications.

As part of this project, MedLEE was extended to support report types, such as resident signout notes and pathology reports, that were not already supported.

Queries to detect events

Queries for event detection are based on inclusion and exclusion criteria applied to the coded data and the coded output of MedLEE. The system is able to apply the criteria to 30,000 discharge summaries in 4 minutes.

Verifying events

To create reference standards for each of the studies, clinical experts applied defined criteria (varied with the study) to the electronic and paper patient charts. Performance was quantified in terms of positive predictive value, sensitivity, and specificity (wen possible). Expert reliability was estimated.

Case-based reasoning on an errors database

Reported errors are collected into an errors database. When new errors are reported, patient safety personnel need to determine whether similar errors have occurred and what has been done about them. We studied the use of case-based reasoning to identify relevant past errors when a new error is detected, based on defined parameters collected during the error reporting process. (This work was done in collaboration with AHRQ health systems grant U18 HS11905.)

Cognitive studies

To shed light on the source of errors and to provide clues to facilitate error detection and, in the future, prevention, we (led by Vimla Patel) performed a number of cognitive analyses as part of this project. A cognitive team used direct observation and review of documents to collect data and cognitive analysis. The details of each are given in the results section below.

Errors terminology

Identifying and characterizing medical errors requires an error terminology. As part of this project, we (led by Suzanne Bakken) developed an error terminology to be incorporated into Logical Observation Identifiers, Names, and Codes (LOINC), a national standard terminology.

Results

Extension of MedLEE

We extended MedLEE to improve its ability to distinguish comorbidities from complications [3]. We extended its parsing to cover resident signout notes [5] and structured pathology reports.

Explicit voluntary error reporting in the medical record

We studied the rate at which physicians document medical errors in the electronic medical record (i.e., explicitly state that an error occurred) [11,12]. We used keywords searches using terms such as “mistake,” “error,” “incorrect,” “inadvertent,” and “iatrogenic.” We found that physicians document errors in the record at a rate similar to voluntary error reporting. For example, use of a keyword search on the electronic record is estimated to have detected .3% to 1.9% of medication errors. The positive predictive value for different keywords ranged from 3% to 24% in discharge summaries, 0% to 36% in resident signout notes, and 0% to 22% in outpatient visit notes. “Mistake,” “inadvertent,” and “iatrogenic” were the most predictive terms studied. The predictive value was best if the word occurred in the hospital course section of the discharge summary.

Conflicts in the record

We detected conflicts in the medical record that might possibly signal that a medical adverse event has occurred during a patient’s hospital stay by comparing admission and discharge diagnoses [19]. We culled a subset of the inpatient records for the years 1990-1999 when patients were discharged with (but not admitted for) aspiration pneumonia, acute myocardial infarction, catheter-related infection, pulmonary embolism, or stroke. We manually reviewed the discharge summaries and signout notes for a random sample of these records and showed that we can obtain a positive predictive value of 0.41 with a sensitivity of 0.84, and specificity of 0.48, for detecting true errors in a cohort of all patients with the relevant diagnoses (e.g., anyone with a myocardial infarction). The estimated ROC area of 0.66 reveals that the predictive value of the method is only moderate and would need to be paired with other techniques to be effective.

We also studied the amount of useful information recorded in electronic discharge summaries and resident signout notes, using manual chart review as a gold standard [19]. We found that manual review of the electronic notes achieved a sensitivity of 0.57, a specificity of 0.96, a positive predictive value of 0.95, and an ROC area of 0.77. This demonstrates that the electronic notes do contain useful patient safety information and that natural language processing (if it works) has the potential to improve error detection.

Specific event detection—NYPORTS

We studied the detection of errors reportable under NYPORTS. In one study, aspiration pneumonia in the setting of conscious sedation, perioperative myocardial infarction, and foreign body retention in surgery were assessed (manuscript in preparation). For each event type, a query was performed based on the include/exclude criteria. The queries were performed on the electronic medical record, using a combination of coded data (ICD9, laboratory data, pharmacy data), narrative data (e.g., discharge summaries and radiology reports), and MedLEE parsed data. We found that natural language processing (NLP) improved detection of aspiration pneumonia over coded data (sensitivity .75 and positive predictive value .14 with NLP and coded data; and sensitivity .50 and positive predictive value .13 with coded data alone). NLP had little improvement for perioperative myocardial infarction (sensitivity .82 and positive predictive value .31 for coded data), most likely because coded laboratory troponin levels were highly accurate in identifying infarction patients. For detection of retained foreign bodies, it shifted the detection curve (sensitivity .44 and positive predictive value .57 with NLP alone; and sensitivity .67 and positive predictive value .19 with coded data alone). More analyses are pending.

In a separate study, we wrote queries for all 45 NYPORTS events based on MedLEE output for discharge summaries alone [24]. The results were impressive. Of 57,452 total electronic discharge summaries analyzed, 704 of 1,590 summaries returned by the system were true NYPORTS events, and manual reporting in that period totaled 294 NYPORTS events. Consequently, for the system as a whole, the estimated positive predictive value was 0.44 (95% CI 0.42-0.47), and sensitivity was 0.27 (CI 0.22-0.31). Estimated overall system specificity was 0.99 (CI: 0.98-0.99) based on an estimated prevalence for each event type of less than 1%. This compares to previous predictive values for specific event reporting based on narrative reports of 7 to 12% [25,26]. One study of nonspecific event reporting [27] achieved a positive predictive value of 52%, but the estimated prevalence of errors in the underlying sample was 45%, signifying only minor improvement over random sampling.

Case-based reasoning on an errors database

We applied case-based reasoning [4] to an errors database [23]. Data were collected to assess the performance of HAWK, a case-based reasoning retrieval system that operates on MERS-TM, a transfusion errors database. Given a new case, similar old cases are retrieved according to a set of rules. Based on manual review to ascertain case similarity, the basic HAWK retrieval system (which used expert-authored rules) achieved an ROC area .96. That is, it is very effective in retrieving similar cases. We also tested several machine learning methods to improve performance (using learned rules instead of expert-authored rules), but performance via machine learning at best equaled and in some cases lagged far behind expert-authored rules.

Cognitive studies

We (led by Vimla Patel) addressed the cognitive issues surrounding the reporting and cause of medical events. We performed a number of studies funded under this CLIPS grant to elucidate cognitive causes of errors that may, in turn, be used to improve detection and prevention [7,8]. The group studied the electronic recording and presentation of clinical information from a cognitive point of view, looking at various levels of clinician expertise [17]. The group found that structured data (as opposed to narrative data) resulted in better recall and better inferences for novice and intermediate level clinicians. This points to a need for either structured data entry or effective natural language processing to structure the data.

The group carried out a series of studies related to medical errors in the use of infusion pumps. The group conducted a heuristic evaluation of an infusion pump interface and showed that this complemented traditional outcomes and error detection studies. In another study, the group found that different stakeholders (administrators, engineers, nurses, and physicians) interpreted error causation differently and that there was a greater tendency to assign human blame to errors when errors were presented retrospectively.

The group looked at the administrative decision-making process for selecting devices such as infusion pumps and how those decisions could affect medical error rates [16]. The study showed that, nominally, the decisions are made collaboratively; in truth, there is little evidence of coordination, and human factors are largely ignored. Therefore, the attempt to reduce medical errors by detecting them and feeding the information back to decision makers will be thwarted by the current decision-making process.

Errors terminology

To promote generalization, we (led by Suzanne Bakken) addressed the terminologic aspects of data mining for patient safety purposes through a number of analyses, and we incorporated our work into the national standard. We examined the extent to which the aspects of a root cause analysis as implemented in the Medical Event Reporting System – Total Hospital (MERS-TH) (e.g., consequent events, contributing factors) could be represented using the semantic structure of clinical measurements in the Logical Observation Identifiers, Names, and Codes (LOINC) Database, a public domain coding system. Based on the positive results of the analysis, a LOINC representation was created for the MERS-TH root cause analysis event description items, and 26 items were subsequently approved for inclusion in LOINC at the January 2003 meeting. This work was shared with the Institute of Medicine Patient Safety Data Standards Committee.

Other project work

In addition, the following related work was partially funded by the project. We studied the abstraction and filtering of diagnoses from narrative reports [20] to better improve detection accuracy. We applied similar techniques to study the achievement of residency competencies during training [18]. We studied event monitors, which can be used to deploy event detection logic in real time [9]. We participated in a mobile computing project aimed to improve patient safety by making data and detected event more readily available [6].

Collaborations

Parts of the work were carried out in collaboration with other institutions, including David W. Bates and colleagues from Brigham and Women's Hospital and R. Scott Evans and colleagues from Intermountain Health Care on the methodologies of detecting adverse events using information technology [14,15], as well as Jiajie Zhang from University of Texas on cognitive studies.

Limitations

The main limitation of our work is that it has been performed in a single center. For example, it remains to be proven that queries for medical events will work as well in other institutions. To address this, a P41 infrastructure grant has been submitted to the NIH to disseminate the MedLEE system and its queries, including the error detection queries developed during this grant.

Project publications (ordered by publication date)

1. Hripcsak G, Austin JHM, Alderson PO, Friedman C. Use of natural language processing to translate clinical information from a database of 889,921 chest radiographic reports. Radiology 2002;224:157–63.
2. Bakken S, Cimino JJ, Hripcsak G. Promoting patient safety and enabling evidence-based practice through informatics. In: Measuring and Improving Health Care Quality: Towards Meaningful Solutions to Pressing Problems: Nursing's Contributions to the State of the Science; 2002 Apr 18–20; Philadelphia.
3. Chuang JH, Friedman C, Hripcsak G. A comparison of the Charlson comorbidities derived from medical language processing and administrative data. Proc AMIA Symp 2002: 160–4.

4. Bu D, Hripcsak G. Case-based reasoning for medical risk stratification: contrast associated nephropathy (poster). Proc AMIA Symp 2002: 986.
5. Stetson PD, Johnson SB, Scotch M, Hripcsak G. The sublanguage of cross-coverage. Proc AMIA Symp 2002: 742–6.
6. Scotch M, Hripcsak G. The Companion Project — a wearable computing device for the patient for reducing medical errors and improving patient care (poster). Proc AMIA Symp 2002: 1157.
7. Patel VL, Branch T, Arocha JF. Errors in interpreting quantities as procedures: the case of pharmaceutical labels. International Journal of Medical Informatics 2002;65:193–211.
8. Kubose T, Patel V. Dynamic adaptation to critical care medical environment: error recovery as cognitive activity. To appear in the Proceedings of Cognitive Science Society, August, 2002.
9. Tao Y, Wang D, Shortliffe EH, Lussier YA. Extended Attributes of Event Monitor Systems for Criteria-Based Notification Modalities. Proc AMIA Symp 2002;762–6.

10. Hripcsak G, Bakken S, Stetson PD, Patel VL. Mining complex clinical data for patient safety research: a framework for event discovery. J Biomed Inform 2003;36:120–30.
11. Cao H, Stetson P, Hripcsak G. Assessing explicit error reporting in the narrative electronic medical record using keyword searching (poster). Proc AMIA Symp 2003:803.
12. Cao H, Stetson P, Hripcsak G. Assessing explicit error reporting in the narrative electronic medical record using keyword searching. J Biomed Inform 2003;36:99–105.
13. Murff HJ, Patel VL, Hripcsak G, Bates DW. Detecting adverse events for patient safety research: a review of current methodologies. J Biomed Inform 2003;36:131–43.
14. Bates DW, Evans RS, Murff H, Stetson PD, Pizziferri L, Hripcsak G. Detecting Adverse Events Using Information Technology. J Am Med Inform Assoc 2003;10:115–28.
15. Bates DW, Evans RS, Murff H, Stetson PD, Pizziferri L, Hripcsak G. Policy and the Future of Adverse Event Detection Using Information Technology. J Am Med Inform Assoc 2003;10:226–8.
16. Keselman A, Patel VL, Graham MJ, Zhang J, Johnson T. Institutional decision making to select patient care devices: an analysis to identify threats to patient safety. Proc AMIA Symp 2003.
17. Sharda P, Das A, Patel V. Specifying Design Criteria for Electronic Medical Record Interface Using Cognitive Framework. Proc AMIA Symp 2003.

18. Hripcsak G, Stetson PD, Gordon P. Using the FCIM curricular guide and administrative codes to assess internal medicine resident breadth of experience. Acad Med 2004;79:557–63.
19. Yu A, Stetson PD, Hripcsak G. Detection of adverse events using conflicts in the electronic medical record (poster). Medinfo 2004:1923.
20. Cao H, Chiang MF, Cimino JJ, Friedman C, Hripcsak G. Automatic summarization of patient discharge summaries to create problem lists using medical language processing (poster). Medinfo 2004:1540.
21. Bakken S, Hripcsak G. An informatics infrastructure for patient safety and evidence-based practice in home healthcare. Journal of Healthcare Quality 2004;26:24–30.
22. Bakken S, Cimino JJ, Hripcsak G. Enabling patient safety and evidence-based practice through informatics. Medical Care 2004; 42:S49–56.

23. Bu D, Kaplan H, Rabin Fastman B, Hripcsak G. Case retrieval for medical event reporting systems: an evaluation of the HAWK system. Submitted for publication.
24. Melton GB, Hripcsak G. Patient safety: automated detection of adverse events using natural language processing in discharge summaries. Submitted for publication.

Other references (not due to this project)

25. Honigman B, Lee J, Rothschild J, et al. Using computerized data to identify adverse drug events in outpatients. J Am Med Inform Assoc 2001;8(3):254-66.
26. Field TS, Gurwitz JH, Harrold LR, et al. Strategies for Detecting Adverse Drug Events Among Older Persons in the Ambulatory Setting. J Am Med Inform Assoc 2004.
27. Murff HJ, Forster AJ, Peterson JF, Fiskio JM, Heiman HL, Bates DW. Electronically screening discharge summaries for adverse medical events. J Am Med Inform Assoc 2003;10(4):339-50.

Products

The system used to detect NYPORTS events in discharge summaries [24] is available to outside researchers. The system requires that the researcher obtain a license for the MedLEE natural language processing system, which is currently available only on a limited basis due to training and support issues. To address this limitation, we have submitted a P41 infrastructure grant to the NIH that, if funded, would make MedLEE freely available to researchers around the nation.