Semantic Enrichment of Streaming Healthcare Data

Daniel Cotter, MS¹, V. K. Cody Bumgardner, PhD^{2,3}

¹Information Technology, University of Kentucky Medical Center, Lexington, Kentucky;
²Department of Pathology and Laboratory Medicine, Lexington, Kentucky; ³Department of Computer Science, Lexington, Kentucky

Abstract

In the past decade, the healthcare industry has made significant advances in the digitization of patient information. However, a lack of interoperability among healthcare systems still imposes a high cost to patients, hospitals, and insurers. Currently, most systems pass messages using idiosyncratic messaging standards that require specialized knowledge to interpret. This increases the cost of systems integration and often puts more advanced uses of data out of reach. In this project, we demonstrate how two open standards, FHIR and RDF, can be combined both to integrate data from disparate sources in real time and make that data queryable and susceptible to automated inference. To validate the effectiveness of the semantic engine, we perform simulations of real-time data feeds and demonstrate how they can be combined and used by client-side applications with no knowledge of the underlying sources.

Introduction

Healthcare Information Technology and the Interoperability Problem

Since the early 1970s, healthcare information technology has moved toward a comprehensive electronic medical record (EMR) in which almost every aspect of the patient's healthcare has been digitized and retained indefinitely[1], which has vastly improved the efficiency with which patient information can be retained, communicated, and analyzed. At the same time, the healthcare industry has moved from a fee-for-service model to a value-based model, facilitated in part by the existence of such a record and in part by public policy, such as the Health Information Technology for Economic and Clinical Health (HITECH) Act of 2009 [2], which provided financial incentives for the "meaningful use" of electronic medical records.

The realization of a holistic medical record has been slowed by various obstacles, chief among them is the problem of interoperability between systems. The problem of interoperability arises almost as soon as a healthcare organization begins to choose a vendor for their electronic medical record, when they are faced with a choice between an architecture based on a single monolithic system or a so-called best-of-breed approach involving multiple discrete systems, each chosen for its superior performance in a narrow domain. The monolith claims to handle all aspects of healthcare information management; the best-of-breed approach entails a multiplicity of systems, each of which may be superior in its domain but which are not smoothly integrated.

A major difference between the two architectures is how they solve the problem of interoperability. In the case of the monolith, the problem is solved by the system vendor, at least in principle, but at the cost to the customer of a loss of choice. In the best-of-breed approach, the problem of interoperability is shifted onto the customer, who incurs an often hefty cost in the form of a more complex systems architecture and the resulting need for specialized hardware, software, and staff to maintain it.

In a best-of-breed approach, the need for instantaneous intersystems communication is typically handled via an Enterprise Service Bus (ESB)[3], which ensures real-time message delivery to subscribing systems. Additionally, if the data is to be analyzed in combination, rather than in isolation within the silo of a single system, it must be recombined and stored outside of these systems. This is typically done in an Enterprise Data Warehouse (EDW)[4] and requires further specialized hardware, software, and staff. However, most EDWs are based on a batch-loading system that runs during off-peak hours for the previous calendar day's business[4]; thus, while an EDW can be a powerful tool for retrospective analysis, it is unsuitable to real-time applications.

Interoperability is a serious challenge that modern healthcare systems must address in order to adequately serve their patients. In this paper we demonstrate a hitherto underused approach that combines the attractive aspects of both an enterprise service bus and an enterprise data warehouse to arrive at real-time analytics.

```
MSH|^~\&|GHH LAB|ELAB-3|GHH OE|BLDG4|200202150930||ORU^R01|CNTRL-3456|P|2.4<cr>
PID||555-44-4444|EVERYWOMAN^EVE^E^^^L|JONES|19620320|F||153 FERNWOOD DR.^
^STATESVILLE^OH^35292||(206)3345232|(206)752-121|||AC555444444||67-A4335^OH^20030520<cr>
OBR|1|845439^GHH OE|1045813^GHH LAB|1554-5^GLUCOSE||200202150730|||F|||
555-55-5555^PRIMARY^PATRICIA P^^^^MD^^^|F|||444-44-4444^HIPPOCRATES^HOWARD H^^^^MD<cr>
OBX|1|SN|1554-5^GLUCOSE^POST 12H CFST:MCNC:PT:SER/PLAS:QN|^182|mg/dl|70-105|H||F<cr>
```

Figure 1: Example HL7 v2.4 Message

Background

Health Level Seven Version 2 (HL7v2)

HL7v2 is a healthcare messaging standard developed by the standards organization Health Level Seven International. It first emerged in 1988 and today is the most widely used such standard, having been adopted by over ninety-five percent of health systems in the United States and thirty-five countries worldwide [5]. As such, it is something of a universal medium in the field of healthcare interoperability, yet it is terse and, without specialized training and access to the standard reference, cryptic.

Each HL7 message describes an event in a healthcare workflow and breaks down hierarchically into segments, fields, components, subcomponents, repeated components, and so on. There are well over one hundred types of messages and several times as many types of segments in HL7v2. The current version of the specification, for HL7 v2.8, is well over 2,500 pages long and contains nearly one million words. [1] Partly as a consequence of this complexity, health interoperability has become a specialized field, replete with certifications and training and entire careers built on knowledge of HL7v2. An example HL7 message describing the following information is shown in Figure 1

- The PID (Patient Identification) segment contains the demographic information of the patient. Eve E. Everywoman was born on 1962-03-20 and lives in Statesville OH. Her patient ID number (presumably assigned to her by the Good Health Hospital) is 555-44-4444.
- The OBR (Observation Request) segment identifies the observation as it was originally ordered: 15545 GLUCOSE. The observation was ordered by Patricia Primary MD and performed by Howard Hippocrates MD.
- The OBX (Observation) segment contains the results of the observation: 182 mg/dl.

Health Level Seven Fast Healthcare Interoperability Resources (HL7 FHIR)

FHIR [6] is a new open standard for healthcare data developed by the same company that developed HL7v2. However, whereas HL7v2 uses an idiosyncratic data exchange format, FHIR uses data exchange formats based on those already in wide use on the World-Wide Web such as Extensible Markup Language (XML) and JavaScript Object Notation (JSON) [7], as well as the web's familiar transfer control protocols such as HyperText Transfer Protocol Secure (HTTPS) and Representational State Transfer (REST) [7] and system of contextual hyperlinks implemented with Uniform Resource Locators / Identifiers (URL/URI) [8]. This design choice simplifies interoperability and discoverability and enables applications to be built rapidly on top of FHIR by the large number of engineers already familiar with web application design without a steep learning curve.

In contrast to HL7v2, which is based on events in a healthcare workflow such as admit, discharge, and transfer, FHIR is built on the notion of conceptual entities from the healthcare domain, such as Patient, Encounter, and Observation, i.e. *resources*. Currently, FHIR encompasses 143 resources, each of which is described abstractly in the FHIR standard with the attributes Name, Flags, Cardinality, Type, and Description & Constraints. [8]. In a concrete implementation of FHIR, resources are serialized to one of the data exchange formats listed above. An example of an FHIR XML message is shown in Figure 2.

```
<Bundle xmlns="http://hl7.org/fhir">
  <id value="b5be9983-c913-4d54-b2bb-8e848056662a"/>
  <type value="message"/>
  <entry>
    <resource>
      <MessageHeader>
        <id value="CNTRL-3456"/> <!-- ControlID of the v2 message -->
        <meta>
          <tag>
            <system value="urn:oid:2.16.840.1.113883.5.100"/>
            <code value="P"/>
            <display value="Production"/>
          </tag>
        </meta>
        <identifier value="CNTRL-3456"/>
        <timestamp value="2002-02-15T09:30:00-04:00"/>
        <event>
          <system value="http://hl7.org/fhir/message-type"/>
          <code value="observation-provide"/>
        </event>
        <source>
          <name value="GHH LAB"/>
          <endpoint value="urn:GHH-LAB"/>
        </source>
        <destination>
          <name value="GHH OE"/>
          <endpoint value="urn:GHH-OE"/>
        </destination>
        <data> <!-- The payload, the resource that this observation-provide concerns -->
          <reference value="DiagnosticReport/1045813"/> <!-- Filler order number -->
        </data>
      </MessageHeader>
    </resource>
  </entry>
```

Figure 2: Example FHIR Bundle and Header Message

Semantic Web

The term Semantic Web [9] denotes an interconnected machine-readable network of information. In some ways it is analogous to the World-Wide Web, but with some crucial differences. The most important similarity is in the vision for the two technologies: Like the World-Wide Web, the Semantic Web was envisioned as a way for users from different institutions, countries, disciplines, etc. to exchange information openly and in doing so to add to the sum of human knowledge. The difference, however, is in the different emphases put on human readability versus machine readability: Whereas the World-Wide Web was intended to be visually rendered by one of any number of web browsers before being read by humans and therefore prioritizes fault tolerance and general compatibility over precision, the semantic web prioritizes precision and logical rigor in order for the information contained in it to be machine readable and used for logical inference.

The similarities continue in the technologies used to implement the two webs. Information in both the Semantic Web and the World-Wide Web is intended to be accessed using the familiar data exchange protocol Hypertext Transfer Protocol (HTTP) and addressed using Uniform Resource Identifiers (URI) for the Semantic Web and Uniform Resource Locations (URL) for the World-Wide Web that tell the agent/browser how to find linked information. Even the data exchange formats are remarkably similar: The World-Wide Web uses Hypertext Markup Language (HTML)[10], a tree-structured subset of Standard Generalized Markup Language (SGML)[11], whereas the Semantic Web uses a variety of tree-structured formats such as XML, JSON, Terse RDF Triple Language (i.e. Turtle/TTL)[12], etc.

The most significant difference between the World-Wide Web and the Semantic Web is in the type of information that they encode. The Semantic Web delivers a payload of simple logical statements known as triples, each consisting of a subject, predicate, and object, whereas the World-Wide Web delivers a series of directives to the web browser that govern the layout of the rendered page as well as the content of the page, in the form of text, images, videos, scripts, and so on. This difference in payloads corresponds to their different purposes – the payload is delivered in the first case to an intelligent agent and in the second case to a web browser.

In more technical terms, the semantic web can be thought of as a distributed directed graph whose vertices are resources and whose edges are statements describing those resources. In its openness and decentralized nature, it bears some resemblance to the World Wide Web; however, whereas the World Wide Web consists of ad hoc, unsynchronized data presented in a variety of formats, the semantic web is a machine-readable body of information that can be synchronized while still coming from a variety of sources.

Resource Description Framework (RDF)

RDF is the backbone of the semantic web[9]. It is described as a framework, rather than a protocol or a standard, because it is an abstract model of information whose stated goal is "to define a mechanism for describing resources that makes no assumptions about a particular application domain, nor defines (a priori) the semantics of any application domain." [13] Its concrete realization is typically a serialization into one of several formats including XML, JSON, TTL, etc.

The basic unit of information in RDF is a *statement* expressed as a logical triple; that is, a statement of the form <subject> <predicate> <object>, in which the predicate expresses a relationship between the subject and the object: for instance, `bloodPressure :value 120`. The subject must be a *resource*, that is, an object consisting of one or more statements, and the object may be either a *literal*, that is, a simple numeric or textual value, or another resource. The predicate describes some aspect or property of the subject. Because both the subject and the object can be resources, the object may also be described by statements in which it is the subject, leading to a complex graph structure.

A group of statements can be used to perform inference on their resources, thus creating new statements and enriching the semantic universe of the data set. For instance, the canonical syllogism "Socrates is a man; all men are mortal; therefore, Socrates is mortal" can be reproduced in the two statements `Socrates :isA man` and `man :is mortal`, resulting in a synthesized third statement: `Socrates :is mortal`. RDF supports "inference, shared semantics across multiple standards and data formats, data integration, semantic data validation, compliance enforcement, SPARQL [SPARQL Protocol and RDF Query Language (SPARQL)] queries and other uses."

[14].

FHIR/RDF

One of the several formats into which FHIR can be serialized is RDF. However, because RDF was designed as an abstract information model and FHIR was designed for operational use in a healthcare setting, there is the potential for a slight mismatch between the models. This comes up in two ways: One, RDF makes statements of fact, whereas FHIR makes records of events. The example given in the FHIR documentation is the difference between "patient x has viral pneumonia" (statement of fact) and "Dr. Jones diagnosed patient x with viral pneumonia" (record of event). Two, RDF is intended to have the property of *monotonicity*, meaning that previous facts cannot be invalidated by new facts. The example given for this mismatch is "a modifier extension indicates that the surrounding element's meaning will likely be misunderstood if the modifier extension is not understood." The potential for serious error resulting from this mismatch is small, but it is worth bearing in mind when designing information systems.

SPARQL Protocol and RDF Query Language (SPARQL)

RDF has an associated query language that can be used to search for matching statements, known as SPARQL. Although syntactically and semantically based on Structured Query Language (SQL), the information model over which it searches is RDF's directed graph of resources and statements, not the familiar relations stored in a relational database.

The syntax is beyond the scope of this paper, but in general SPARQL queries outline the shape of the graph they wish to find. For an example SPARQL query that searches for blood pressure readings over 120 b.p.m., see Figure 3.

Method

At a high level, the semantic enrichment engine is designed to take healthcare data in a variety of formats as input and store it in a triplestore database that users can query. In this way, the engine acts as both a collector, receiving messages from numerous sources, and a bus for delivering messages to multiple sources, as well as a real-time analytics platform. For example, a message from a vital signs monitor and from a registration system can be coalesced into a new stream containing blood pressure, temperature, and laboratory values for use in a machine learning model to predict sepsis.

To support future large-scale operations, a multi-protocol message passing system was used for inter-module communication. This modular design also allows different components to be swapped out seamlessly, provided they continue to communicate via the expected interface. Routines were developed to simulate input data based on the authors' experience with real healthcare data. The reasons for this choice were twofold: One, healthcare data can be high in incidental complexity, requiring one-off code to handle unusual inputs, but not necessarily in such a way as to significantly alter the fundamental engineering choices in a semantic enrichment engine such as this one. Two, healthcare data is strictly regulated, and the process for obtaining access to healthcare data for research can be cumbersome and time-consuming.

A simplified set of input data, in a variety of different formats that occur frequently in a healthcare setting, was used

```

5 // Search for blood pressure readings over 120
6 public class HighBloodPressureQuery extends Query {
7     public HighBloodPressureQuery(String outputTopicName) {
8         super(outputTopicName);
9     }
10
11     sqlQuery = "PREFIX fhir: <http://hl7.org/fhir/>\n" +
12         "PREFIX loinc: <http://loinc.org/rdf#>\n" +
13         "PREFIX owl: <http://www.w3.org/2002/07/owl#>\n" +
14         "PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>\n" +
15         "PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>\n" +
16         "PREFIX sct: <http://snomed.info/id#>\n" +
17         "PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>\n" +
18         "\n" +
19         "SELECT\n" +
20         "  ?quantity_value_value\n" +
21         "  ?observation_component_value_quantity_index\n" +
22         "  ?observation_effective_date_time_value\n" +
23         "  ?observation_subject_reference_reference_value\n" +
24         "WHERE {\n" +
25         "  ?observation a fhir:Observation .\n" +
26         "  ?observation fhir:Observation.component.valueQuantity ?observation_component_value_quantity .\n" +
27         "  ?observation_component_value_quantity fhir:Quantity.system ?quantity_system .\n" +
28         "  ?quantity_system fhir:value \"UOM\" .\n" +
29         "  ?observation_component_value_quantity fhir:Quantity.unit ?quantity_unit .\n" +
30         "  ?quantity_unit fhir:value \"mm[Hg]\" .\n" +
31         "  ?observation_component_value_quantity fhir:Quantity.value ?quantity_value .\n" +
32         "  ?quantity_value fhir:value ?quantity_value_value .\n" +
33         "  ?observation_component_value_quantity fhir:index ?observation_component_value_quantity_index .\n" +
34         "  ?observation fhir:Observation.effectiveDateTime ?observation_effective_date_time .\n" +
35         "  ?observation_effective_date_time fhir:value ?observation_effective_date_time_value .\n" +
36         "  ?observation fhir:Observation.subject ?observation_subject .\n" +
37         "  ?observation_subject fhir:Reference.reference ?observation_subject_reference_reference .\n" +
38         "  ?observation_subject_reference_reference fhir:value ?observation_subject_reference_reference_value .\n" +
39         "  FILTER (xsd:decimal(?quantity_value_value) > 100)\n" +
40         "  FILTER (?observation_effective_date_time_value > \"\" + Utils.secondsAgo(1, Utils.XSD_DATETIME_FMT) + \"\"^^xsd:dateTime)\n" +
41         "}"
42
43
44

```

Figure 3: Example SPARQL Query

for simulation. In a production setting, the Java module that generates simulation data would be replaced by either a data source that directly writes to the input message queue or a Java module that intercepts or extracts production data, transforms it as needed, and writes it to the input message queue. A component-level view of the systems architecture is shown in Figure 4

Class Hierarchy

The project was written in Java, with each major component in its own package. There is a top-level class named `ActiveMQEnabled` that handles common tasks, such as connecting to the message broker, logging, event handling, and other such functionality. Each type of component in the pipeline - input, encoder, store, query, output, and application - is a subclass of `ActiveMQEnabled` as well as a superclass to specific types of those components. Most components are able both to send and receive messages, with certain exceptions: for example, inputs can only send and outputs can only receive. Stores can both receive and send, but in the concrete implementation in this project, the TDB store only receives (queries are better handled as timed polls, rather than being event-driven).

Inputs

In the first stage of the module, simulated inputs represent a variety of healthcare entities and arrive in a variety of formats: patients in a pipe-delimited list, encounters as FHIR messages, and observations as HL7v2 messages. As discussed in the Background section, all of these are widely used input formats in modern health systems and

Figure 4: Semantic Enrichment Engine Architecture

realistically represent the heterogeneous message exchanges that are likely to occur in a real healthcare setting. Each input is configurable with regard to message timing and frequency, and the vitals signs can be made to simulate various conditions such as hypertension or hypothermia. An example of a generate vital sign is shown in Figure 5

Encoder

The encoder stage itself has two stages. In the first, input messages arriving at queues named according to the convention "INPUT.ENTITY.FORMAT" are retrieved, parsed, and transformed into internal representations of common domain objects, in this case Patient, Encounter, and Observation. In the second stage, these internal representations are transformed into internal representations of RDF graphs of FHIR resources and written out to the next message queue. By decoupling the parsing phase from the RDF-generating phase, the number of parsing and generating routines required for N sources and M resource types is reduced from $N \times M$ to $N + M$. This also allows parsing and generating jobs to be written in parallel and by different developers using the common internal representations as an intermediate layer. For instance, one developer could be writing the code to parse an HL7 ADT (admit/discharge/transfer) message while another developer was writing the code to turn this message into Patient, Encounter, and Observation resources. (Note that a single HL7 message can be used to create multiple FHIR resources [15]). An example of a POJO to FHIR/RDF message encoder class is shown in Figure 6

Store

The store stage writes RDF-encoded statements to a triplestore database (TDB). For this implementation, the database was Apache Jena Triplestore Database (TDB) [16], which operates as a local on-disk database, although it could equally be a distributed in-memory cache or other implementation in production. It is at this point that the incoming messages are truly conformed to a universal model, as TDB does not record any information relating to encoding. An example of a RDF to TDB (RDB Database) class is shown in Figure 7

```

132      /*
133       * The OBR segment is contained within a group called ORDER_OBSERVATION,
134       * which is itself in a group called PATIENT_RESULT. These groups are
135       * reached using named accessors.
136       */
137      ORU_R01_ORDER_OBSERVATION orderObservationGroup = patientResultGroup.getORDER_OBSERVATION();
138
139      // Populate the OBR
140      OBR obrSegment = orderObservationGroup.getOBR();
141      obrSegment.getObr1_SetIDOBR().setValue(Utils.randomAlphaIdentifier( count: 10));
142      obrSegment.getObr4_UniversalServiceIdentifier().getCe1_Identifier().setValue("28562-7");
143      obrSegment.getObr4_UniversalServiceIdentifier().getCe2_Text().setValue("Vital Signs");
144      obrSegment.getObr4_UniversalServiceIdentifier().getCe3_NameOfCodingSystem().setValue("http://loinc.org");
145      String localDateTime = DateTimeFormatter.ofPattern( pattern: "yyyyMMddhhmmss", Locale.ENGLISH).format(LocalDateTime.now());
146      obrSegment.getObr7_ObservationDateTime().getTs1_Time().setValue(localDateTime);
147
148      /*
149       * The OBX segment is in a repeating group called OBSERVATION. You can
150       * use a named accessor which takes an index to access a specific
151       * repetition. You can ask for an index which is equal to the
152       * current number of repetitions, and a new repetition will be created.
153       */
154
155      buildVitalSignOBX(message, orderObservationGroup, vitals.getSystolicBloodPressure(), rep: 0);
156      buildVitalSignOBX(message, orderObservationGroup, vitals.getDiastolicBloodPressure(), rep: 1);
157      buildVitalSignOBX(message, orderObservationGroup, vitals.getMeanBloodPressure(), rep: 2);
158      buildVitalSignOBX(message, orderObservationGroup, vitals.getBodyTemperature(), rep: 3);
159      buildVitalSignOBX(message, orderObservationGroup, vitals.getPulseRate(), rep: 4);
160      buildVitalSignOBX(message, orderObservationGroup, vitals.getSP02(), rep: 5);
161
162      // Return the message (remember, the MSH segment was not fully or correctly populated)

```

Figure 5: Java Simulated HL7 Message

Query

The query stage polls the triplestore database for RDF graphs matching specified criteria, for instance, low blood pressure combined with low body temperature and high pulse rate, indicating hypothermia, or patients with blood pressure readings over a certain threshold, indicating hypertension. It passes matching patients on to the output stage for data capture or immediate use in applications.

SPARQL queries against FHIR/RDF (see Figure 3), can often be complex and verbose, simply because a high level of detail was required to represent healthcare data unambiguously in FHIR, and an equally high level of detail was required to extract it unambiguously.

As a means of simplifying the work required to query the data, We considered a two-phase design in which the first layer would extract the relevant data from the TDB database in great detail before using RDF's CONSTRUCT syntax to build a simplified representation of the data for use by the second layer. This idea has potential, but after a few tries at writing the code to implement it, there was too much loss of detail for it to be worth pursuing in this iteration. In the end, the default option of writing a detailed, if verbose, RDF query once was deemed a better option than the added complexity and potential loss of fidelity of the two-layer approach.

Output

In the output stage, the results of the queries in the previous stage are written out to an output destination such as a text file or a screen. This differs from the Application stage in that the input was intended to be written immediately to an output sink such as a file or screen on the local computer. Its use in this project was limited to debugging.

Application

In the application stage, a variety of applications (complex event processors, common data models, machine learning models, etc.) receive the outputs of the queries from the prior stages and use them as inputs to particular applications.


```

44 Model encodePatient(@NotNull Patient patient) {
45     Model model = this.getEncoderModel();
46
47     Resource root = model.getResource(model.expandPrefix( s: "fhir:Patient/" + patient.identifier));
48
49     root.addProperty(
50         model.createProperty(model.expandPrefix( s: "rdf:type")),
51         model.createResource(model.expandPrefix( s: "fhir:Patient")));
52
53     root.addProperty(
54         model.createProperty(model.expandPrefix( s: "fhir:nodeRole")),
55         model.createResource(model.expandPrefix( s: "fhir:treeRoot")));
56
57     root.addProperty(
58         model.createProperty(model.expandPrefix( s: "fhir:Resource.id")),
59         model.createResource()
60             .addProperty(
61                 model.createProperty(model.expandPrefix( s: "fhir:value")),
62                 patient.identifier));
63
64     root.addProperty(
65         model.createProperty(model.expandPrefix( s: "fhir:Patient.identifier")),
66         model.createResource()
67             .addProperty(
68                 model.createProperty(model.expandPrefix( s: "fhir:Identifier.value")),
69                 model.createResource()
70                     .addProperty(
71                         model.createProperty(model.expandPrefix( s: "fhir:value")),
72                         patient.identifier)));

```

Figure 6: POJO to FHIR/RDF Encoder

A high-level view of how the semantic encoder might be used in clinical workflow is shown in Figure 8

Several applications presented themselves as potentially benefiting from a semantic enrichment engine such as this one. One such application was complex event processing (CEP), in which streams of data are analyzed in search of *events* in real time[17]. From simple events more complex events can be derived, so that a number of individually innocuous events may add up to either an opportunity or a threat event. In a healthcare setting, this could mean monitoring patient vital signs and flagging them as high, low, or normal, then analyzing the combination of vital signs for a condition or set of conditions. Additionally, a patient's individual health conditions, such as comorbidities, recent procedures, and so on could be used to inform the meaning of the instantaneous vital signs as they are received. Using data from the TDB store, I was able to write several queries in Esper, a well-known complex event processing engine[18], to detect conditions that were initially simulated by the vital signs input, such as hypothermia or hypertension. To some extent, the RDF queries used to feed Esper overlapped with the capabilities of Esper itself, although Esper's query language EPL is much more versatile than SPARQL for event processing.

Another such project was the Observational Medical Outcomes Partnership (OMOP) Common Data Model (CDM)[19]. This is an analytical database intended to collate data from multiple partner data sources and conform it to a common representation, using standardized vocabularies such as LOINC[20] and SNOMED-CT[21] in order to facilitate collaborative research. Using data queried from the TDB store, I was able to build several data-loading jobs to populate an OMOP-CDM database. This application takes advantage of the semantic enrichment engine's ability to conform data from disparate sources, since by the application stage all the data has been conformed to FHIR/RDF and is ready to be loaded to the OMOP database with only one transformation (from FHIR/RDF to OMOP schemas).

```

13 public class TDBStore extends ActiveMQEnabled {
14     private Dataset dataset = TDBFactory.createDataset( dir: "./tdb/");
15
16     public TDBStore(String inputTopicName) {
17         super(inputTopicName, outputTopicName: null);
18     }
19
20     @Override
21     protected List<String> processInputText(String inputMessageText) {
22         // Open data set transaction in write mode
23         Model model = dataset.getDefaultModel();
24         dataset.begin(ReadWrite.WRITE);
25         try {
26             // Deserialize message to model
27             Reader messageReader = new StringReader(inputMessageText);
28             model.read(messageReader, s: null, s1: "TURTLE");
29             dataset.commit();
30         } catch (Exception e) {
31             e.printStackTrace();
32         } finally {
33             dataset.end();
34             // FUTURE: log the database write
35         }
36
37         return super.processInputText(inputMessageText);
38     }
39 }

```

Figure 7: FIHR/RDF to TDB Storage Class

Validation

Health Level Seven International (HL7) provides a FHIR validator, which was useful for ensuring that the FHIR generated by the encoder was correctly formed. ShEx (Shape Expressions) [22] language is a language used for describing the expected shape of RDF and testing it for conformity to that shape. Its syntax is similar to Turtle and SPARQL, while its semantics resemble those of regular expression languages such as RelaxNG [23]. I were limited in our ability to validate FHIR conformance due to limitations of the FHIR validation tool (vague error messages, program crashes, etc.)

Challenges

Our needs are twofold and, at first, apparently contradictory. The first was to store data from disparate sources so that the sources could be joined up and benefit from synergies among the different semantic components embedded in the data. The second was to answer queries about the data over a finite time range. The challenge is that the mechanism that was to trigger the execution of a query, the receipt of a message from the store, happened with such frequency that the query engine quickly became overloaded and unable to respond in a timely fashion to new requests. This necessitated a redesign of parts of the encoder module and the query engine, such that each resource was timestamped when it was encoded and each query specified a time range within which to return results. Prior to this redesign, the query engine was querying the triple store each time a message arrived without specifying a time bound, resulting in

Figure 8: Semantic Engine Use Clinical Workflow

a constantly increasing number of results that eventually would overmatch the system's capabilities.

Another challenge was that RDF does not easily support streams[24]. With each query, all matching results are returned, not only the new results since the last query. This means the result size of the query increases monotonically until the system is overwhelmed. To design around this, we timestamped each entity as it arrived and used this timestamp as a filter in the subsequent queries. This worked well and is not unlike what CEP systems do natively.

Conclusion

The semantic enrichment engine designed described in this paper has broad applicability in healthcare operations and research. The data exchange standards, protocols, databases, query languages, and so forth used to implement this system are freely available. This system has characteristics of both an enterprise service bus and an enterprise data warehouse, but augments the analytical capability of the former and addresses the high latency of the former. We expect the system can be used to inform artificial intelligence for inference, populate structured databases with enriched data streams, and derive new data for use in machine learning training.

References

1. Graham Benson, Tim; Grieve. *Principles of Health Interoperability*. Springer-Verlag, 2016.
2. Ashish K. Adler-Milstein, Julia; Jha. Hitech act drove large gains in hospital electronic health record adoption. *Health Affairs*, 2017.
3. <https://www.ibm.com/cloud/learn/esb>. Accessed: 7/21/2019.
4. Margy Kimball, Ralph; Ross. *The Data Warehouse Toolkit, 3rd Ed*. Wiley, 2013.
5. https://www.hl7.org/implement/standards/product_brief.cfm?product_id=185. Accessed: 8/22/2019.
6. <https://www.hl7.org/fhir/>. Accessed: 6/28/2019.
7. <https://www.hl7.org/fhir/overview-dev.html>.
8. <https://www.hl7.org/fhir/resourcelist.html>. Accessed: 6/11/2019.
9. Jim Allemang, Dean; Hendler. *Semantic Web for the Working Ontologist*. Morgan Kaufman, 2011.
10. <https://www.w3.org/html/>. Accessed: 6/30/2019.

11. <https://www.w3.org/TR/REC-html40-971218/conform.html>. Accessed: 6/30/2019.
12. <https://www.w3.org/TR/turtle/>. Accessed: 6/28/2019.
13. <https://www.w3.org/TR/PR-rdf-syntax/>. Accessed: 8/1/2019.
14. <https://www.hl7.org/fhir/rdf.html>. Accessed: 3/15/2019.
15. <https://fhirblog.com/2015/03/18/more-fhir-messaging-adt-messages>. Accessed: 4/21/2019.
16. <https://jena.apache.org/documentation/tdb/index.html>. Accessed: 7/14/2019.
17. David C. Luckham. *Event Processing for Business: Organizing the Real-Time Enterprise*. John Wiley & Sons, Inc., 2012.
18. <https://esper.espertech.com/>. Accessed: 7/15/2019.
19. <https://www.ohdsi.org/data-standardization/the-common-data-model/>". Accessed: 8/27/2019.
20. <https://loinc.org/>. Accessed: 6/15/2019.
21. <https://www.snomed.org>. Accessed: 6/15/2019.
22. <http://shex.io/shex-semantic-20170327/>. Accessed: 7/1/2019.
23. <https://relaxng.org/>. Accessed: 7/1/2019.
24. <https://lists.w3.org/Archives/Public/public-lod/2015Jan/0190.html>. Accessed: 8/6/2019.