



HHS Public Access

Author manuscript

J Biomed Inform. Author manuscript; available in PMC 2016 June 13.

Published in final edited form as:

J Biomed Inform. 2012 August ; 45(4): 763–771. doi:10.1016/j.jbi.2012.01.009.

Building a robust, scalable and standards-driven infrastructure for secondary use of EHR data: The SHARPn project

Susan Rea^{a,*}, Jyotishman Pathak^b, Guergana Savova^c, Thomas A. Oniki^d, Les Westberg^e, Calvin E. Beebe^b, Cui Tao^b, Craig G. Parker^a, Peter J. Haug^{a,f}, Stanley M. Huff^{d,f}, and Christopher G. Chute^b

^aHomer Warner Center for Informatics Research, Intermountain Healthcare, Murray, UT, USA

^bDivision of Biomedical Statistics and Informatics, Department of Health Sciences Research, Mayo Clinic, Rochester, MN, USA

^cChildrens Hospital Boston and Harvard Medical School, Boston, MA, USA

^dIntermountain Medical Center, Intermountain Healthcare, Murray, UT, USA

^eAgilex Technologies, Inc., Chantilly, VA, USA

^fDepartment of Biomedical Informatics, University of Utah, Salt Lake City, UT, USA

Abstract

The Strategic Health IT Advanced Research Projects (SHARP) Program, established by the Office of the National Coordinator for Health Information Technology in 2010 supports research findings that remove barriers for increased adoption of health IT. The improvements envisioned by the SHARP Area 4 Consortium (SHARPn) will enable the use of the electronic health record (EHR) for secondary purposes, such as care process and outcomes improvement, biomedical research and epidemiologic monitoring of the nation's health. One of the primary informatics problem areas in this endeavor is the standardization of disparate health data from the nation's many health care organizations and providers. The SHARPn team is developing open source services and components to support the ubiquitous exchange, sharing and reuse or 'liquidity' of operational clinical data stored in electronic health records. One year into the design and development of the SHARPn framework, we demonstrated end to end data flow and a prototype SHARPn platform, using thousands of patient electronic records sourced from two large healthcare organizations: Mayo Clinic and Intermountain Healthcare. The platform was deployed to (1) receive source EHR data in several formats, (2) generate structured data from EHR narrative text, and (3) normalize the EHR data using common detailed clinical models and Consolidated Health Informatics standard terminologies, which were (4) accessed by a phenotyping service using normalized data specifications. The architecture of this prototype SHARPn platform is presented. The EHR data throughput demonstration showed success in normalizing native EHR data, both structured and narrative, from two independent organizations and EHR systems. Based on the demonstration, observed challenges for standardization of EHR data for interoperable secondary use are discussed.

*Corresponding author Address: Homer Warner Center for Informatics Research, Intermountain Medical Center, 5171 Cottonwood St., Murray, UT 84107, USA. Fax: +1 801 442 0748. Susan.Rea@imail.org (S. Rea).

Keywords

Electronic health records; Meaningful use; Health information exchange; High-throughput phenotyping; Natural language processing

1. Introduction

Strategic Health IT Advanced Research Projects (SHARP) were established under the direction of the Office of the National Coordinator for Health Information Technology (ONC) [1,2]. SHARP research and development is focused on achieving breakthrough advances to address well-documented problems that have impeded adoption of the electronic health record (EHR):

- Security and privacy of health information.
- User interfaces that support clinical reasoning and decision-making.
- Shared application and network architectures.
- Secondary use of EHR data to improve health [3,4].

The health improvements envisioned by the secondary-use SHARP consortium (SHARPN) include care process and outcomes measurement, more efficient biomedical research and support for public health activities. Critical to these secondary-use scenarios is the transformation of health information into standards-conforming, comparable, and consistent data. Traditionally, a patient's medical information, such as medical history, exam data, hospital visits and physician notes, are stored inconsistently and in multiple locations. SHARPN will enable the use of EHR data for secondary purposes by creating tangible, scalable, and open-source tools, services and software for large-scale health record data normalization and sharing [5].

One year into the design and development of the SHARPN framework, an integrated set of services and components was implemented as a prototype platform for natural language processing (NLP) and normalization pipelines to a persisted data store to cultivate high-throughput phenotyping. A demonstration of the prototype platform, focused on the throughput of normalized data, was conducted. Actual data from approximately ten thousand patient records were input from two independent EHRs: the Mayo Clinic and Intermountain Healthcare. EHR data corresponded to input data for a set of diabetes phenotyping rules used in practice. EHR data were sourced from both structured data and provider clinical narratives. Data were normalized to consistent, standardized forms in the SHARPN NLP and normalization components and stored in a shared database. A rules engine queried cases in the database for input data defined for a diabetes phenotyping algorithm. The target for this data throughput demonstration was the successful population of standardized data in the phenotyping data model.

In this paper, we present an overview of the architecture of this prototype SHARPN platform, describe the data flow from EHR to standardized data inputs for phenotyping, and discuss

the challenges for shared, standardized secondary EHR data usage as observed in this data throughput demonstration.

2. Background: the SHARPn framework

2.1. Data normalization and phenotyping architecture

The prototype platform supports a process flow implemented for this demonstration project. The SHARPn architecture consists of an adaptable framework of components and services. An overview of the SHARPn prototype data normalization and phenotyping architecture is shown below:

An overview of all processes and data flows (Fig. 1) is given in brief:

1. Intermountain Healthcare (IHC) structured clinical data were generated in Health Level Seven (HL7) [6] 2.x format. Other data input formats are supported in the architecture. Encounter ICD-9-CM [7] codes and attributions were prepared in tabular text files.
2. IHC data were pushed to Mirth Connect [8] on the IHC side of its firewall.
3. Mirth Connect is an open source product created by Mirth Corporation to transform messages among formats and/or route them from one location to another. Mirth Connect creates Nationwide Health Information Network (NwHIN) [9] Document Submission (XDR) messages from HL7 2.x and ICD-9-CM data files.
4. Mirth Connect routes the XDR message to the IHC NwHIN Aurion Gateway [10]
5. Aurion is an open source implementation of an NwHIN gateway. The gateway forwards the message using NwHIN protocols and services with secured two-way Transport Layer Security (TLS) with Security Assertion Markup Language (SAML) to the SHARPn NwHIN Aurion Gateway. This was implemented in a cloud service, residing behind the Mayo Clinic firewall.
6. The Aurion gateway on the SHARPn cloud passes the IHC XDR messages to Mirth Connect on the cloud.
 - 6a Mayo Clinic lab and ICD-9-CM data were generated as described for IHC above. These and clinical text documents are pushed to Mirth Connect on the SHARPn cloud behind Mayo Clinic's firewall.
7. Mirth Connect routes all incoming data to an Apache Unstructured Information Management Architecture (UIMA) [11] pipeline on the SHARPn cloud.
8. Normalization and NLP services are implemented in UIMA pipelines.
9. UIMA components route the normalized data back to Mirth Connect.
10. Mirth Connect routes the data instances to an open source SQL database, MySQL, where all data are persisted.

11. A rules engine queries the SQL database for normalized input data defined for a diabetes phenotyping algorithm.

The three main functions for the data normalization and phenotyping platform are described below:

2.2. EHR data normalization

2.2.1. Common information model—Clinical concepts must be normalized if decision support and analytic applications are to operate reliably on secondary EHR data. One EHR may store “resting heart rate” as a simple name-value pair (name = “resting heart rate” and value = “60 bpm”), while another EHR may store it in a more complex, *post-coordinated* fashion (name = “heart rate”, value = “60 bpm”, qualifier = “resting”). EHRs may contain more or less description or detail (*granularity*) about clinical events and observations. Further, the contents may be encoded using different terminologies, free text, or a mixture of both. A decision support application will need to take those differences into account when retrieving a clinical concept such as “resting heart rate”, and *normalize* the data to a common representation before processing.

The Clinical Element Model (CEM) is Intermountain Healthcare’s strategy for detailed clinical models.[12–19] A CEM provides a single logical structure for a data element, to which all incoming data of that type are normalized. CEMs are computable: they are written in the Constraint Definition Language (CDL) that can be compiled into objects that are passed between the layers of a system’s technology stack. CEMs *do not dictate the physical implementation of those objects or of the stack components*. For example, CEMs may be compiled into XML Schema Definitions (XSDs), java or C# programming language classes, or Semantic Web Resource Description Framework (RDF) to provide structured model definitions in particular computing environments [20]. Correspondingly, CEM instances may be XML instances, java or C# objects, or RDF instances. CEMs, like HL7 versions 2 and 3 models, prescribe that codes from controlled terminologies be used as the values of many of their attributes. Thus, CEMs facilitate interoperability between systems at a syntactic and semantic level while allowing for disparate architectures within those systems.

CEMs are agnostic of the origination of the data. Instances may be populated from textual reports by NLP services or from structured EHR data. All data stored in the EHR is encompassed in the scope of the CEM, including but not limited to:

- Allergies.
- Problem lists.
- Laboratory results.
- Medication and diagnostic orders.
- Medication administration.
- Physical exam and clinical measurements.
- Signs, symptoms, diagnoses.

- Clinical documents.
- Procedures.
- Family history, medical history and review of symptoms.

Fig. 2 shows a high-level abstraction of a medication order model and an instance of conformant data. The medication order CEM might prescribe that a medication order have attributes of “orderable item”, “dose”, “route”, and so on. Further, it may dictate that an orderable item’s value must be a code in the “orderable item value set” – a set list of controlled codes appropriate to represent orderable medications – and that the dose be represented as a “physical quantity” data type, which contains a numeric value and a code for a unit of measure, where that unit of measure code comes from a “dose unit” value set.

2.2.2. Terminology mapping services—Structured, coded terminology plays a critical role in CEMs. In specifying the content of data instances, CEMs also specify the places in the instances where coded concepts must be used. The CEMs used by SHARPN will specify standard terminologies conformant with consolidated health informatics (CHI) [21]. Terminology services – especially mapping services – are essential to SHARPN success. An overview of the envisioned SHARPN normalization services are depicted in Fig. 3.

SHARPN will align the terminology services it uses with the evolving Common Terminology Services (CTS) effort [22]. CTS defines a standard set of service interfaces for accessing and maintaining structured terminology. The second version of this effort (CTS2) is currently moving through the Object Management Group (OMG) standardization process.

2.3. Natural language processing of health care documentation text

A substantial part of the information pertaining to a given patient is recorded in the unstructured, free-text part of the EHR. Information extraction (IE) is an application of NLP aiming at converting the free-text into structured information. Normalization targets can vary from simple mappings to a terminology to more template-like structures. The CEMs provide such templates. Traditionally, IE and NLP have been applied to the domain of medicine for specific use cases where only a limited set of signals are to be discovered, such as discovering pneumonia from chest X-ray reports [23,24] and peripheral arterial disease from radiology notes [25]. These approaches result in excellent performance but the application is limited to only the specific use cases for which they were developed. Our SHARPN philosophy takes a different approach. The goal is to discover a relevant set of summary information for a given patient in a disease- and use-case agnostic way, which is then stored and indexed for retrieval. The main principle, thus, is *process once, use multiple times*.

In the SHARPN project, the core NLP software is the clinical Text Analysis and Knowledge Extraction System (cTAKES) [26]. As is true for the SHARPN platform, cTAKES is built within the Apache Unstructured Information Management Architecture (UIMA) which provides a solid basis for scalability, expandability and collaborative software development. One of the main UIMA concepts is that of a type system, which defines the annotations and their structure as generated by the pipeline. An agreed-upon basic type system ensures a software development foundation. The SHARPN NLP type system, which continues to be

developed and evaluated, [27] implements the templates which are derived from the CEMs and represent the NLP normalization targets. The current cTAKES pipeline processes clinical notes and identifies the following clinical named entity mentions: drugs, diseases/disorders, signs/symptoms, anatomical sites, and procedures. The dictionary lookup algorithm uses a knowledge base as a resource created off SNOMED Clinical Terms[®] (SNOMED CT[®]) [28] and RxNorm [29] terms and their synonyms. Document input can be plain text or Clinical Document Architecture (CDA) [30] compliant XML documents.

2.4. High-throughput phenotyping

The SHARPn view of phenotyping includes inclusion and exclusion criteria for clinical trials, numerator and denominator criteria for clinical quality metrics, epidemiologic criteria for outcomes research or observational studies, and trigger criteria for clinical decision support rules, among others. Our use implies the algorithmic recognition of any cohort within an EHR for a defined purpose. These purposes were inspired by the algorithmic identification of research phenotypes in the National Human Genome Research Institute (NHGRI) funded electronic MEDical Records and GENomics (eMERGEs) consortium [31] primarily to facilitate EHR-derived genomics studies. The underlying principles of using well-defined algorithms across diagnosis assertions, laboratory values, medication use, and NLP-derived observations adheres to the practices demonstrated in eMERGE [32].

Phenotyping algorithms in essence comprise a complex set of inclusion and exclusion criteria that are coupled using a set of logical operators. This lends to representing the algorithm criteria using a rules-based formalism. For SHARPn, we chose the open-source Drools production rule system [33]. Supported by the JBoss community, Drools comprises an inference engine, complex event processing and workflow system that provides a rich, declarative programming environment. The inference engine, also known as a “production rules system”, represents clinical knowledge as simple “conditions” and “actions”. The Drools system excels in that the rules engine and the workflow engine are tightly integrated and workflow can be used to visually represent complex relationships between logically related groups of rules. Development of a robust, common data model and scalable approaches to implement computable phenotyping algorithms within the Drools system are current SHARPn objectives.

The prototype SHARPn architecture includes persistence of the data. One can regard the persistence function as a data store for normalized representations. This is necessary because phenotyping constitutes a question/answer process, where within an algorithm one ascertains whether a particular patient does or does not manifest definitional criteria. This persistence layer can be realized by many data base technologies including: XML databases, object databases, RDF triple stores and SQL databases. For convenience in the development phases, we opted for SQL database technology. The ultimate goal is to decouple the algorithms from the physical implementation of the persistence layer.

3. Methods: data throughput demonstration

After one year of design and development effort, a SHARPn platform was implemented for internal use. A demonstration of end to end throughput of high volume, real patient data was

conducted in June, 2011. Approximately ten thousand patient records from two independent EHR systems, Mayo Clinic and Intermountain Healthcare were used to test the end to end flow of actual data. The use case was a common diabetes mellitus phenotyping algorithm. The algorithm was modified to normalized data specifications. The demonstration was focused on the acquisition and normalization of disparate data, and successful data flow from EHR to the population of input data in the rules engine.

The SHARPN platform infrastructure, software artifacts and research and development processes are designed and conducted with appropriate consideration of health information security and privacy requirements. Institutional review board policies are followed, and data are shared under appropriate data sharing agreements. In particular, site to site EHR data transmission for this demonstration involved exclusively de-identified data. Communication between the SHARPN organizations and the SHARPN cloud are secured in accordance with the NWHIN specifications. There is ongoing collaboration with the SHARP consortium tasked with EHR data security solutions.

3.1. Use case – diabetes mellitus phenotyping algorithm

A phenotyping algorithm for a medically managed diabetes mellitus cohort was the use case for the data throughput demonstration. The algorithm was adapted from the eMERGE Northwestern Type 2 Diabetes Mellitus phenotyping algorithm [34] by the Southeast Minnesota Beacon Program [35] for purposes of group practice reporting.

A simplified algorithm (Fig. 4) was derived to validate data throughput. The normalized EHR data to be used as input to the algorithm for this demonstration were:

- All ICD-9-CM codes in time period 1/1/2009–12/31/2010.
- All ambulatory Medications ordered or noted as taken (RxNorm).
- Lab values (LOINC® [36]), results and date collected in time period 1/1/2009–12/31/2010.

3.2. Source EHR data and acquisition

3.2.1. Mayo Clinic—10,000 patients seen in 2010 at the Mayo Clinic were randomly selected from the Mayo research data warehouse, Enterprise Data Trust (EDT). The selection criteria were: (a) valid research authorization, (b) not deceased, (c) 18 years as of June, 2011, (d) minimum of two visits in the last year with ICD-9-CM codes, and (e) patient had at least 3 years of laboratory and medication data available. For selected patients, three years of prior clinical data were accessed, ending on June 30, 2010. The data sets selected were clinical notes, laboratory results and ICD-9-CM codes from the EDT. Unlike the de-identified Intermountain datasets, the Mayo patient data was processed locally and used to assess the platform's processing throughput and performance capabilities. No specific patient results were used or released.

Fig. 5 depicts the data acquisition and flow. The cTAKES pipeline processed CDA formatted clinical documents from the Mayo Clinic. The focus for this demonstration was discovery of all noted medications relevant to the patient. The clinical notes of 10,000 patients from

2008–2010 were pulled from the Mayo EDT, resulting in 360,452 documents. Each document was then processed through cTAKES, matching RxNorm terms and their synonyms. This resulted in 3442,000 Drug CEMs, which were passed to Mirth to deposit into the CEM database for phenotyping consumption.

The laboratory data was formatted into HL7 laboratory results messages using HL7 2.x syntax. The results were then normalized in a configurable UIMA pipeline. The ICD-9-CM data was extracted from the Mayo EDT and processed from tabular representations in text files by the AdminDiagnosis UIMA Processor.

3.2.2. Intermountain Healthcare—One hundred cases were randomly selected from a curated pool of Intermountain test cases established for EHR development and validation activities. The test cases were developed from data stored in the Intermountain Enterprise Data Warehouse (EDW). The test data are real patient data, with all protected health identifiers of patients and providers removed and replaced with realistic false data. The encounter dates and other dates were shifted, using a random seed, so that the temporality of events is maintained. The test data were de-identified in compliance with HIPAA Guidelines.

All Medication List entries, ICD-9-CM codes with sequence numbers for the associated ambulatory encounters, and Intermountain clinical laboratory results for a period of time 7/1/2008–6/30/2011, to allow overlap with time period of interest for the classification algorithm, were extracted for the 100 test cases. Medication and laboratory data were prepared in HL7 2.7 format. This was accomplished by taking example Pharmacy Order and Unsolicited Laboratory Result HL7 2.7 messages from the Intermountain production interfaces. Using these formats as templates, the false identifiers and true EDW data for the test cases were inserted into HL7 2.7 message and pushed to the framework where they were processed just as Mayo Clinic's. (Fig. 5)

The Medication List may also contain text medication entries. For this demonstration, only coded medication entries were selected. The medications are encoded using First DataBank [37] (FDB) orderable drug codes (GCN). This code and the linked text description were sent in the HL7 input messages.

Although ICD-9-CM codes can be messaged in HL7 Diagnosis (DG1) segments, neither Mayo Clinic nor Intermountain Healthcare uses HL7 operationally for ambulatory ICD-9-CM codes. It was decided to prepare the ICD-9-CM data in a simple tab delimited, columnar text file. The unique ambulatory encounter identifiers, the shifted and de-identified encounter dates, the ICD-9-CM codes, their sequence numbers, and the provider service for the encounter were included. These data were processed like those from Mayo Clinic. (Fig. 5)

3.3. Data normalization processes

3.3.1. Clinical Element Models—The data throughput demonstration required three CEMs:

- Ambulatory medication order.

- Standard lab panel.
- Administrative diagnosis.

These CEMs can be viewed at <http://www.intermountain-healthcare.org/cem>.

The CEMs were compiled into XSDs, which formed the basis of data instances processed by the SHARPN framework components. HL7 message structures, the tabular input files for administrative diagnosis codes, and the output of the NLP processes were mapped to these XSDs in UIMA Common Analysis Structure (CAS) components. UIMA CAS components render a common data structure for the application. Incoming HL7 and tabular data and NLP output were processed to populate XML instances that conformed to the CEM XSDs.

3.3.2. Terminology mapping—The UIMA HL7 pipelines that processed structured medications and labs had components to translate site specific coded data to standardized terminologies. (Fig. 6) Mayo internal lab codes and Sunquest codes from Intermountain were translated to LOINC codes. Intermountain FDB ordered medication (GCN) codes were translated to RxNorm Ingredient or Multiple Ingredient codes. We utilized a time tested and basic mechanism to achieve the translation step. The crosswalk tabular data (Mayo internal to LOINC, Sunquest Intermountain to LOINC and FDB GCN to RxNorm) were loaded into memory as lookup resources at initialization time, which then were used conditionally to perform the crosswalk at execution time, based on the origination of the data. These components served as proxies for the standards-based SHARPN terminology services currently in development.

3.4. Data input to diabetes phenotyping rules implemented in Drools

Once the data were transformed and normalized to a consistent form, they were stored into a simple relational database, specifically open source MySQL. The CEM normalized data instances are represented as XML document instances. To facilitate access to that content, an index data record was added to support query access to stored clinical element data instances. Using these index records, the Drools data acquisition components were able to identify and access specific data stored in the normalized database.

The data throughput end goal was to populate a simple data model accommodating the three data types in the use case: ICD-9-CM codes, medications and laboratory results. The majority of data were successfully processed. Neither specific data transformation errors nor accuracy of the phenotyping rules were evaluated for this data throughput trial.

4. Results and discussion

The data throughput trial demonstrated successful functionality for three design goals: a secure, cloud-based transport architecture; a common syntax; and standardized semantics. These are requirements for interoperability of EHRs [38]. Both NLP-derived and structured data were normalized into common concept instances, which enabled a phenotyping algorithm to operate on patient data regardless of the concept's origin in the EHR. We believe the conceptual merging of structured and text-based EHR data is critically important to secondary use of EHR data as well as for patient data exchange for care purposes,

decision support and the general case for retrieving a patient's information in the EHR. Recent studies have shown the value in using both text and structured data for more accurate disease classification [39–41]. The CEM's concept-based formalism unites them. The CEMs can be consumed at the concept level, or they can be constrained by their attributions if the researcher wishes to process NLP-derived versus structured source data separately.

SHARPN is focused on the generation of publicly available services and components to enable liquidity of standards-conforming, comparable, and consistent data for broad scale, high-throughput secondary use of the EHR. The architecture supports the definition and instantiation of terminology standards, as they are declared, into normalization components and services. The standardized data elements thus generated conform to a conceptual information model, and may be implemented in physical objects and physical information models as needed. We have described and demonstrated this functionality, and applied three terminology standards: ICD-9-CM, RxNorm, and LOINC. The implementation of this prototype platform and data throughput demonstration suggested challenges to our secondary EHR data liquidity goals:

4.1. Platform architecture

A normalization and secondary use platform of the intended scalability requires a robust and secure architectural foundation. Data security, end to end connectivity and reliable data flow were critical to the normalization and phenotyping services. In the data throughput demonstration, architectural issues were discovered as expected. The data exchange and host processing platforms will require ongoing expertise from computer scientists with an understanding of health information exchange domain requirements.

4.2. Error handling

Error checking and handling is a requirement for the framework. The prototype platform did not implement this functionality. Errors are expected to occur in multiple components. Some places that errors will occur are: (1) transmission related errors, (2) transformation errors, (3) EHR data quality errors, and (4) data persistence errors. Error recognition and proper handling will be highly important features of ubiquitous normalization and secondary use of EHR data [38,42].

4.3. Acquiring source EHR data

The pilot showed that HL7 2.x messages could be sent as the payload of an NWHIN Document Submission message, which could then be transformed and normalized using the UIMA pipeline. However, it is envisioned that there would be multiple formats for messages sent to the SHARPN cloud. HL7 messaging may cover the data that are typically interfaced between the EHR and ancillary systems, such as clinical lab, provider text documentation, and pharmacy. HL7 includes message structures that accommodate a variety of health information, such as diagnosis codes, allergies, and general clinical observations. But if organizations' operations do not include messaging those data, it is arguably no easier to acquire source data using a standardized HL7 representation than using other representations that can be mapped to CEMs.

The demonstration forced us to rethink the strategy of using HL7 Admit/Discharge/Transfer (ADT) messages to communicate ICD-9-CM codes. Since the organizations' business practices did not use HL7 to communicate ambulatory encounter codes, a text file format was populated by query at each organization and pushed to Mirth. Operationally, ICD-9-CM codes are messaged from providers to payors in ANSI standard X-12 messages [43], required under HIPAA for claims. These existing messaging processes and structure could be considered to harvest ICD-9-CM data in an automated manner. The acquisition of ICD-9-CM codes must consider their source. A practice management system may contain codes submitted or entered by providers, whereas other workflow processes may occur to select a subset of the codes or even modifications to submit as actual claims for billing purposes. These workflow processes differ among organizations, as can be stated about processes of populating EHR data in general.

There is untallied variation in the style, organization and formatting of clinical text documents among organizations and EHRs. The Mayo Clinic EHR environment implements a number of HL7 standards for clinical document organization. The CDA standard defines XML tags which mark the beginning and end of document sections. This unambiguous sectioning enables selective extraction of medications with variable levels of confidence depending on the section in which they occur in. A productive strategy is to limit the medication extraction to a set of sections such as Impression/Report/Plan, History of Present Illness, Current Medications, Admission Medications and Discharge Medications. The improvement in the accuracy of NLP information extraction methods when standardized clinical document styles are used remains a significant question. The current lack of standard styles in many clinical documents remains a challenge for normalization.

Design and development of SHARPn source EHR data acquisition tools require analysis of the origin, storage, and standard extraction or messaging mechanisms that may be tapped. In future iterations of the framework, Mirth Connect can take messages in many different forms, using many different protocols and transform them. SHARPn teams continue to develop acquisition strategies, focusing on the most useful EHR data categories. Meaningful use leaders and the National Quality Foundation have selected priority data types to support in EHRs and for quality measures [44,45].

4.4. Normalization models

The CEMs used in the demonstration were based on the models created to populate EHR content. Based on feedback from the informatics community, SHARPn will instead design CEMs to cover general requirements for a spectrum of secondary use cases, realizing there will be needs for additional or different data requirements for particular uses. The clinical element models can be revised or extended as those needs arise. It is a challenge to create canonical models of EHR data for secondary use, as there are a variety of purposes to be served. However, large scale normalization of data for shared secondary use depends upon common models.

4.5. Terminology services

The prototype platform did not implement runtime terminology services – simple table lookup mechanisms were implemented. Future SHARPN work will focus on calling standardized terminology services accessing a robust terminology server, as discussed in Section 2.2.2. Nevertheless, important insights regarding terminology mapping and services were gained during the pilot.

4.5.1. Mapping of local terminologies—A significant effort in the implementation of an EHR is mapping an institution’s local codes to the terminologies stored and used within the EHR. Specialists or integration consultants are often required for the work. The SHARPN effort proved to be no exception as effort was required to complete the provider institutions’ mappings from local laboratory observation codes to LOINC, even for the small demonstration project that was conducted. A large-scale data normalization project will need to take into account the considerable effort of mapping institutions’ terminologies to standard terminologies. Compounding the problem is that many institutions will not be able to produce on demand the code sets sent in their messages in a form consumable by a terminology server, or at all. “Smart” tools are called for – tools that can analyze months-worth of an institution’s messages and data and from them deduce the local code sets that need to be mapped to standards, and then even postulate the mappings.

Further, the inclusion of CEMs in the normalization solution suggests that not just terminology mapping solutions but model mapping solutions are required.

4.5.2. Medication terminology—Mapping of medication codes is a particular type of mapping that introduces its own unique complexities. Healthcare provider institutions, vendors of drug information code sets, and the RxNorm standard all have multiple classes or hierarchies of codes, i.e., ingredients, generic drugs, branded drugs, and therapeutic classes. There are inevitably small differences in the structures of the parties’ hierarchies and classes that cause mappings to be difficult. The mappings included with RxNorm realize much of the mapping work required between RxNorm and vendors; however, the work of mapping individual providers’ codes and hierarchies to RxNorm will remain.

Furthermore, choices between hierarchies and classes need to be made. For example, Intermountain Healthcare chose to map its clinical drug data to one of the two different orderable drug code sets maintained by their drug information vendor. That class of drug code can map to RxNorm ingredient codes, RxNorm generic orderable drug codes, or RxNorm branded orderable drug codes. For the demonstration, we had to choose which of those to map to, and communicate to all parties involved: the team performing NLP, the team normalizing data, and the team executing phenotyping algorithms. The choices depend largely upon use cases. Decision support algorithms, for example, need to *know* when a patient is on any drug whose formulation contains a particular ingredient and are uninterested in the particular branding, formulations, dosage forms, or strengths of those drugs. Ordering applications, on the other hand, are very interested in brands, formulations, forms, and strengths. Given that different use cases demand different hierarchies of drug codes, it is inevitable and essential that support be provided to navigate between hierarchies.

4.5.3. Terminology versioning and updating—The SHARPN demonstration did not deal with multiple versions of terminologies or updates to terminologies, but it became apparent that any robust data normalization effort will need to do so. Controlled terminologies are not static, but undergo constant change as concepts are replaced by/redirected to other concepts, concepts are obsoleted, and new concepts are added. Generally, when an incoming code is mapped to a standard terminology, the latest changes in the standard terminology must be recognized and accessed through the mapping services. If a code has been made obsolete in the standard, a mapping to that code must be adjusted to point to the code that replaces the obsolete code and the services must navigate to the new code. If new codes are added to the standard, mappings must be updated and the services must access the new mappings in order to navigate to the newly added codes. If the meaning of a code changes, which occasionally occurs in some terminologies, mappings involving that code must be adjusted so as to continue to produce consistent results.

In some use cases, specific versions of terminologies must be recognized and accessed. Incoming codes may not always be from the latest version of a terminology, but they still need to be understood by the services and mapped correctly. Reporting requirements may dictate that certain versions of standard terminologies be sent out, requiring mappings from the internally-stored codes to various versions of outgoing terminologies. One example of this is found in standardization efforts such as Healthcare Information Technology Standards Panel (HITSP) and HL7 v3, which support the notion of dynamic and static bindings. In dynamic binding, a data element is bound to the current contents of a specified code set, e.g., the HITSP Procedure value set is bound to codes in the procedure hierarchy in whatever is the most current version of SNOMED CT. In contrast, in static binding, a data element is bound to a specific version of a code set, e.g., the HITSP Diagnosis Priority value set is bound to the 2.5.1 version of the HL7 Diagnosis Priority code system. Consequently, static binding requires terminology servers and services that maintain multiple versions of terminologies.

Terminology changes may occur frequently. RxNorm, for instance, is released monthly in an attempt to keep up with the frequently changing medication realm. Changing terminologies will clearly require considerable maintenance effort in order to make SHARPN data normalization and interoperability successful.

4.6. Specification of standardized phenotyping algorithms

Another challenge highlighted by the demonstration was the difficulty of expressing this simple phenotyping algorithm's data specifications in terms of normalized objects and standard terminologies. Cohort identification algorithms' data are often expressed in abstract clinical terms, such as 'fasting glucose'. This expression maps to several LOINC codes. Existing data specifications must be translated to standardized data concept specifications. The computable specification of standardized phenotyping algorithm data and rules among them, are a recognized challenge with active engagement in the health informatics community.

5. Conclusions

A prototype platform was implemented from the developing SHARPN framework to perform secure transport, data normalization and common phenotyping services on disparate EHR data. A demonstration of data throughput from two large organizations was conducted to test the design and inform future development. The demonstration was successful in executing the intended end to end data processing flow, merging text and structured data into instances of clinical concepts, and normalizing disparate data to common objects with standard terminologies. We have shared the emerging architecture and observed challenges for standardization of EHR data for interoperable secondary use.

Acknowledgments

This manuscript was made possible by funding from the Strategic Health IT Advanced Research Projects (SHARP) Program administered by the Office of the National Coordinator for Health Information Technology. The contents of the manuscript are solely the responsibility of the authors.

The authors would like to thank the following individuals that assisted in this project or provided the computing infrastructure that made it possible. *Mayo Clinic*: Jay Doughty, Sridhar Dwarkanath, Donna Ihrke, Vinod Kaggal, James Masanz; *Childrens Hospital Boston*: Pei Chen; *Homer Warner Center for Informatics Research, Intermountain Healthcare*: Jeffrey Ferraro, John Holmen, PhD, Herman Post, Darin Wilcox. We would also like to thank Lacey A. Hart, MBA, PMP[®], Director, PMO Health Sciences Research, SEMN Beacon Program Manager, and SHARP Executive Officer, for providing administrative oversight and project management services.

References

1. The Office of the National Coordinator for Health Information Technology. US Department of Health and Human Services; 2011. <http://healthit.hhs.gov/portal/server.pt/community/healthit_hhs_gov_home/1204>[cite 10.08.11]
2. Strategic Health IT Advanced Research Projects (SHARPs) Program. US Department of Health and Human Services; 2011. <http://healthit.hhs.gov/portal/server.pt/community/healthit_hhs_gov_sharp_program/1806>[cite 28.08.11]
3. Hersh WR. Adding value to the electronic health record through secondary use of data for quality assurance, research, and surveillance. *Am J Manag Care*. 2007; 13:277–8. [PubMed: 17567224]
4. Safran C, Bloomrosen M, Hammond WE, Labkoff S, Markel-Fox S, Tang PC, et al. Toward a national framework for the secondary use of health data: an American medical informatics association white paper. *J Am Med Inform Assoc*. 2007; 14:1–9. [PubMed: 17077452]
5. Chute, CG.; Pathak, J.; Savova, GK.; Bailey, KR.; Schor, MI.; Hart, LA., et al. The SHARPN Project on secondary use of electronic medical record data: progress, plans, and possibilities. American medical informatics association 2011 annual symposium; Washington (DC). 2011; p. 248-256.
6. V2 Messages. Health Level Seven International; 2011. <<http://www.hl7.org/implement/standards/v2messages.cfm>>[cite 1.11.11]
7. ICD-9-CM Official Guidelines for Coding and Reporting. The Centers for Medicare and Medicaid Services (CMSs) and the National Center for Health Statistics (NCHSs); 2009.
8. Mirth Connect. Mirth Corporation; 2011. <<http://www.mirthcorp.com/community/mirth-connect>>[cite 1.11.11]
9. Sullivan, T., editor. Mapping the future of NwHIN. Government Health IT, MedTech, Media; 2011.
10. Westberg, L. Aurion Federal Gateway WSDL Interfaces. 2011. <<http://wiki.aurion-project.org/display/AurionMain/Aurion+Federal+Gateway+WSDL+Interfaces>>[cite 1.11.11]
11. Apache UIMA. The Apache Software Foundation; 2011. <<http://uima.apache.org/>>[cite 1.11.11]
12. Coyle JF, Mori AR, Huff SM. Standards for detailed clinical models as the basis for medical data exchange and decision support. *Int J Med Inform*. 2003; 69:157–74. [PubMed: 12810120]

13. Huff, SM. Practical modeling issues: representing coded and structured patient data in EHR systems. AMIA 2010 annual symposium; Washington (DC). 2010;
14. Huff SM, Rocha RA, Bray BE, Warner HR, Haug PJ. An event model of medical information representation. *J Am Med Inform Assoc.* 1995; 2:116–34. [PubMed: 7743315]
15. Huff SM, Rocha RA, Coyle JF, Narus SP. Integrating detailed clinical models into application development tools. *Stud Health Technol Inform.* 2004; 107:1058–62. [PubMed: 15360974]
16. Huff SM, Rocha RA, Solbrig HR, Barnes MW, Schrank SP, Smith M. Linking a medical vocabulary to a clinical data model using abstract syntax notation 1. *Methods Inf Med.* 1998; 37:440–52. [PubMed: 9865042]
17. Oniki, T. A detailed clinical model development environment. 12th International HL7 interoperability conference; Orlando (Florida, US). 2011;
18. Parker CG, Rocha RA, Campbell JR, Tu SW, Huff SM. Detailed clinical models for sharable, executable guidelines. *Stud Health Technol Inform.* 2004; 107:145–8. [PubMed: 15360792]
19. Steiner DJ, Coyle JF, Rocha BH, Haug P, Huff SM. Medical data abstractionism: fitting an EMR to radically evolving medical information systems. *Stud Health Technol Inform.* 2004; 107:550–4. [PubMed: 15360873]
20. Tao, C.; Parker, CG.; Oniki, TA.; Pathak, J.; Huff, SM.; Chute, CG. An OWL meta-ontology for representing the clinical element model. American medical informatics association annual, symposium; 2011;
21. Office of the National Coordinator for Health Information Technology (ONC) Presidential Initiatives: Consolidated Health Informatics US Department of Health & Human Services. 2006. <<http://www.hhs.gov/healthit/chiinitiative.html>>[cite 15.08.11]
22. Common Terminology Service 2. Mayo Clinic; 2011. <<http://informatics.mayo.edu/cts2>>[cite 15.08.11]
23. Christensen, LM.; Haug, PJ.; Fiszman, M. Proceedings of the ACL-02 workshop on natural language processing in the biomedical domain. Vol. 3. Philadelphia (Pennsylvania): Association for Computational Linguistics; 2002. MPLUS: a probabilistic medical language understanding system; p. 29-36.
24. Fiszman M, Chapman WW, Aronsky D, Evans RS, Haug PJ. Automatic detection of acute bacterial pneumonia from chest X-ray reports. *J Am Med Inform Assoc.* 2000; 7:593–604. [PubMed: 11062233]
25. Savova GK, Fan J, Ye Z, Murphy SP, Zheng J, Chute CG, et al. Discovering peripheral arterial disease cases from radiology notes using natural language processing. *AMIA Annu Symp Proc.* 2010; 2010:722–6. [PubMed: 21347073]
26. Savova GK, Masanz JJ, Ogren PV, Zheng J, Sohn S, Kipper-Schuler KC, et al. Mayo clinical Text Analysis and Knowledge Extraction System (cTAKES): architecture, component evaluation and applications. *J Am Med Inform Assoc.* 2010; 17:507–13. [PubMed: 20819853]
27. Wu, S.; Kaggal, V.; Savova, G.; Liu, H.; Dligach, D.; Zheng, J., et al. Generality and reuse in a common type system for clinical natural language processing. *Managing Interoperability and Complexity in Health Systems (MIXHS 2011)* in conjunction with the 20th ACM international conference on information and knowledge management; Glasgow (UK). 2011;
28. Unified Medical Language System[®] (UMLS[®]) SNOMED Clinical Terms[®] (SNOMED CT[®]). US National Library of Medicine; 2011. <http://www.nlm.nih.gov/research/umls/Snomed/snomed_main.html>[cite 17.01.12]
29. Unified Medical Language System[®] (UMLS[®]) RxNorm. US National Library of Medicine; 2011. <<http://www.nlm.nih.gov/research/umls/rxnorm/>>[cite 17.01.12]
30. Dolin RH, Alschuler L, Boyer S, Beebe C, Behlen FM, Biron PV, Shabo Shvo A. HL7 clinical document architecture: release 2. *J Am Med Inform Assoc.* 2006;30–9. [PubMed: 16221939]
31. McCarty CA, Chisholm RL, Chute CG, Kullo IJ, Jarvik GP, Larson EB, et al. The eMERGE network: a consortium of biorepositories linked to electronic medical records data for conducting genomic studies. *BMC Med Genomics.* 2011; 4:13. [PubMed: 21269473]
32. Kullo IJ, Fan J, Pathak J, Savova GK, Ali Z, Chute CG. Leveraging informatics for genetic studies: use of the electronic medical record to enable a genome-wide association study of peripheral arterial disease. *J Am Med Inform Assoc.* 2010; 17:568–74. [PubMed: 20819866]

33. Drools Business Logic integration Platform. JBoss Community; 2011. <<http://www.jboss.org/drools>>[cite 15.08.11]
34. eMERGE. eMERGE Network Phenotype Library; 2011. <https://www.mc.vanderbilt.edu/victr/dcc/projects/acc/index.php/Library_of_Phenotype_Algorithms>[cite 13.03.11]
35. Southeast Minnesota Beacon Program. 2011. <<http://semnbeacon.wordpress.com/>>[cite Aug 10, 2011]
36. Logical Observation Identifiers Names and Codes (LOINC®). Regenstrief Institute, Inc; 2012. <www.loinc.org>[cite 17.01.12]
37. First DataBank. <www.firstdatabank.com> [cite 17.01.12]
38. Hammond WE, Bailey C, Boucher P, Spohr M, Whitaker P. Connecting information to improve health. *Health Aff (Millwood)*. 2010; 29:284–8. [PubMed: 20348075]
39. Li L, Chase HS, Patel CO, Friedman C, Weng C. Comparing ICD9-encoded diagnoses and NLP-processed discharge summaries for clinical trials pre-screening: a case study. *AMIA Annu Symp Proc*. 2008:404–8. [PubMed: 18999285]
40. Liao KP, Cai T, Gainer V, Goryachev S, Zeng-treitler Q, Raychaudhuri S, et al. Electronic medical records for discovery research in rheumatoid arthritis. *Arthritis Care Res (Hoboken)*. 2010; 62:1120–7. [PubMed: 20235204]
41. Murff HJ, FitzHenry F, Matheny ME, Gentry N, Kotter KL, Crimin K, et al. Automated identification of postoperative complications within an electronic medical record using natural language processing. *JAMA*. 2011; 306:848–55. [PubMed: 21862746]
42. Chan KS, Fowles JB, Weiner JP. Review: electronic health records and the reliability and validity of quality measures: a review of the literature. *Med Care Res Rev*. 2010; 67:503–27. [PubMed: 20150441]
43. ASC X12. The Accredited Standards Committee; 2011. <<http://www.x12.org/>>[cite 15.08.11]
44. Quality Data Model. National Quality Forum; 2011. <http://www.qualityforum.org/Projects/h/QDS_Model/Quality_Data_Model.aspx>[cite 15.08.11]
45. Standards and Certification. Office of the National Coordinator for Health Information Technology; 2011. <http://healthit.hhs.gov/portal/server.pt/community/healthit_hhs_gov__standards_and_certification/1153>[cite 15.08.11]




Fig. 1. SHARPN data normalization and phenotyping architecture.

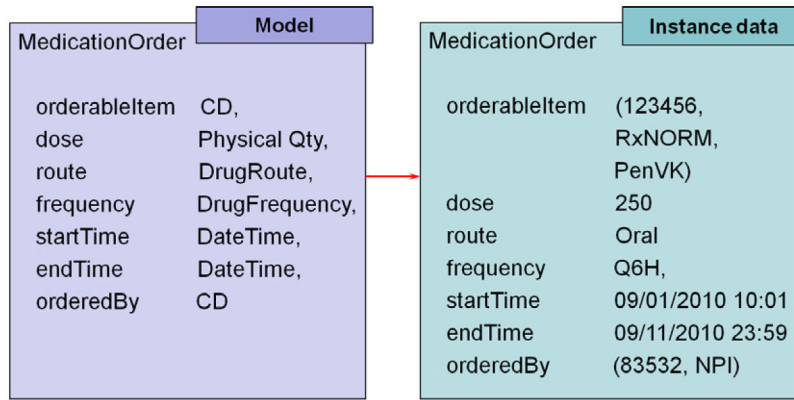


Fig. 2.
Partial CEM model and instance data for a medication order.

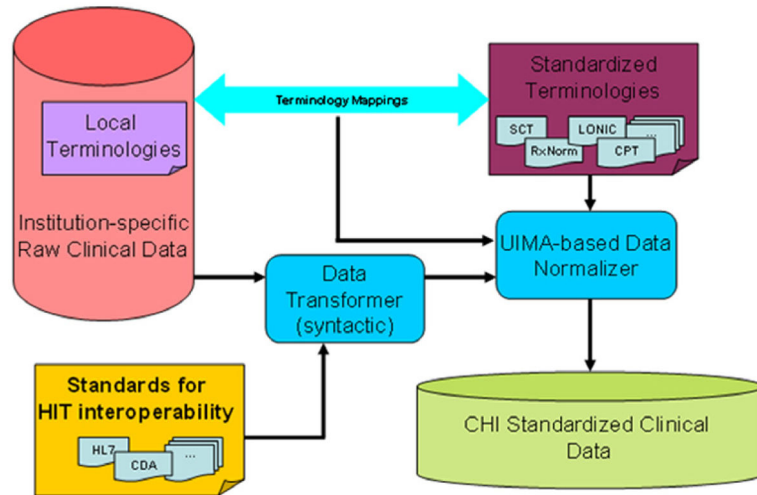


Fig. 3.
SHARPN normalization services overview.

Patient meets AT LEAST One of the following Criteria:

At least 2 face-to-face outpatient visits [1/1/09-12/31/10] with a diabetes ICD-9-CM code *[list of codes given]*

OR

Any medications ordered in *[list of brand and generic medication names]*

OR

At least 2 face-to-face outpatient visits [1/1/09-12/31/10] with a capillary glucose order in the measurement period OR with an abnormal blood glucose or HbA1c result *[abnormal thresholds given]*, not glucose tolerance test and not part of pregnancy screen.

Fig. 4.
Simplified medicinally managed diabetes classification algorithm.

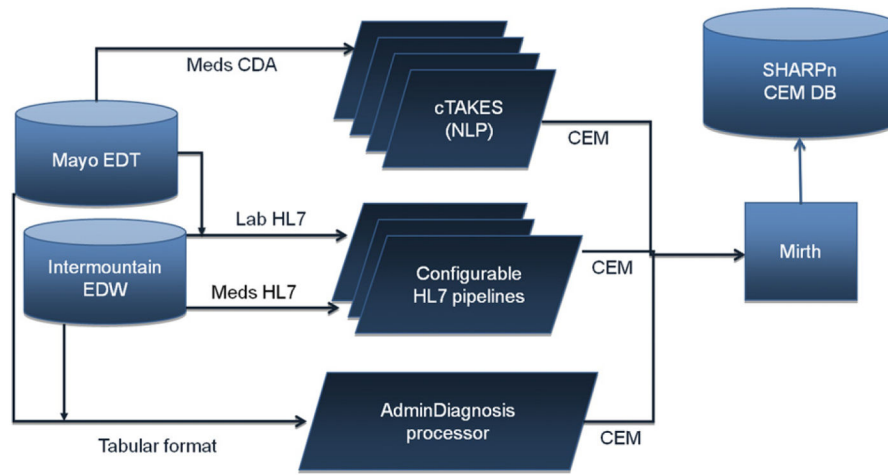


Fig. 5. UIMA normalization processes accessed via mirth.

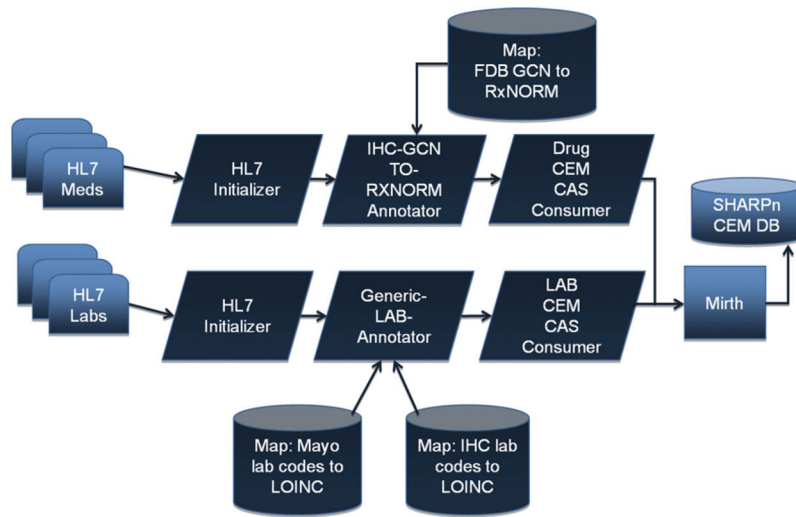


Fig. 6. UIMA terminology mapping processes.