

Active Semantic Electronic Medical Record

A. Sheth¹, S. Agrawal², J. Lathem¹, N. Oldham², H. Wingate², P. Yadav²,
and K. Gallagher²

¹ LSDIS Lab, University of Georgia
Athens, Georgia 30602
{amit, lathem}@cs.uga.edu

² Athens Heart Center
Athens, Georgia 30606
{subodh, noldham, ppyadav, kgallagher}@athensheartcenter.com

Abstract. The healthcare industry is rapidly advancing towards the widespread use of electronic medical records systems to manage the increasingly large amount of patient data and reduce medical errors. In addition to patient data there is a large amount of data describing procedures, treatments, diagnoses, drugs, insurance plans, coverage, formularies and the relationships between these data sets. While practices have benefited from the use of EMRs, infusing these essential programs with rich domain knowledge and rules can greatly enhance their performance and ability to support clinical decisions. Active Semantic Electronic Medical Record (ASEMR) application discussed here uses Semantic Web technologies to reduce medical errors, improve physician efficiency with accurate completion of patient charts, improve patient safety and satisfaction in medical practice, and improve billing due to more accurate coding. This results in practice efficiency and growth by enabling physicians to see more patients with improved care. ASEMR has been deployed and in daily use for managing all patient records at the Athens Heart Center since December 2005. This showcases an application of Semantic Web in health care, especially small clinics.

Keywords: Ontology, Rules, Electronic Medical Record (EMR), Electronic Health Record, Clinical Decision Support, RDQL, Web Services, Semantic Annotation, Active Semantic Document.

1 Introduction

The most cumbersome aspect of health care is the extensive documentation which is legally required for each patient. For these reasons, physicians and their assistants spend about 30% of their time documenting encounters. Paper charts are slowly being phased out due to inconvenience, inability to mine data, costs and safety concerns. Many practices are now investing in electronic medical records (EMR) systems which allow them to have all patient data at their fingertips. Although current adoption by medical groups (based on a 2005 survey [1]) is still below 15% with even less adoption rate for smaller practices, the trend is clearly towards increasing adoption. This trend will accelerate as regulatory pressures such as “Pay-4-Performance” become mandatory thus enhancing the ROI sophisticated systems can

achieve. This paper focuses on the first known development and deployment¹ of a comprehensive EMR system that utilizes semantic Web and Web service/process technologies. It is based on substantial collaboration between practicing physicians (Dr. Agrawal is a cardiologists and a fellow of the American Cardiology Association, Dr. Wingate is an emergency room physician) at the Athens Heart Center and the LSDIS lab at UGA. More specifically, we leverage the concept and technology of Active Semantic Documents (ASDs) developed at the LSDIS lab. ASDs get their *semantic* feature by automatic semantic annotation of documents with respect to one or more ontologies. These documents are termed *active* since they support automatic and dynamic validation and decision making on the content of the document by applying contextually relevant rules to components of the documents. This is accomplished by executing rules on semantic annotations and relationships that span across ontologies.

Specifically, Active Semantic Electronic Medical Record (ASEMR) is an application of ASDs in health care which aims to reduce medical errors, improve physician efficiency, improve patient safety and satisfaction in medical practice, improve quality of billing records leading be better payment, and make it easier to capture and analyze health outcome measures. In ASMER, rules specified in conjunction with ontologies play a key role. Examples of the rules include prevention of drug interaction (i.e., not allowing a patient to be prescribed two severely interacting drugs, or alerting the doctor and requiring his/her to make specific exceptions when low or moderate degree of interactions are acceptable) or ensuring the procedure performed has supporting diagnoses. ASDs display the semantic (for entities defined in the ontologies) and lexical (for terms and phrases that are part of specialist lexicon, specific items related to the clinics, and other relevant parts of speech) annotations in document displaced in a browser, show results of rule execution, and provide the ability to modify semantic and lexical components of its content in an ontology-supported and otherwise constrained manner such as through lists, bags of terms, specialized reference sources, or a thesaurus or lexical reference system such as WordNet [11]. This feature allows for better and more efficient patient care and because of the ability of ASDs to offer suggestions when rules are broken or exceptions made.

ASEMR is currently in daily and routine use by the Athens Heart Center (AHC) and eight other sites in Georgia. ASEMRs have been implemented as an enhancement of AHC's Panacea electronic medical management system. Panacea is a web-based, end-to-end medical records and management system, and hence it is used with respect to each patient seen at AHC. This has enhanced the collaborative environment and has provided insights into the components of electronic medical records and the kinds of data available in these systems. The preliminary version was implemented during Summer 2005 and tested in early fall. The current version was deployed and has been fully functional since January 2006. Parts of ASMER we will focus on in this paper are:

- the development of populated ontologies in the healthcare (specifically cardiology) domain

¹ Preliminary deployment in September 2005, full deployment in January 2006.

- the development of an annotation tool that utilizes the developed ontologies for annotation of patient records
- the development of decision support algorithms that support rule and ontology based checking/validation and evaluation.

The remainder of this paper is organized as follows. Section 2 makes a case for semantics through a motivating scenario (for brevity, only one example is given). Section 3 describes the knowledge and rules representation. The application is detailed in Sections 4 and the implementation details are given in Section 5. Section 6 evaluates the approach and provides statistics which support the growth of the practice since the use of the EMR. Section 7 lists related work and Section 8 concludes with future work.

2 Motivating Scenario and Benefits

In addition to the complexity of today's healthcare, medical practitioners face a number of challenges in managing their practices. One of the challenges is the need to improve the quality of care, adhere to evolving clinical care pathways, reduce waste and reduce errors (with associated need to develop and report quality of care measures). Another challenge is that of medical billing. Let's investigate the latter further. Each insurance company follows Local Medical Review Policy (LMRP) which are policies specifying which diagnosis justify the medical necessity of a procedure. If the appropriate codes are not given in accordance with these LMRPs, the insurance will not pay for the charge. Because of these rigid requirements many claims are rejected and the amount of time for receiving a payment is prolonged and in many cases the physicians are not reimbursed for their services. If correct coding compliance is enforced by the system at the point of charge entry on the superbill (the bill of all charges and diagnoses for a visit) the problem of procedures without supporting diagnosis codes is eliminated. Table 1 contains a partial list of ICD9CM codes that support medical necessity for CPT 93000 EKG which were taken from the Centers for Medicare and Medicaid Services [3]².

The primary diagnosis code selected for the EKG must be one of the supporting diagnosis codes listed above. There are additional complex rules such as certain ICD9CM codes should not be selected together and certain procedures should not be billed for in the same claim. In section 4.2, we will present our approach which uses a combination of OWL ontologies and rules to validate data in a superbill to ensure coding compliance by presenting the appropriate subset of linking diagnosis codes when a procedure is selected. Due to the creation of more accurate and compliant claims, this approach has the potential to eliminate coding errors which would result in improved financials.

² ICD9-CM stands for "The International Classification of Diseases, 9th Revision, Clinical Modification"; these codes are used to denote the diagnosis. CPT (Current Procedural Terminology) codes are used to denote treatments. Payment is done based on the treatment, but the bill must contain acceptable diagnosis for that treatment.

Table 1. Medical Necessity for EKG

ICD9CM	Diagnosis Name
244.9	HYPOTHYROIDISM
250.00	DIABETES MELLITUS TYPE II
250.01	DIABETES MELLITUS TYPE I
242.9	HYPERTHYROIDISM
272.2	MIXED HYPERLIPIDEMIA
414.01	CAD-NATIVE
780.2-780.4	SYNCOPE AND COLLAPSE - DIZZINESS AND GIDDINESS
780.79	OTHER MALAISE AND FATIGUE
785.0-785.3	TACHYCARDIA UNSPECIFIED - OTHER ABNORMAL HEART SOUNDS
786.50-786.51	UNSPECIFIED CHEST PAIN - PRECORDIAL PAIN
786.59	OTHER CHEST PAIN

In addition to greater facilitation of billing process, physicians benefit from the clinical decision support that can be provided by a system which has rich domain understanding through the use of ontologies and rules. Patients benefit as well as this ability allows better patient care, increased safety and satisfaction. Checks such as preferred drug recommendations lead to prescription drug savings for patients leading to improved satisfaction. The most important benefit we seek from ASEMRS with its proactive semantic annotations and rule-based evaluation is the reduction of medical errors that could occur as an oversight. Ultimately the proof of these support features will be manifest by improved outcome data for example better Medpar scores (medicare beneficiary morbidity and mortality data) for Physicians.

3 Knowledge and Rules Representation

We employ a combination of OWL [6] ontologies with RDQL[9] rules in order to supply the document with rich domain knowledge. The rules provide additional domain knowledge and compensate for the limitations of the OWL language.2.1 Ontologies. A more complex rule specification (and corresponding rule processing) capabilities may be needed in future, but for our current purpose this was more than adequate and this choice also provided efficient implementation alternative.

We utilize three ontologies to represent aspects of the domain. The practice ontology contains concepts which represent the medical practice such as facility, physician, physician assistant, and nurse. The infrastructure of the medical practice is given by the concepts and relationships. The practice ontology was created in conjunction with experts in the medical field. Parts of our own databases were the source for populating this ontology.

The Drug ontology contains all of the drugs and classes of drugs, drug interactions, drug allergies, and formularies. Capturing such information reduces medical errors and increases patient safety. Furthermore, prescribing drugs from the

formularies of the patient's insurance plans improves patient satisfaction. License content (Gold Standard Media) equivalent to physician's drug reference was the primary source for populating this ontology which is shown, in part, in figure 1.




Fig. 1. Partial View of Drug Ontology

The Diagnosis/Procedure ontology includes concepts such as medical conditions, treatments, diagnoses (ICD-9), and procedures (CPT). Licensed SNOMED (Systematized Nomenclature of Medicine--Clinical Terms) [8] content is used for populating this ontology. A key enhancement involved linking this ontology to the drug ontology. This allows powerful decision support by giving the system specialized domain knowledge. We will use this representation to enable the system to suggest treatments and drugs based on the patient's condition or diagnosis. User or user group specific frequently used codes lists are supported by this ontology. This allows customizability such that each area of the practice will be given procedures and diagnosis codes which frequently apply to their area.

For example, procedures such as Dipiridamol injections and Muga scans are generally administered in the area of Nuclear medicine and should therefore the remainder of the clinical staff should not be bothered with those procedures cluttering their view. Each area has customizable frequent lists such as Nuclear, Pacemaker Evaluation, Echocardiograph, etc.

Medical records of patients are automatically annotated using the ontologies listed above and are displayed in a browser. Drugs, allergies, physicians and facilities (e.g., physicians or facilities the patient is referred to), treatments, diagnosis, etc. are automatically annotated. The physician has the ability to pull up a contextual list or even a visual subset of the relevant ontology and pick alternative choices. In some

cases, alternatives are provided in ranked order list (e.g., other physicians with the same specialty in the same area and accepting the same insurance as the patient).

3.1 Rules

ASEMRs support active features by executing relevant rules over semantic annotations to support the following initial sets of capabilities:

- drug-drug interaction check,
- drug formulary check (e.g., whether the drug is covered by the insurance company of the patient, and if not what the alternative drugs in the same class of drug are),
- drug dosage range check,
- drug-allergy interaction check,
- ICD-9 annotations choice for the physician to validate and choose the best possible code for the treatment type, and
- preferred drug recommendation based on drug and patient insurance information

The benefits of combining the use of ontologies and rules are two-fold. First, the rules allow the system to make decisions. Second, using rules the system can become declarative to the extent that additional relationships and facts can be added at any time without changing the code. For example, if the relationship “cancels_the_effect” is added to the ontology coupled with a rule indicating which drug or combinations of drugs cancel the effect of drugX, then the capability of the system is enhanced without any code modifications. This allows for a great deal of extensibility and flexibility such that one could even define classes of drugs, such as blood thinners, which cancels the effects of other classes of drugs. **Rules allow for more flexibility, enhanced reasoning power and extensibility.**

4 Application

The following section details two components which utilize semantic web technologies and are currently deployed and in use by at least eight beta sites. The evaluation section contains an analysis of the effect of this semantic health record application on one practice.

4.1 Active Semantic Documents

Physicians are required to thoroughly document each patient encounter. Reports usually contain a problem list, family history, history of present illness, review of symptoms, impressions and plans. Data acquisition and data entry is a painstaking process which usually results in late hours for the physician. One alternative is dictation. While dictation maybe faster for the physician, it has many negative drawbacks including lack of structured data for analysis and mistakes in transcription that have to be corrected. It is clear from our experience that a better solution is an application which “understands the domain” thus facilitates the structured entry of

data by offering relevant suggestions in a customizable point and click interface 1 generating complete and coherent reports. The Active Semantic Documents (ASD) EMR both expedites and enhances the patient documentation process. The support and speed provided by them enables physicians and physician assistants to complete all of their patient documentation while the patient is still in the room allowing the physician to provide better care with a greater volume of patients.

The annotation view pictured in figure 2 is an Active Semantic Document. The annotations facilitate the creation of the document by performing annotations of ICD9s, words and sentences within the report, and drugs. Three drug related annotations can be seen in figure 2 under the “Current Medications” section. The drug Coumadin has a level three interaction warning. Holding the cursor over this warning displays the name of the drug with which it interacts. The yellow *F* annotation warns that the drug is not covered under the patient’s insurance formulary. The annotations can also be extended to also semantically enhance the monograph. The green *A* annotation warns that the patient is allergic to this drug. Clicking on the Explore button allows the user to write a prescription for this drug, change quantities, or view the monograph for this drug. Exploring the drug allows for semantic browsing, querying for such details as how many patients are using this class of drug, and for performing decision support. Figure 3 shows the exploration of the drug Tasmar.

Edwards, Christopher MR # 111111 DOB: 2/12/1975 Age: 31 Sex: M Phy: Wingate, Harry Ins: BCBS - PPO Notes Recent Docs

Problem List Medications Encounter Sheet Soap Note Annotation View Print Report Visit On 5/5/2006

8. Shortness of breath (86.00) E

A. Really no breath C

9. Nonorganic Sleep Disorder (307.4) E

10. 2005-Dr.Gibson treated for mild hypertension (401.1) C

11. Cad - Native (414.01) C

12. Coronary Atherosclerotic Vessel (414.00) E

Chief Complaint: Evaluation of angina and chest pain.

History of Present Illness: Mr. Edwards is a 31 year old patient of Dr. Subodh Agrawal. His last office visit was on 4/13/2006. Since his last visit new problems have developed. He is taking his medications as prescribed. The patient reports chronic moderate cramping chest pain located in the epigastric area. The chest pain is associated with itchy. He was evaluated at Emory emergency room. He was treated with CABG. The patient was then transferred to St. Joseph. He underwent an echocardiogram. He denies chest pain, edema, fatigue, and palpitations.

Current Medications	Medications After Visit
	Coumadin tablets 1 mg, 1 tab [3] F E
	Penicillin tablets 250 mg, 1 tab A F E
	Viagra 50 mg, 1 tab [3] E

Pharmacy: CVS - Lawrenceville Phone: 770-963-8183

Allergies: PENICILLIN, TYLENOL

Fig. 2. An application of Active Semantic Documents

4.2 Coding of Impressions

Section 2 described a scenario in which the complexity of medical billing is remedied by enforcing correct coding at the point of data entry by the nurse, physician, or assistant. As a patient is seen, orders and diagnoses are marked by the healthcare

provider on an ‘encounter sheet’ or ‘superbill’. It is imperative at this time that a diagnosis which supports medical necessity for a procedure be given in order to facilitate the billing process. This application employs a novel semantic approach for entering charges into the encounter sheet based on domain knowledge taken from the procedure and diagnosis ontology. This application allows for diagnoses to be taken directly from the documentation described in the previous section. Furthermore, when orders are placed the subset of diagnoses codes which are defined to support medical necessity for that order are shown. This method ensures that the charges will be entered correctly at the very beginning of the process. The semantic encounter sheet is shown in figure 4. As users select orders from the right column, the left column automatically populates with the linking diagnosis codes which support medical necessity. The doctor is required to validate this choice, and ontology enables him/her to easily consider alternatives.




Fig. 3. Exploration of the neighborhood of the drug Tasmar




Fig. 4. Semantic Encounter Sheet

5 Implementation Details

The Panacea database holds all information about a patient and the patient's visits including the patient demographics, medications before the visit, medications added during the visit, past and present problems, diagnoses, treatment, doctors seen, insurance information, and a text description of the visit. The method of data entry and data storage ensures that it is well structured and can trivially be converted into a single large XML. It is important to note that the text description is not simply stored as one large string but as a tree structure which can be lexically annotated far faster and with better accuracy compared with using natural language processing. A detailed discussion of this is out of the scope of this paper.

After the XML is created annotations must be applied in order to assert the rules. Since the structure and schema of the XML is known a priori, annotation is simply performed by adding metadata to the correct tags. The correct tags are identified using XPath. This approach has a much higher accuracy than most types of semantic annotation techniques. This is a result of knowing the structure of the XML prior to the annotation.

The module that creates the XML and the module that annotates the XML are separate entities on different servers and implemented in different languages. This was necessary as the legacy code is in ASP and most wide spread tools for XML and ontology querying are written in Java. The two modules communicate by passing the XML from the ASP to the Java server via a REST based web service. The addition of Web 2.0 technologies such as REST services allows much of the requests to generate from the client instead of the server. This gives the application the ability to mask latency and allow easy integration in to client side scripting. This solution offers much more than fixing the heterogeneity created by the two languages. This solution also offers scalability and extensibility. Allowing the memory and IO intensive ontology querying to be done independently of the application server frees up resources which may be used elsewhere.

After annotation a third module applies rules to the annotations. The rules used are written in RDQL. A rule either checks for the existence of an edge or its absence. For example, an 'interaction' relationship should not exist between two drugs or there should be a relationship, 'covered', between a drug and patient's insurance. When these rules are broken metadata is added to the previously added annotations in the form of properties. Once all of the annotations have been applied and the rules are asserted, the annotated XML makes its way back to the client where an XSLT is applied. The XSLT turns the XML into HTML which can be made interactive and presented to the user for review and edits.

Currently Panacea annotates doctors, problems, diagnosis, drugs, and patient demographics semantically. The rest of the document is annotated lexically. Queries that could be ran against these annotation include but are not limited to:

- drug-drug interaction check,
- drug formulary check (e.g., whether the drug is covered by the insurance company of the patient, and if not what the alternative drugs in the same class of drug are),
- drug dosage range check,

- drug-allergy interaction check,
- ICD-9 annotations choice for the physician to validate and choose the best possible code for the treatment type, and
- preferred drug recommendation based on drug and patient insurance information

Figure 5 depicts the architecture of the Active Semantic Document component of Panacea.




Fig. 5. ASEM Architecture

6 Deployment and Evaluation

At AHC, the main site of deployment, the application accommodates between 78 and 80 patient encounters per day, most of which are seen within a four hour time frame. The AHC, with two physicians, two to four mid-level providers, eight nurses, and four nuclear and echo technicians, relies on Panacea/ASEM for fully Web-based paperless operations for all functions except for billing (which is currently under development). The semantically annotated document creation in conjunction with workflow solutions such as patient tracking has allowed the AHC to operate in ‘real time’ mode such that the physicians and their assistants are able to complete all documentation for the patient’s visit during the encounter. Prior to deploying ASEM, majority of charts were completed in Panacea after patient hours, often requiring mid-level providers to complete them over the weekend.

As a result of Panacea deployment first, followed by its ASEM extension, the AHC has greatly increased the volume of patients which they are able to care for, and

importantly, without increasing its clinical staff. Figure 6 shows the growth of the AHC since March of 2004. This data was obtained by querying the database for the number of appointments scheduled. The development of Panacea began in the year 2002 and the ASEMR was deployed in December 2005; it became fully operational in January 2006. In other words, data prior to December 2005 reflects pre-semantic situation (as Panacea did not have any semantic/ontological/rule support, and the data after January 2006 reflect situation after deploying the semantic technology. The number of clinical staff members and facility remained relatively consistent throughout the entire sample period. The AHC saw growth in 2005 as they scheduled around 1000-1200 patients per month. The patient volume for the year 2006 has started at a consistent growth rate of 25-30%, with March peaking around 1400 appointments scheduled per month. Even with this increase in patient volume, the physician assistants are able to accompany the physicians to the hospital immediately after clinic hours instead of charting until late evening hours. Before the deployment of the new annotation view supported in ASEMR the mid-level providers remained in the office an additional 4-5 hours charting after the clinic closed. Main reason for the work remaining after clinical hours related to the need to insure consistency, completeness and correctness of the patient record (e.g., the CPT and ICD9 codes that form parts of billing information captured as part of coding of impressions). Since ASEMR addressed these issues through semantics and rules. Since the time we completed the training of clinical staff, all charts are completed before the clinic closes, and in most cases a chart is completed while the patient is still in the office.




Fig. 6. Athens Heart Center Practice Growth

Even with this increase in patient volume, the physician assistants are able to accompany the physicians to the hospital immediately after clinic hours instead of charting until late evening hours. Before the deployment of the new annotation view supported in ASEMR the mid-level providers remained in the office an additional 4-5 hours charting after the clinic closed. Figures 7 and 8 show the dramatic change in the number of charts completed on the same day versus the number of charts backlogged at the end of the day for pre-deployment and post-deployment months respectively.




Fig. 7. Chart Completion before the preliminary deployment of the ASMER




Fig. 8. Chart Completion after the preliminary deployment of the ASMER

We have observed improvement in the patient satisfaction such as through the use formulary check as this could reduce patient costs through the check for medication with lower co-payments and insurance coverage, and the benefits associated with the use coding impression on improved billing as the basis of improved medical records and billing. Our next challenge is to measure these improvements and benefits quantitatively as part of an effort to develop and share return on investment (ROI) measures. As an aside, this work has in part enabled us to be an active member of

W3C's interest Group on Semantic Web for Health Care and Life Sciences, and provide the perspective of semantic Web applications and deployments in health care arena with a focus on smaller practices [10].

Given that this work was done in a live, operational environment, it is nearly impossible to evaluate this system in a "clean room" fashion, with completely controlled environment – no doctors' office has resources or inclination to subject to such an intrusive, controlled and multistage trial. Evaluation of an operational system also presents many complexities, such as perturbations due to change in medical personnel and associated training. In this context, we believe we have been able to present convincing evaluation of the benefits of a semantic technology.

7 Related Work

Some other healthcare applications have benefited from the use of ontologies. Chen et al. have experimented with step-wise automation of clinical pathways for each patient, in particular, according to the patient's personal health condition at the time of consultation in [4]. Their approach uses ontologies and web services; however, this approach does not propose the use of rules to supplement domain knowledge to compensate for the limitations of OWL. BioDASH [2] is a Semantic Web prototype of a Drug Development Dashboard that associates disease, compounds, drug progression stages, molecular biology, and pathway knowledge for a team of users. This work mentions use of rule-based processing using off-the-shelf RDF inference engines, and the use of rules to filter and merge data. Kashyap et al present a semantics-based approach to automate structured clinical documentation based on a description logics (DL) system for ontology management in [5]. This paper describes the use of document and domain ontologies. Onto-Med Research Group has designed Onto-Builder [7], a tool designed to support the construction and administration of Data Dictionaries in the field of clinical trials. This standard Data Dictionary is then used in the collection and analysis of clinical trials data. Quality assurance in carrying out clinical trials and uniformity are some benefits to such ontology.

We also note that as a "SW in Use" track paper, we focus on discussing a deployed system demonstrating the use of semantic web, rather than attempt to distinguish research contributions with respect to years of research in AI and decision support in healthcare, some of which took much longer to mature and find operational use than the new newer technologies. The newer technologies encompassing Semantic Web, SOA and Web 2.0 offer many practical advantages, including ease of use, deployment and maintenance, which we have not discussed in detail due to space limitations. Resources such as OpenClinical (<http://www.openclinical.org/>), where this system is also listed, provide extensive complementary material covering research, applications and demonstrations.

8 Conclusion and Future Work

The approach proposed in this paper combines three ontologies with rules in order to enhancing the accuracy of EMRs both by providing clinical decision support and improving the correctness of medical coding therefore reducing the number of rejected claims. We have presented a semantic approach which improves patient care

and satisfaction, and enables healthcare providers to complete all charge entry and documentation before the patient has left the office. At this time, we are unaware of any application similar to ASEMR that is in daily use, especially at small practices in any field of health care. During ISWC 2006, we have planned to organize group visits to AHC (which is 5 minutes from the conference venue) to enable all interested persons to observe the use of ASEMR in person (a canned demo is at <http://lsdis.cs.uga.edu/projects/asdoc/>). This work also demonstrate successful collaboration between academic research and small medical clinics. For business and legal reasons, we are unable to present some details such as error detection and reduction in this paper.

The ASEMR approach can be extended to provide decision support on a deeper level. For example, semantic associations [12] can be discovered to find even obscure relationships between symptoms, patient details, and treatments. Semantic alerts will also be explored in future versions such as when a physician scrolls down on the list of drugs and clicks on the desired drug, any study, clinical trial, or news item about the drug and other related drugs in the same category can be displayed. In addition ontologies can be utilized to find contradictions and mistakes in the medical report. Another key area of extension that we are also working on include coupling this system with a billing system with higher degree of automation (e.g., with better workflow and better validation of billing data) than current state of the art in medical billing.

Acknowledgements. We thank M. Eavenson, C. Henson, and D. Palaniswami at LSDIS for their effort in ontology design and population.

References

1. Agency for Healthcare Research & Quality
<http://ahrq.gov/news/press/pr2005/lowehrpr.htm>
2. E. Neumann and D. Quan, BioDASH: A Semantic Web Dashboard for Drug Development, Pacific Symposium on Biocomputing 11:176-187(2006) Also, <http://www.w3.org/2005/04/swls/BioDash/Demo/>
3. Centers for Medicare and Medicaid Services <http://www.cms.hhs.gov/>
4. H. Chen, D. Colaert, J. De Roo, Towards Adaptable Clinical Pathway Using Semantic Web Technology, W3C Workshop Semantic Web for Life Science, 2004.
5. V. Kashyap, A. Morales, T. Hongsermeier and Q. Li Definitions Management: A semantics-based approach for Clinical Documentation in Healthcare DeliveryIndustrial Track, Proceedings of the 4th International Semantic Web Conference, November 2005
6. D. McGuinness, and F. Harmelen, eds. OWL Web Ontology Language Overview <http://www.w3.org/TR/owl-features/>
7. Open Clinical http://www.openclinical.org/dm_ontobuilder.html
8. SNOMED <http://www.snomed.org/>
9. A. Seaborne, RDQL - A Query Language for RDFW3C Member Submission 9 January 2004, <http://www.w3.org/Submission/RDQL/>
10. W3, Semantic Web in Health-care and Life Sciences, www.w3.org/2005/04/swls/
11. Wordnet <http://wordnet.princeton.edu/>
12. K. Anyanwu and A. Sheth, "The ρ Operator: Discovering and Ranking Associations on the Semantic Web," The Twelfth International World Wide Web Conference, Budapest, Hungary, May 2003, pp. 690-699.