

CM-Extractor: An Application for Automating Medical Quality Measures Abstraction in a Hospital Setting

Mark L. Morsch¹, Joel L. Vengco², Ronald E. Sheffer, Jr.¹, Daniel T. Heinze, PhD¹

¹ A-Life Medical, Inc.
6055 Lusk Boulevard – Suite 200
San Diego, California 92121
mmorsch, rsheffer, dheinze@alifemmedical.com

² Partners HealthCare System Inc.
93 Worcester Street
Wellesley, Massachusetts 02481
jvengco@partners.org

Abstract

In the US, health care providers are required to report evidence-based quality measures to various governmental and independent regulatory agencies. Abstracting appropriate facts from a patient's medical record provides the data for these measures. Finding and maintaining qualified staff for this vital function is a challenge to many healthcare providers. Emerging systems and technologies in large-scale clinical repositories and AI techniques for information extraction have the potential to make the process of collecting measures more consistent, accurate and efficient. This paper presents CM-Extractor, a computerized system that automates the process of quality measures abstraction using natural language processing and a rule-based approach. An evaluation of a deployed system used for hospital inpatient cases is discussed. The results showed that the NLP performed with high accuracy across multiple types of medical documents, and users were able to significantly improve productivity. Challenges remain in the areas of availability of electronic patient data and a model for deploying and supporting solutions on a large scale.

Problem and Task Description

This paper discusses the implementation and evaluation of CM-Extractor, a computerized system that automates the medical data abstraction process. CM-Extractor integrates electronic sources from an enterprise-wide clinical repository with natural language processing (NLP) to extract information from free text and present the results via a flexible user interface to a human reviewer for verification. We consider the motivations for, approach to and results of an implementation in a hospital setting.

In 2002, the Joint Commission on Accreditation of Healthcare Organizations (JCAHO) began collecting a new set of evidence-based quality measures called core measures from accredited hospitals in the US. The latest evolution of JCAHO's ORYX initiative, the core measures, provides standard metrics of patient care for certain disease categories (JCAHO 2005). With detailed specifications for sampling, collecting and reporting, core measures provide a means for measuring performance across facilities. The initial set of core measures covered four medical conditions: acute myocardial infarction

(AMI), heart failure (HF), community acquired pneumonia (CAP) (pneumonia (PN) in the revised measure) and pregnancy. Each core measure consists of a set of pre-defined data elements that represent discrete treatment actions or historical information associated with an inpatient visit. For every case that meets the sampling criteria for one of the medical conditions, the values for the associated data elements must be abstracted from the medical record. Other organizations, such as the Center for Medicare and Medicaid Services (CMS), mandate the collection of evidence-based measures, so hospitals are motivated both ethically and financially to ensure that the abstraction process is dependable, auditable and accurate.

To comply with these new mandates, hospitals must extend current abstraction functions, or set up new ones, to collect, store and report the quality measures data. For the majority of facilities that rely on paper-based records, this is an additional manual process that requires patient records to be pulled from storage, routed to the appropriate resource for abstraction and then returned to storage. For these types of solutions, it can be challenging to maintain high levels of reliability and consistency over time due to the risk of human abstracting errors and administrative errors in document handling and data entry. With the increasing adoption of electronic health records, there is an opportunity to leverage the investment that facilities have in technology infrastructure and develop more accurate and efficient processes.

Abstracting medical documents is a difficult task for both humans and machines. There are inherent difficulties in the task regardless of how it is approached, and there are specific difficulties that primarily affect either humans or machines.

Inherent difficulties with medical abstracting relate to the guidelines for abstracting, the manner in which information is represented in medical documents and the options that are available for formal representation of the abstract.

Data definitions for abstracting, those for JCAHO core measures for example, tend to be lengthy because they frequently define criteria that are not simple values such as the patient's blood pressure. Common elements that

compose a definition include the patient’s primary medical problem, co-morbidities, aspects of treatment, medications, vital signs and laboratory measurements, event timing, etc. Elements of a definition may be either positive (given aspirin) or negative (not given aspirin) and include exclusion criteria such as contraindications (active bleeding at time of arrival). Because definitions are often complex, it is not uncommon for the information needed to satisfy the definition’s elements to be widely spread across a variety of locations within the patient’s chart and may require inference in addition to direct extraction. Also the information in the documentation may be incomplete. In this regard, for each element that must be satisfied by exception or for which there is no documented information (i.e. finding that something did not happen), a complete and time consuming review of the entire chart and results is required.

Due to the complexity of the abstraction task, performance errors (slips and mistakes) that primarily affect humans, and systematic errors (knowledge and rule-based errors) that affect both humans and machines, are of considerable concern (Wickens 1992). The evaluation presented here continues a series of evaluations testing the notion that it is possible to integrate human and machine performance in a way that emphasizes the strengths and minimizes the weaknesses of each (Morris et al. 2000; Morsch et al. 2004).

Application Description

The design of the CM-Extractor was guided by four principles:

- 1) Provide a single source for all relevant information;
- 2) Achieve an intuitive and logical workflow and an unobtrusive user interface;
- 3) Be scalable to other institutions and abstraction initiatives;
- 4) Increase abstractor efficiency and accuracy.

Great care was taken to understand abstraction workflow and the requisite clinical knowledge prior to implementing the CM-Extractor. The user interface and workflow are designed to be flexible enough to account for institutional and even department level differences. Clinical knowledge inherited from clinician designers as well as from JCAHO’s clinical guidelines references is incorporated at all operational levels of the system so that the NLP system and all abstractors across the hospital system work in harmony, ensuring that correct and appropriate concepts are abstracted.

As illustrated in Figure 1, CM-Extractor consists of a database and processing components operating within a centralized data center. It is delivered to the hospital as an Application Service Provider (ASP) tool. Information is exchanged between the hospital and the CM-Extractor data

center over a secure Internet connection. Hospitals submit electronic medical records (transcription notes, orders, results, etc.) to the data center for processing through an automated, scheduled process. These records contain both structured (e. g. lab results and medication orders) and unstructured data (e. g. free text transcriptions) and typically originate from multiple information system within a hospital.

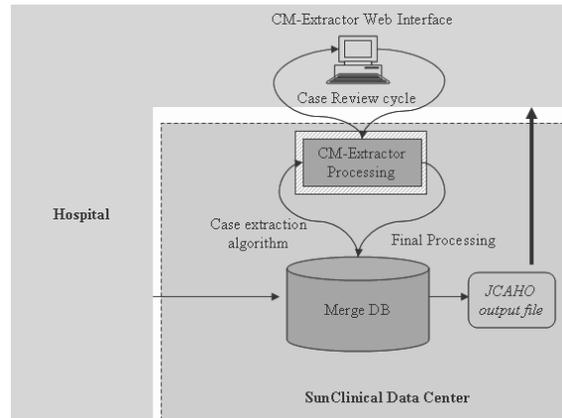


Figure 1: CM-Extractor Workflow Architecture

The merge database is a data warehouse that holds electronics records for all patient visits, referred to as cases, in a hospital. As part of the JCAHO core measures abstraction process, cases are selected and extracted from the merge database based on selection criteria defined by JCAHO. For example, to select an acute myocardial infarction (AMI) case the patient must be 18 years of age or older and have a principle ICD-9 diagnosis code from a list of specified codes, see Table 1 (JCAHO 2005). Cases that match the selection criteria are processed through the CM-Extractor system, where the NLP engine and expert rules process both the free text and the structured data, assigning values to the required data elements.

Code	ICD-9 Description
410.01	Anterolateral wall, AMI-initial episode
410.11	Other anterior wall, AMI-initial episode
410.21	Inferolateral wall, AMI-initial episode
410.31	Inferoposterior wall, AMI-initial episode
410.41	Other inferior wall, AMI-initial episode
410.51	Other lateral wall, AMI-initial episode
410.61	True posterior wall, AMI-initial episode
410.71	Subendocardial, AMI-initial episode
410.81	Other specified sites, AMI-initial episode
410.91	Unspecified site, AMI-initial episode

Table 1: Principle Diagnoses for AMI Cases

Note that automated ICD-9 coding was not part of this application, although it is within the capabilities of the NLP engine used. Hospital abstractors retrieve the

processed medical records for validation, quality assurance, approval and reporting through a single web-based user interface from any client workstation in the hospital. The finalized report for submission to JCAHO is loaded over the Internet from the data center to the hospital or directly to a JCAHO-approved reporting application.

Data warehousing plays a significant role in supporting the implementation of CM-Extractor. Giving the NLP engine access to electronic information is the most important prerequisite for achieving optimum results. A second important requirement is having hospital records integrated in a data warehouse. In native form, data from hospital transaction systems are not easily aggregated. Additionally, uniform codification and nomenclature is desirable in order to provide a standard way to view, report on and analyze data. For the implementation described here, SunClinical Data Institute created unified clinical and financial data warehouses by linking disparate systems within the hospital enterprise. This allowed for all of the electronic documents from each patient visit to be sent to the NLP engine at one time through an XML interface using a standardized representation of structured data and a

standard encapsulation of unstructured data. This level of integration facilitated the development and implementation process by producing very high quality data and eliminating the need for multi-system integration.

Another important aspect of the tool is its ability to collect and display a variety of information from multiple clinical, laboratory and administrative systems, providing the abstractor with a single source from which to abstract. Typically, abstractors pull information from paper charts, the billing system, the clinical order entry system, the electronic transcription system and so on. All of these systems are usually stand-alone applications or paper sources that the abstractor must open separately in order to acquire answers. Having all of the information on a single readable screen (Figure 2), as provided by CM-Extractor, decreases confusion and time spent searching multiple sources for information. The abstractor can quickly review and validate the information for submission. Measures of information certainty, as assessed by the system, direct the abstractor and provide an added level of streamlining and reliability.

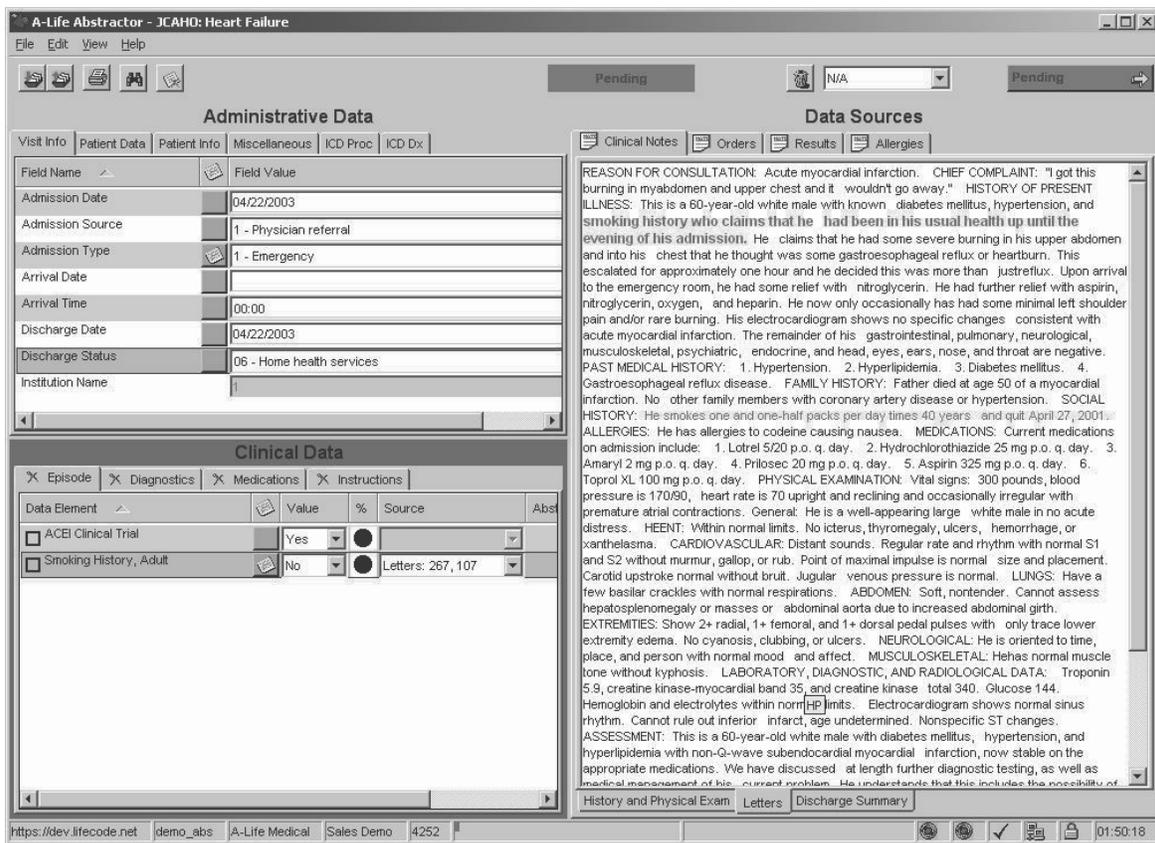


Figure 2: The CM-Extractor provides access to patient demographics (Administrative Data), all available electronic clinical documents, orders, and results (Data Sources), and results of the abstraction (Clinical Data). This example is not a real patient, but actual patient records were used in the implementation and evaluation.

Uses of AI Technology

A-Life Medical's LifeCode® NLP engine forms the basis for the CM-Extractor automated abstracting capability. A-Life's commercial products use sophisticated natural language processing to extract and normalize information from free text clinical documents. LifeCode's linguistic analysis relies on morphological, syntactic, semantic and pragmatic analysis to extract and synthesize complex concepts from free text medical documents (Heinze et al. 2001). The LifeCode technology is comparable to earlier work in the field, most notably (Sager et al. 1994), (Freidman et al. 1994) and (Freidman et al. 1999). These systems, along with LifeCode, have been able to achieve high accuracy in information extraction tasks on medical texts. The LifeCode engine was selected for CM-Extractor based on its proven accuracy and a demonstrated ability to scale to a production environment. The current primary application for the LifeCode engine is automated medical coding of diagnosis and procedure codes with over 300 sites in production.

To implement CM-Extractor, the NLP engine was enhanced in several important ways. We present the following that are most relevant in their use of AI technology.

As implemented in CM-Extractor, LifeCode has the capability to perform NLP across multiple documents of a full inpatient chart, spanning administration notes, history and physical exams, nurse notes, consultations, labs, orders, surgical notes and discharge summaries. This entails recognizing the document context as well as being able to associate information across documents. This is accomplished in a two step approach. The first step extracts and encodes discrete facts using multi-stage pattern analysis that transforms the free text from multiple documents into streams of multi-word tokens. These tokens are assigned meaning by matching against large knowledge bases of medical concepts using a vector-based calculation. The components of each vector are the individual terms within a token. Meanings are refined by merging or dividing the multi-word tokens governed by specific rules of syntax and semantics and then repeating the vector computation. The result is lists of facts categorized by original context, including the type of source document, which encode medical information at the phrase and sentence level. The pattern analysis, semantic analysis and vector computation are described in detail in (Heinze et al. 2001). The second step is a pragmatic analysis that merges the individual concepts using a rule-based approach to arrive at appropriate values for the desired data elements. These rules are designed to implement the inclusion and exclusion criteria from the data dictionary of the JCAHO core measures specification (JCAHO 2005). An example of these criteria is shown in Table 2 for the data element 'Adult Smoking History'. It

is interesting to note in this application that single data elements, such as 'Adult Smoking History', require information from multiple documents to be synthesized in order to assign appropriate values. This stems from the property of medical documents, where the type and extent of the information content is related to the purpose of the specific medical encounter being documented. For instance, a surgical note may just mention that a patient is a smoker while a history and physical exam note would typically have more detail on the type of tobacco used, pattern of usage and duration of use.

Inclusion Criteria
• positive smoker, type of product not identified
• positive tobacco use, type of product not identified
• history of cigarette use without mention of a time frame, if no indication that patient quit
• history of smoking (type of product not identified), without mention of a time frame, if no indication that patient quit
• history of smoking within one year prior to arrival, type of product not identified
• history of tobacco use (type of product not identified), without mention of a time frame, if no indication that patient quit
• history of tobacco use within one year prior to arrival, type of product not identified
• recent smoker
Exclusion Criteria
• chewing tobacco use only
• cigar smoking only
• cigarette smoking within one year prior to arrival or any of the other inclusion terms described using one of the following qualifiers: cannot exclude, cannot rule out, may have, may have had, may indicate, possible, suggestive of, suspect or suspicious
• illegal drug use only (e.g., marijuana)
• oral tobacco use only
• pipe smoking only
• remote smoker (smoked in the past, but greater than one year ago)

Table 2: Inclusion and Exclusion Criteria for 'Adult Smoking History' defined as cigarette smoking within one year prior to hospital arrival

As part of the pragmatic analysis, LifeCode employs a type of self-assessment to determine the engine's confidence in its own performance and to accommodate

partial, corroborating and contradictory information. In the user interface, extracted elements are marked with a visual indicator that directs the abstractor to areas where a review is recommended. For each extracted element value, the system assigns one of four levels of 'certainty': high, medium-high, medium and low. The levels are a qualitative judgment of how well the information extracted from the data sources matches the specified data element definition. The supporting rule-based architecture combines the assignment of value and certainty level into one rule per combination based on expert knowledge. So if a data element has two possible values, yes or no, then there are eight (two values x four certainty levels) potential rules covering every combination of value and certainty level. It is not required in this model that every combination of data element value and certainty level be defined; only those that are meaningful. For instance, for an adult patient to have smoking history with high certainty, statements affirming a smoking history must be present in the past medical history section of at least one of four types of notes (emergency department note, history and physical exam note, consultant note or discharge summary), and no contradictory statements should be present in any history section. Similar to (Wilcox and Hripcsak 2003), we have found expert medical knowledge, which underlies the majority of the pragmatic analysis, to be a significant factor in improving performance.

Because many of the core measures have time specific requirements, LifeCode provides extensive date and time handling capabilities. This includes both the capability to associate dates and times with specific information and to assess relative times and intervals. For example, for the data element 'Administration of Aspirin at Arrival' a determination must be made if aspirin was administered to the patient within 24 hours of arrival to the hospital. As part of its lexical analysis capabilities, the system is able to extract the time of administration and calculate the relative time between arrival and administration.

In addition to its free text processing capabilities, LifeCode also handles structured data, including the linguistically impoverished context of orders, in which a list notation indicates which labs or medications were ordered by whom and at what time. LifeCode categorizes medications based on free text and on sections and document types. These categories include the structured lab and medication lists, current medications, medications the patient was on prior to arrival, medications the patient was prescribed during a visit, medications the patient is allergic to, transfer medications and discharge medications. Handling the combination of free text and structured data from an electronic medical record has also been addressed by (Hazlehurst et al, 2005).

Application Use and Payoff

CM-Extractor was deployed and evaluated over a ten month period at Alamance Regional Medical Center (ARMC) in Burlington, NC. Actual patient cases were used during the deployment and evaluation. Alamance is a 350 bed facility with state of the art facilities, medical equipment and advanced information technology infrastructure.

Evaluations were held periodically through the deployment period as the system was further developed. During this deployment, the AMI and CAP core measures were the primary focus for development and evaluation. The goals of the evaluations were to measure the accuracy of the NLP engine's extraction and the productivity improvement of the abstractors. The CM-Extractor users were nurses with previous experience abstracting for clinical purposes who were trained in the use of the CM-Extractor user interface. They used the CM-Extractor user interface to view the results of the NLP engine and review the electronic sources, completing and, if needed, correcting the abstracted results. Portions of the medical records that were not available electronically were supplied as paper records. Both the original and edited data values were stored in the CM-Extractor database. The edited values assigned by the human reviewers were considered correct for the purposes of determining accuracy. The user interface collected productivity data using the 'edit time' feature of the systems that counts the total time that a case is open in the CM-Extractor user interface.

Detailed results for the final evaluation, which covered the CAP core measure, are described below. This evaluation included twenty cases. The electronic sources available included both structured and unstructured data. The structured data contained orders, results and allergies. The unstructured data contained admission notes, history and physical exams, progress notes, consultant notes and discharge summaries. The cases selected for evaluation were pulled from the hospital's data warehouse using the selection criteria based on diagnosis coding as specified by JCAHO. The selected evaluation cases had not been used in developing or testing the system.

For evaluation purposes, the hospital utilized a secure Internet-based deployment of CM-Extractor that interfaced with the SunClinical Data Warehouse. Each user was assigned a unique user name and password to access the system. All data links were encrypted, and the system maintained a complete audit trail of user access to records and changes to extracted data items. An automated lockout feature prevented the unintended viewing of patient data if the user interface was left unattended.

Based on the productivity and accuracy data collected and interviews with the reviewers, CM-Extractor offers significant benefits. The productivity increase per case,

as compared to a manual paper-based process was over 3 times for AMI cases (from over 30 minutes to 10 minutes per case) and 1.5 times for CAP cases (from about 15 minutes to less than 10 minutes per case). Note, the time per case for the manual paper-based process was based on interviews with the nurse abstractors. This time savings does not take into account the reduction in effort to pull, route and re-file paper records. Additional productivity increases should be obtained with more electronic documentation. At the time of the deployment, Alamance had more than 60% of their medical records in electronic format.

The simplicity of the user interface was a major factor in the enhanced productivity. Abstractors learned to use two interfaces: one for selecting and retrieving cases, and a second to verify and complete the extracted items and view the electronic sources. An interactive search feature of the electronic sources helped abstractors find specific information quickly. Also, once the reviewer completed a case, one button click would mark the completed case ready for submission and move to the next case.

The accuracy of the NLP engine was measured using a single statistic, similar to precision, calculated as (correct values / possible values). Possible values are defined as the number of data elements where the NLP engine assigned a value. This statistic reflected a goal of the project to emphasize correctness over completeness. Because of the importance of accurately reporting quality measures, too many incorrect values were not acceptable to hospital users and eroded confidence in the system. Throughout all evaluations, NLP accuracy consistently exceeded 90%. Table 3 shows the detailed result for the data elements comprising the final CAP evaluation; in this case the overall accuracy was 97.8% which was the best score measured during the deployment period.

Data Element	# Correct	# Possible
Arterial Blood Gas Done	14	14
Antibiotic Received	15	15
Blood Culture Collected after Arrival	9	9
Comfort Measures Only	4	4
Pneumonia Working Diagnosis on Admission	11	12
Pulse Oximetry Done	9	9
Smoking Counseling, Adult	11	12
Smoking History, Adult	15	15
Total	88	90

Table 3: Results for the Community Acquired Pneumonia core measure evaluation

The human abstractors also used paper sources while using CM-Extractor, so these statistics do not penalize the system for information found external to electronic documentation. Interestingly, human abstractors found an additional 29 data elements from paper sources in this

evaluation. With this high accuracy, abstractors indicated that they quickly gained confidence in the NLP results.

Key factors in the decision to implement a solution like CM-Extractor are the financial cost and benefit. Complying with reporting requirements for JCAHO and other agencies is a non-reimbursed cost for hospitals. CM-Extractor offers labor savings for abstractors, particularly at facilities where personnel can be redeployed in roles that are more valuable than reporting. In addition to reducing labor, CM-Extractor's consistent and compliant application of the JCAHO guidelines proves beneficial in terms of the quality of reporting. Because CM-Extractor can be used to produce reporting data in near real-time, hospitals can be proactive in addressing quality of care issues. Additionally, standardizing the abstraction process with the CM-Extractor ensures that regulatory reports generated for the given hospital are accurate, representative and optimal. The CM-Extractor is also accessible by various parties within the hospital, resulting in a collaborative environment. Work efficiency and accuracy of reporting are benefited by messaging, comment functionality, case status information and flagging in the CM-Extractor, all of which provide a means for abstractors and managers to work together more effectively.

Application Development

The CM-Extractor application was developed initially over a nine month period, prior to deployment at ARMC and then was further refined during deployment. The development team consisted of two software engineers, two NLP developers, a product manager and several part-time medical domain experts (primarily medical students) that filled the roles of analysts and testers. Once the system was deployed at ARMC, further development was done collaboratively with the users who had significant clinical and abstracting expertise.

The system was developed as a multi-tier Web-based application. The goal was to make the system easy to use and deploy and capable of running on existing PCs within the hospital. The user interface was implemented in Visual Basic.Net and communicated to a centralized SQL Server database over an SSL encrypted link. The NLP engine and other processing components were all implemented in C++. The starting point for NLP development was the emergency medicine coding version of the LifeCode engine. The major enhancements were support for multiple types of medical documents, replacement of the medical coding expert rules with those necessary for quality measures reporting and extension of the medical knowledge bases to cover the required data elements for JCAHO core measures. The medical knowledge bases were written in a specialized language and are compiled directly into C++ data structures,

allowing the dictionaries to be extended without recompiling the core engine.

Deployment and Maintenance

The most significant challenge faced in this deployment was accessing and integrating electronic data sources. During the deployment, the complete patient record was not available electronically, so human abstractors needed to use some paper sources. Even with this limitation, the results were very good. A potential challenge for some facilities would be the quality of the medical documentation. The documentation at ARMC was of very high quality and consistency, and minimal work was needed to adapt the CM-Extractor. In facilities with lower quality documentation, an additional level of implementation effort on the NLP engine could compensate for some aspects of documentation quality.

Once CM-Extractor was deployed, the ASP delivery model and a modular design supported efficient maintenance. The NLP engine and other processing components were installed in a secure, centralized data center. Only the Web-based user interface was installed at the hospital. This allowed the back-end components to be updated with minimal disruption of on-line users. Updates that did require system down time, such as Web server software upgrades, database schema changes or hardware maintenance, were scheduled in advance with users. The user interface application itself was designed to automatically check for updates and start the install process if a new version was available. Site-specific data import and export modules were available, supporting the specific interfacing requirements of different hospital information systems. As was mentioned above, the modularity of design extends to the NLP engine where the knowledge bases, which hold the medical dictionaries and related ontologies, can be modified and deployed without changing core algorithms.

Deployment ended for this facility when it was determined that there was not enough demand or readiness in the general marketplace for this type of solution. So while a technical success and well received by users, most hospitals' technology infrastructure does not seem to be ready to support an application like CM-Extractor. However, with the trend toward computerized patient records and the massive amount of free text in medical records, the authors believe that AI tools for automated abstraction, like CM-Extractor, will be further developed and deployed within the next few years.

Conclusion

Automated medical coding solutions that utilize NLP technology have achieved commercial acceptance, particularly for physician billing applications (AHIMA 2004). However, NLP solutions are not in significant use

for clinical applications, such as quality reporting, outcomes analysis and pharmaceutical research. CM-Extractor is one of the first applications to help solve one on the biggest challenges facing hospitals today, the need for detailed, timely and accurate clinical abstractions of medical records.

The evaluation of the CM-Extractor product demonstrated several benefits to a hospital in meeting the increasing demands for abstraction. The productivity of the human abstractors increased significantly, NLP results were highly accurate and the user interface was found to be efficient and easy-to-use. Possibly as important as the measurable benefits was the introduction of a standardized workflow making the process more reliable and repeatable over the long term. These positive results are balanced against the challenges of developing a cost-effective business model for delivering this type of solution and integration issues that limit the availability of electronic sources. Wider adoption of electronic health records, standard interchange formats and a broader offering that supports multiple abstracting and coding tasks within a hospital all seem to be ingredients of a future solution.

Acknowledgements

The authors gratefully acknowledge the staff of Alamance Regional Medical Center, particularly Carol Hudson, Aaron Rose and Edith Apple for their invaluable services in the development and evaluation of CM-Extractor.

References

- AHIMA e-HIM™ Work Group on Computer-Assisted Coding. 2004. Delving into Computer-assisted Coding (AHIMA Practice Brief). *Journal of the American Health Information Management Association* Nov/Dec; 75(10):48A-H.
- Friedman, C.; Alderson, P.; Austin, J.; Cimino, J.; and Johnson, S. 1994. A General Natural Language Text Processor for Clinical Radiology. *Journal of the American Medical Informatics Association* 1(2):161-174.
- Friedman, C.; Knirsch, C.; Shagina, L.; and Hripcsak, G. 1999. Automating a Severity Score Guideline for Community-Acquired Pneumonia Employing Medical Language Processing of Discharge Summaries. *Proceedings of the AMIA Annual Symposium*, 256-260. Washington, DC.
- Hazlehurst, B.; Frost, H. R.; Sittig, D.; and Stevens, V. 2005. MediClass: A System for Detecting and Classifying Encounter-based Clinical Events in Any Electronic Medical Record. *Journal of the American Medical Informatics Association* 12(5):517-529.

Heinze, D.; Morsch, M.; Sheffer, R.; Jimmink, M.; Jennings, M.; Morris, W.; and Morsch, A. 2001. LifeCode: A Deployed Application for Automated Medical Coding. *AI Magazine* 22(2): 76-88.

Joint Commission on Accreditation of Healthcare Organizations. 2005. Specification Manual for National Implementation of Hospital Core Measures. Available at: http://www.jcaho.org/pms/core+measures/aligned_manua1.htm. Accessed January 11, 2006.

Morris, W.; Heinze, D.; Warner Jr., H.; Morsch, A.; Sheffer, R.; Jennings, M.; Morsch, M.; and Jimmink, M. 2000. Assessing the accuracy of an automated coding system in emergency medicine. *Proceedings of the AMIA Annual Symposium*, 595-599. Los Angeles, CA.

Morsch, M.; Byrd, D.; and Heinze, D. 2004. Factors in Deploying Automated Tools for Clinical Abstraction and Coding. *Proceedings of Information Technology in Health Care (CD-ROM Proceedings)*. Portland, OR.

Sager, N.; Lyman, M.; Bucknall, C.; Nhan, N.; and Tick, L. 1994. Natural Language Processing and the Representation of Clinical Data. *Journal of the American Medical Informatics Association* 1(2):142-160.

Wickens, C. D. 1992. *Engineering psychology and human performance*. New York: HarperCollins.

Wilcox, A., and Hripcsak, G. 2003. The Role of Domain Knowledge in Automating Medical Text Report Classification. *Journal of the American Medical Informatics Association*. 10(4):330-338.