Workshop on Semantic Cities: Beyond Open Data to Models, Standards and Reasoning, AAAI14, Quebec City, 28 July 2014.

An Ontology for Open 311 Data

Soroosh Nalchigar, Mark S. Fox
Department of Computer Science, University of Toronto
Department of Mechanical and Industrial Engineering, University of Toronto
soroosh@cs.toronto.edu, msf@eil.utoronto.ca

Abstract

A major challenge in the analysis of city data is the integration of data from different sources. This paper defines an ontology, called Open 311 Ontology, that provides a unified terminology and a reference model for representing the 311 data. We illustrate how the ontology can be used to map and integrate data from multiple cities, and for answering competency questions.

1. Introduction

A fundamental aspect of a smart city is to integrate and combine the data coming from various sources and places. Data integration is a challenging task, partially due to differences in the schema and content of data sources. Consider the existing 311 data of cities as an example. In order to combine these data for analysis, there are some issues that have to be tackled in advance. The first one is to find and map equivalent attributes between existing datasets. Is the attribute “Responsible Agency” in San Francisco’s dataset equivalent to Toronto’s “Division”, “Section-unit”, or both? Is Toronto’s “Service Request Name” equivalent to “Request Type” in San Francisco’s dataset? The second issue is to define a mapping between values of equivalent attributes. Provided that Toronto’s “Service Request Name” is equivalent to San Francisco’s “Request Type”, how the values of these attributes could be integrated and unified? Is San Francisco’s “Sign Repair” equivalent to Toronto’s “Sign Maintenance” or “Missing/Damaged Signs” or both? Without finding and mapping equivalent attributes and values across 311 datasets, it is not possible to integrate, merge, and analyze the data (Fox, 2013). To make these issues more clear, we note that each cities’ dataset have a different number of attributes. Even worse, that Toronto’s 311 is using 371 different names for describing the service request types. It is while New York, San Francisco, and Chicago are using 120 and 25, and 12 different names (unique within each dataset) for representing service requests, respectively.

In this paper, we focus on open 311 datasets. 311 is the name and the telephone number of city agencies that provide non-emergency municipal services to the public. The main goal of 311 systems is to enhance accessibility of city services, increase cities effectiveness in responding to public inquiries, and hence to improve city life. In this paper, our focus is enabling longitudinal analysis of cities’ 311 data (i.e., changes over time for a single city), and transversal analysis (i.e., comparison of two or more cities). To tackle this problem, our approach is to develop a formal 311 foundation ontology, referred to as *Open 311 Ontology*, with which specific city's 311 data models and instances can be defined. The ontology is aimed to provide a unified and complete terminology that could be utilized for integration of existing open data applications, enabling city data analytics, and hence facilitating the current movements towards smart and data-driven cites.

The rest of this paper is organized on the basis of the ontology development methodology defined in (Gruninger & Fox, 1995). Section 2 provides an overview of the existing datasets of four cities and describe their data schemas. Section 3.1 presents two motivating scenarios to highlight the need for the ontology and results in a set of competency questions in Section 3.2. The main concepts and properties of the open 311 ontology are presented in Section 3.3. Thereafter, Section 3.4 describes other existing ontologies that are related and used in design of Open 311 Ontology. In Section 3.5 the main classes of the ontology are represented through formal axioms, in order to make the definitions and restrictions clearer and also to enable representation of the schema in a formal language. Finally, in Section 4, the ontology is evaluated by checking its ability to answer competency questions and also by illustrating the possibility of mapping and representing existing data to it.

2. Analysis of Published 311 Data

In this paper we use datasets of 311 departments of four cities, namely Toronto, New York, San Francisco, and Chicago. To choose these cities, we considered factors such as availability of 311 data as well as existence of enough instances of service requests for understanding the domain and creating the unified terminology. This section describes the datasets of each of the cities.

Toronto. This dataset[1] includes 6 fields. **Service Request Name** is the unique title of an individual service request. **Problem Code** is a unique identifier of the service request names. **Creation Date** indicates the date and time in which the corresponding request instance is submitted to 311. The fields **Division** and **Section-Unit** represent the responsible City division and the 311's section or unit under which the service request is listed. Finally, **Internet Self Serve** shows if the service request is reported via the web. Table 1 shows a service request record in this dataset.

Creation Date	04-26-2010 16:02:36
Service Request Name	Residential: Garbage Bin: Exchange to Medium
Division	Solid Waste Management Services
Section – Unit	Collections
Problem Code	SWBNMTC-26
Internet Self Serve	Yes

Table 1: A service request record in Toronto's 311 dataset

San Francisco. Dataset[2] of this city includes 15 fields. In this dataset, fields such as **Category** and **Responsible Agency** are equivalent to the fields Service Request Name and Section–Unit, respectively, in Toronto's dataset. Clearly this dataset has more fields than Toronto, such as **Status, Address,** and **Point** (latitude and longitude coordinates). Table 2 shows an example of a service request record in this dataset.

New York. This dataset[3] includes 52 fields. **Created Date**, **Closed Date,** and **Agency** are equivalent to Opened, Closed, and Responsible agency attributes form San Francisco's dataset, respectively. Other fields such as **Complaint Type**, **Latitude,** and **Longitude** have obvious equivalents, but with a different name, in the San Francisco dataset. Some of the fields that are appearing only in this dataset are **Due Date**, **Facility Type**, **Cross Street**. It should be noted this dataset has some attributes whose value is unspecified, NA, or missing value for the whole dataset. Those fields are not considered in design of the ontology.

[1] http://www.toronto.ca/311
[2] http://data.sfgov.org
[3] http://nycopendata.socrata.com

CaseID	2441829
Opened	06-03-2013
Closed	06-03-2013
Status	Closed
Work Status	New
Responsible Agency	311 Supervisor Queue
Address	2329 Castro St, San Francisco, CA, 94131
Category	Street and Sidewalk Cleaning
Request Type	Sidewalk_Cleaning
Request Details	Furniture
Source	Voice In
Supervisor District	9
Neighborhood	Inner Mission
Updated	06-03-2013 6:49
Point	(37.750540724, -122.419933447)

Table 2: A service request record in San Francisco's 311 dataset

Chicago: Data of this city[4] has 15 fields and are provided in separate files, where each file includes requests of a specific type (e.g., tree debris, garbage carts, etc.). All the fields in this dataset have an equivalent field in either or both San Francisco and New York datasets. However, they a different name. For example the **Completion Date** here is equivalent to the Closed field in San Francisco.

Due to page constraints, we do not present an example of New York and Chicago datasets here. Interested readers are referred to the URL of the dataset given in footnote.

3. Open 311 Ontology

3.1 Usage Scenarios

In order to illustrate and motivate the need for Open 311 Ontology, this section provides two hypothetical use case scenarios. These scenarios are later used to indicate how the ontology would be helpful in these cases.

Customer inquiries. The contact center of the city 311 gets numerous calls from customers who have inquiries about their previously reported service requests. Usually, the customers call to check the status of their request and to get updates on that, having the unique reference number of their submitted service request. To answer those inquiries, the contact center needs to access the stored data of service requests. To this end, the city 311 needs to keep

[4] http://data.cityofchicago.org

records of the date and time in which the request was submitted as well as its latest status (open, closed, etc.).

Performance management. Every day, the 311 center of the city receives thousands of service requests from the crowd, through various channels such as email, smart phone apps, and phone calls. The board of directors in the city 311 understands that in the current rapidly changing business environment, getting insights from raw data and making data-driven decisions are of great importance. Towards these ends, the city 311 has developed a standard reporting system that addresses the information needs of boards of directors and answers their business questions. Among others, the board wants to know what the busiest 311 agencies are, i.e., which agencies are receiving highest number of service request. This information would help them to assign more employees to busy agencies, balance the workload, and hence reduce the time it takes to address the requests. Also, each service request is about a different subject, e.g., garbage bins, graffiti, roads, etc. The board of directors is interested to know what the most reported service topics are. These will help them in aggregating messages arising from the crowd and use it to gain insights about the city problems. Beside these reports, the 311 board of directors is interested in comparisons and cross-city analyses. They like to know how other cities are different from them in term of environmental pollutions and crime. In particular, they like to know which cities are having more reports about dead animals as well as reports about illegal issues. In order to generate these reports, the city 311 needs accurate, relevant and timely data, which often is hard to achieve.

3.2 Competency Questions

Competency questions are essential for evaluation of an ontology (Gruninger & Fox, 1995). Based on usage scenarios, we define a set of competency questions in three categories. The first category examines the ability of the ontology in property value retrieval:

QC-1: What is the submission date of a given service request with the unique code "XYZ"?
QC-2: What is the status of a given service request with the unique code "XYZ"?

The second category of competency questions focus on aggregation type of questions and includes:

QC-3: What are top five busiest 311 agencies in terms of number of submitted service requests?
QC-4: How many service requests about "Subject1" are reported since the beginning of the year?

Finally, the last category of competency questions focuses on cross city comparison type of questions and contains:

QC-5: Which cities have more than 1000 reports categorized as "illegal issues"?
QC-6: What are top three cities with most number of reports of the subject "dead animals"?

3.3 Concepts and Properties

In this section we illustrate the construction of Open 311 Ontology and explain the primitive concepts as well as object and data properties of the ontology[5]. This ontology is expressed in OWL language and is implemented in the Protégé ontology editor. At the core of the Open 311 Ontology is the class ServiceRequest, which is the class of all service requests submitted to the city 311. A ServiceRequest contains following data properties: AddressType, Borough, CloseDate, CommunityBoard, CrossStreet, Details, DueDate, EventID, EventZip, Intersection, LocationType, Neighborhood, OpenDate, Source, Status, UpdateDate, Ward.

Along with these data properties, the class ServiceRequest has following object properties:

- **Has311Type:** whose range is the class of 311Type and identifies the category of the service request.
- **isHandledBy:** whose range is the class Agency, represent the 311 agency that handles the service request.
- **isSubmittedTo:** whose range is the class Division, showing the 311 responsible division to which the service request submitted.
- **hasSPS:** whose range is the class SpsPoint, identifying the exact location of the service requests.

Another important class in our ontology is 311Type which, as the name suggests, represents the type to which each instance of ServiceRequest belong. As we explained in the Introduction, one of the main issues in integration of open 311 data is that each city has its own vocabulary for describing the service request types. For example a service request about a damaged street signs is recorded as "Sign Maintenance" in Toronto dataset, while it appears as "Sign Repair" and "Street Sign - Damaged" in San Francisco and New York datasets, respectively. To tackle this problem and to propose a unified way of representing problem types, following object properties along with their range classes are defined to be connected to the class 311Type:

- **has311Subject:** whose range is the class 311Subject and presents what the corresponding 311Type is about. The subclasses of the class 311Subject include but not limited to TransportationRoutes, RoadSybmol, GarbageContainer, etc. Each of these has its own subclasses.
- **Need311Action:** whose range is the class 311Action and represent the action that the Agency needs to undertake in response to the ServiceRequest. The class

[5] Available at: `http://ontology.eil.utoronto.ca/o311o.owl`

311Action has subclasses such as Replace, Repair, Remove, Reinstall, Install, Inspect, etc.

- **has311MessageCategory:** whose range is the class 311MessageCategory and includes subclasses such as Compliant, Report, Inquiry, etc.

These definitions allow each city 311 to connect its own service request types with their specific naming values to our ontology and thereafter populate it with real instances of service requests. In this way, our ontology facilitates integration of data across various cities and supports reasoning about service requests as well as querying, and analysis of the integrated data.

3.4 Related Ontologies

We build the Open 311 Ontology "one brick at a time" using some existing foundational ontologies such as Organization Ontology, Time Ontology, and GeoNames Ontology. In the next subsections we briefly introduce these ontologies and indicate how they are related to Open 311 Ontology.

Organization Ontology. Organization ontology[6], defined by Fox et al. (1996), focuses on organization structure, roles, authority and empowerment. It is developed as part of the TOVE Project (Fox, 1992) and supports reasoning in industrial environments. One of the core classes in this ontology is Organization, defined as a set of constraints on the activities performed by agents. This class contains following data and object properties: hasName, hasGoal, and consistsOf. Figure 1 depicts how the Open 311 Ontology is related to this ontology. In this figure we specialize the class Organization to the classes Agency. This will allows our classes Agency to inherit the properties of the Organization as defined in Organization ontology, e.g., hasName.

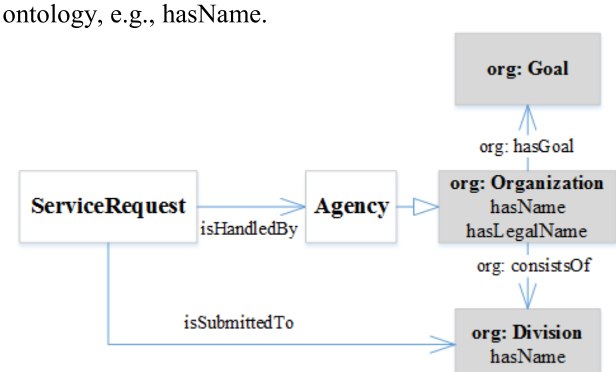

Figure 3: Open 311 Ontology in relation to Organization Ontology.

Placename Ontology. The service requests submitted to 311 are associated with a geographic area, which would be a city and country. Therefore, a requirement for the Open 311 Ontology is the ability to represent the geographic area to which the service request is related. The Schema.org[7] ontology provides classes such as sc:Place, sc:City, and sc:Country and includes properties such as sc:address, sc:map, and sc:review. The GeoName geographical database includes over 8.3 million placenames covering all countries. Beyond names of places in various languages, this database integrates geographical data such as latitude, longitude, population and postal codes from various sources. All the placenemase are instantiations of the GeoNames Ontology[8] which integrates numbers of ontologies including Schema.org. The most fundamental class in GeoNames Ontology is the class gn:feature which, among others, includes the following properties: name, altenativeName, countryCode, and population.

Figure 2 shows how the Open 311 Ontology is connected to GeoNames and Schema.org ontologies. This figure indicates that the object property hasCity connects the class ServiceRequest to the class sc:City which inherits all the properties of the class gn:Feature.

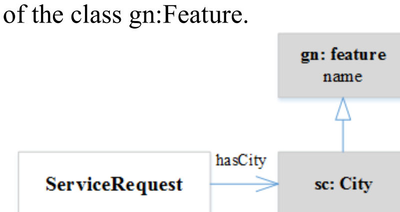

Figure 1: Open 311 Ontology in Relation to GeoNames and Schema.org ontologies

International Contacts Ontology. The current data of service requests includes the address for which the request is made. The address usually includes number, street name, as well as the postal code. Hence, the Open 311 Ontology requires representing the address. This will allow the ontology to be refined and better represent the location of the service request. International Contact (iContact) Ontology[9] provides basic classes and properties for the representation of street addresses, phone numbers and emails. One of the important classes in this ontology is ic:Address that includes following properties: hasStreet, hasUnitNumber, hasPostalCode, hasStreetDirection, and

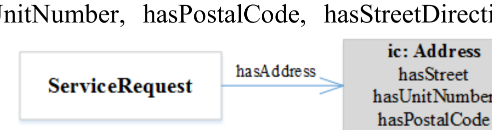

Figure 2: Open 311 Ontology in Relation to iContact Ontology

[6] This ontology is available at `http://ontology.eil.utoronto.ca/organization.owl`. In this paper, the prefix "`org:`" is used to show the classes as well as data and object properties of this ontology.

[7] This ontology is available at `http://schema.org`. In this paper, the prefix "`sc:`" is used to show the classes as well as data and object properties of this ontology.

[8] This ontology is available at `http://www.geonames.org/ontology/ontology_v3.1.rdf`. In this paper, the prefix "`gn:`" is used to show the classes as well as data and object properties of this ontology.

[9] This ontology is available at `http://ontology.eil.utoronto.ca/icontact.owl`. In this paper, the prefix "`ic:`" is used to show the classes as well as data and object properties of this ontology.

hasStreetType. Figure 3 shows how the Open 311 Ontology is related to iContact Ontology.

Time Ontology. The service requests are associated with a set of temporal information such as the submission date, the closing date, etc. It is important for Open 311 Ontology to capture and represent the time attributes of service requests. To do that, we use the Time Ontology[10] for representing temporal properties of service requests. Time Ontology provides a standard set of classes and relations for representing facts about topological relations among instants and intervals, as well as information about durations and datetime information. One of the main classes in this ontology is DateTimeInterval that is connected to the class DateTimeDescription through the object property hasDatetimeDescription. The class DateTimeDescription includes various data properties such as second, minute, hour, day, month, year, etc. Figure 4 illustrates that the class ServiceRequest form Open 311 Ontology is connected to Time Ontology through four different object properties, namely hasOpenDate, hasClsoeDate, hasUpdateDate, and hasDueDate.

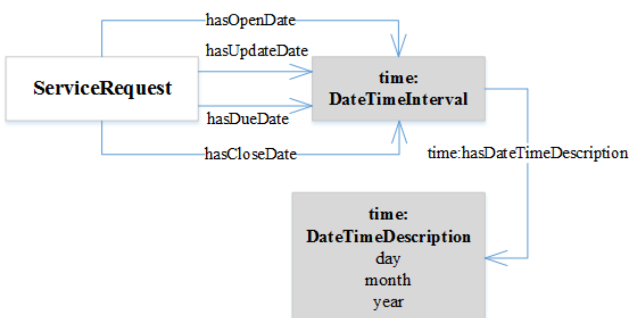

Figure 5: Open 311 Ontology in Relation to Time Ontology

Other related ontologies. In the Open 311 Ontology, the class ServiceRequest is connected to the class 311Type via the object property hasType. The class 311Type is connected to the class 311Subject through the object property has311Subject (See the section Concepts and Properties). One of the subclasses of the class subject is TransportationRoutes, meaning that a service request could be about a transportation route such as an expressway. In order to provide enough expressivity, we needed to define subclasses of the TransportationRoutes. Several ontologies exist in the literature that are constructed to solve specific problems regarding cities and urban areas, e.g., Towntology ontology (Keita et al., 2004) and CityGML ontology (Gröger et al., 2007). Although have been created with a specific task in mind, they could be used to identify some of the subclasses of TransportationRoutes. Figure 5 shows how our ontology is connected to the Towntology, CityGML, and DBpedia ontologies.

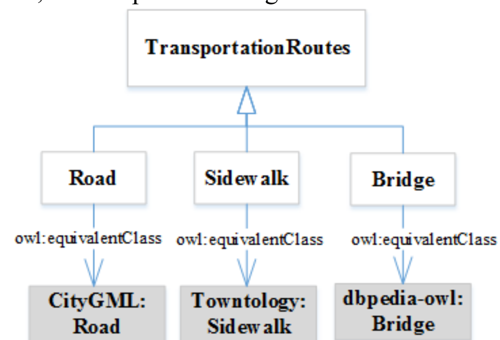

Figure 4: Equivalent Classes in Other Ontologies

It should be noted that within Open 311 Ontology, the class TransportationRoutes has other subclasses that were identified by careful review of 311 city datasets, e.g., Expresseway, Boulevard. Moreover, there are other classes in our ontology, such as Plants, Animal&Insects that are connected to some other ontologies, but they are not explained here due to lack of space. Interested readers are referred the OWL file of our ontology for further details[11].

3.5 Axioms

In this section, using Description Logic (DL), we present the axioms that formalize the ontology and support the reasoning about the classes. Following formulation represent the definition of a ServiceRequest[12]:

```
ServiceRequest ≡
  311Thing ⊓
  =1 has311Type.311Type ⊓
  =1 isHandledBy.Agency ⊓
  =1 hasAddress.Address ⊓
  =1 hasCity.City ⊓
  =1 hasOpenDate.DateTimeInterval ⊓
  ≤1 hasCloseDate.DateTimeInterval ⊓
  ≤1 hasDueDate.DateTimeInterval ⊓
  ≥0 hasUpdateDate.DateTimeInterval ⊓
  =1 EventID.string ⊓=1 Source.string ⊓
  =1 Status.string ⊓≤1 AddressType.string ⊓
  ≤1 Borough.string ⊓≤1 CommunityBoard.string ⊓
  ≤2 CrossStreet.string ⊓ ≤1 Deatils.string ⊓
  ≤1 Intersection.string ⊓ ≤1 LocationType.string ⊓
  ≤1 Neighborhood.string ⊓ ≤1 Ward.string
```

Moreover, the class 311Type is defined in terms of following formulation:

```
311Type ≡
    311Thing  ⊓
    ≥1 has311Subject.311Subject ⊓
    ≥1 need311Action.311Action ⊓
    =1 has311MessageCategory.311MessageCategory ⊓
    =1 311TypeCode.string ⊓
    =1 311TypeName.string
```

---

[10] This ontology is available at `http://www.w3.org/2006/time`. In this paper, the prefix "`time:`" is used to show the classes as well as data and object properties of this ontology.

[11] Available at: `http://ontology.eil.utoronto.ca/o311o.owl`

[12] To represent the "exactly one" cardinality in these formulations, we contract the ≥1 and ≤1 constructors to =1, due to space limitations.

4. Evaluation

This sections evaluates the ontology in two parts. The first part evaluates the ability of ontology to represent the data that is needed to answer the competency questions of the Section 3.2. The second part evaluates the ontology by illustrating how data of each city are mapped and represented in the ontology.

4.1 Answering the Competency Questions

This Section presents the competency questions and shows how the SPARQL query language (Prud'Hommeaux & Seaborne, 2008) could be used to retrieve the relevant data from the ontology and to answer the questions[13].

QC-1: What is the submission date of a given service request with the unique code "XYZ"?

In order to answer the first competency question, we need to retrieve the date in which the given service request was submitted to the city 311. Following query finds the answer:

```
SELECT ?day ?month ?year
WHERE{
  ?ServiceRequest o311o:EventID "XYZ".
  ?ServiceRequest o311o:hasOpenDate ?DTInterval.
  ?DTInterval time:hasDateTimeDescription ?DTD.
  ?DTD time:day ?day.
  ?DTD time:month ?month.
  ?DTD time:year ?year
 }
```

In our ontology, the ServiceRequest class is connected to the class DateTimeInterval (imported from Time Ontology) via the object property hasOpenDate. In Time Ontology, the DateTimeInterval class is connected to the class DateTimeDescription through the object property hasDateTimeDescription. The data that is required to answer the first competency question are represented as properties of the class DateTimeDescription.

QC-2: What is the status of a given service request with the unique code "XYZ"?

Following query answers the question:

```
SELECT ?status
WHERE {
  ?ServiceRequest o311o:EventID "XYZ".
  ?ServiceRequest o311o:Status ?status
}
```

In the ontology, the class ServieRequest has the data property of Status whose value is of type string. The answer to the second competency question could be obtained from this data property.

QC-3: What are top five busiest 311 agencies in terms of number of received service requests?

The answer to the third competency question is obtained by following SPARQL query:

```
SELECT ?Name (COUNT (?ServiceRequest) AS ?Total)
WHERE{
  ?ServiceReuqest o311o:isHandledBy ?Agency.
  ?Agency org:hasName ?Name
}
GROUP BY ?Name
ORDER BY ?Total
LIMIT 5
```

In our ontology, the object property isHandledBy connects the class ServiceRequest to the class Agency. The class Agency has the data property of hasName which is a unique string representing the name of agency that handles the service request. In order to compute the answer to the forth competency question, this query counts total number of service requests that are submitted to the city agencies. Then, by ordering and finding the top 5 instances of the class Agency, the answer to the third competency question is found.

QC-4: How many service requests about "Subject1" are reported since the beginning of the year?

Regarding the forth competency question, we need to retrieve and count service requests of the given subject that are reported in the current year. To do that, following SPARQL query is used:

```
SELECT (COUNT(?ServiceRequest) AS ?Total)
WHERE{
 ?ServiceRequest o311o:has311Type ?311Type.
 ?311Type o311o:has311Subject ?Subject.
 ?Subject a o311o:"Subject1".
 ?ServiceRequest o311o:hasOpenDate  ?DateTimeInterval.
 ?DateTimeInterval time:hasDateTimeDescription ?DTD.
 ?DTD time:year ?Year.
 FILTER (?Year == 2014)
}
```

In our ontology, each instance of the class ServiceRequest is associated with its 311Type. Moreover the class ServiceRequest is connected to the class DateTiemInterval from Time Ontology, to keep the time information in which a request is submitted. Within Time Ontology, the class DateTimeInterval is connected to the class DateTimeDescription through the object property hasDateTimeDescription. The first step in answering this competency question is to retrieve the set of all instances of the class ServiceRequest whose 311Type instance has the subject that is the given in the competency questions. Having this set, the next step is to exclude those service requests which are not submitted in the current year and to count total number of instances that are remained.

QC-5: Which cities 311 has received more than 1000 reports categorized as illegal issues?

Following query computes the answer to this competency question:

```
SELECT ?City (COUNT (?ServiceReuqest) AS ?Total)
WHERE{
  ?ServiceReuqest o311o:hasCity ?City.
  ?ServiceReuqest o311o:has311Type ?311Type.
  ?311Type o311o:has311MessageCategory ?Category.
  ?Category a o311o:IllegalIssue
}
GROUP BY ?City
HAVING COUNT (?ServiceReuqest) > 1000
```

---

[13] All the queries assume that the namespace prefix `o311o` refers to the IRI `http://ontology.eil.utoronto.ca/o311o.owl`. Also, it is assumed that the prefix `time` refers to the IRI `http://www.w3.org/2006/time`.

In our ontology, the class 311Type is connected to the class 311MessageCategory via the object property has311MessageCategory. The class 311MessageCategory has various subclasses one of which is Complaint which has IllegalIssue as a subclass. To compute the answer of this competency question, for each city, all the instances of ServiceRequest whose category is illegal issue are retrieved and counted. This will result in a list of cities along with their corresponding number of illegal issue reports. The last step is to exclude those cities which have less than 1000 service requests of the specified category.

QC-6: What are top three cities with most number of reports of the subject "dead animals"?

The answer to this question results from following query:

```
SELECT ?City (COUNT (?ServiceRequest) AS ?Total)
WHERE{
  ?ServiceReuqest o311o:hasCity ?City.
  ?ServiceReuqest o311o:has311Type ?311Type.
 ?311Type o311o:has311Subject ?Subject.
  ?Subject a o311o:DeadAnimal
}
GROUP BY ?City
ORDER BY ?Total
LIMIT 3
```

In the ontology, the class 311Type is connected to the class 311Subject via the object property has311Subject. The class 311Subject has various subclasses one of which is Pests. The class Pests has two subclasses, namely Animal and Insects. The DeadAniumal class is a subclass of the class Animal. Similar procedure to previous question is used here to answer this competency question. The only difference is that in this question we look for ServiceRequest instances that are connected to the DeadAnimal class via the object property has311Subject.

4.2 Mapping Datasets to the Open 311 Ontology

In this section, we illustrate the possibility of mapping/representing existing datasets to/in the Open 311 Ontology. Figures 6 and 7 show how Tables 1 and 2 (See Section 2) are represented in the Open 311 Ontology. It should be mentioned that we also have mapped data sets of

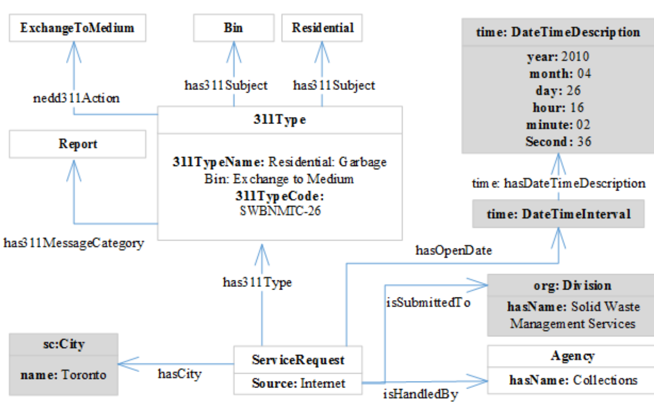

Figure 6: Mapping Toronto's Dataset to Open 311 Ontology

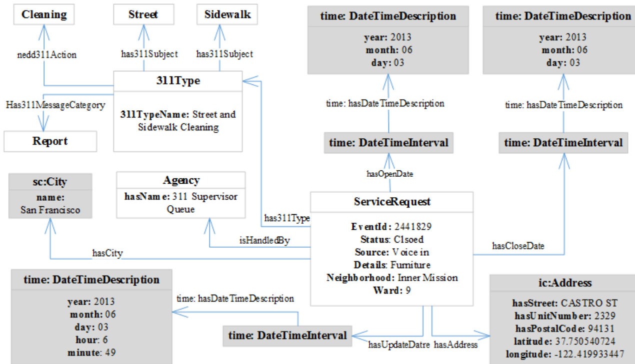

Figure 7: Mapping San Francisco's Dataset to Open 311

the cities New York and Chicago to our ontology. However, because of space limitations the examples are not presented here.

5. Conclusion

This paper describes the Open 311 Ontology. Existing representations of 311 data, as found in the four cities sampled, differ in the both their data models and their content. By providing an ontology for 311 data, we make it possible to merge 311 data from across cities, thereby enabling transversal data analytics. This research is part of a broader research agenda whose long term goal is to transform and improve how people participate in the organizational decision making. The next steps would be to use the Open 311 Ontology for aggregating messages that comes from the crowd and infer the city knowledge from that. Also, to perform crowd-based reasoning for inferring solutions to city problems is another interesting venue.

References

Fox, M. S. 1992. The TOVE project towards a common-sense model of the enterprise. *Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*. Springer Berlin Heidelberg, 25-34.

Fox, M. S. 2013. City Data: Big, Open and Linked. Working Paper, Enterprise Integration Laboratory, University of Toronto.

Fox, M. S., Barbuceanu, M., and Gruninger, M. 1996. An organisation ontology for enterprise modeling: preliminary concepts for linking structure and behaviour. *Computers in industry* 29.1: 123-134.

Gruninger, M., Fox, M. S. 1995. Methodology for the design and evaluation of ontologies, In *Proceedings of the Workshop on Basic Ontological Issues in Knowledge Sharing at IJCAI*.

Gröger, G., Kolbe, T., and Czerwinski, A. 2007. Candidate OpenGIS® CityGML Implementation Specification (City Geography Markup Language). *Open Geospatial Consortium Inc, OGC*.

Keita, A., Laurini, R., Roussey, C., and Zimmerman, M. 2004. Towards an ontology for urban planning: The towntology project. In *CD-ROM Proceedings of the 24th UDMS Symposium, Chioggia* (Vol. 12, No. 1).

Prud'Hommeaux, E., and Seaborne, A. 2008. SPARQL query language for RDF. *W3C recommendation*, *15*.