A Foundation Ontology for Global City Indicators

Mark S. Fox
Mechanical and Industrial Engineering
University of Toronto
msf@eil.utoronto.ca

Abstract

This paper defines a generic, reusable ontology that forms a foundation for representing the definitions of ISO37120 Global City Indicators and their instantiation on the Semantic Web by cities. It makes two contributions. The first is to the field of city indicators. The ontology makes it possible to:

- create precise definitions of indicators thereby reducing the ambiguity of their interpretation and making them computationally accessible;
- represent the "data behind the data" or supporting data, enabling drilling down; and
- achieve interoperability, namely the ability to access, understand, merge and use measurement data available from datasets spread across the Semantic Web by providing a standard representation using OWL 2.

The second contribution is the successful integration of seven basic ontologies and their extensions. The ontology spans analytical models (e.g., ratios), statistical models (e.g., population measurements), geo-spatial models (e.g., city boundaries), temporal models (e.g., time periods) and description logic models (e.g., definitions of students, teachers, etc.). It also integrates meta-knowledge such as provenance, validity and trust.

Keywords: City Indicator, Ontology, Semantic Web, ISO 37120.

1. Introduction

Cities are moving towards policy-making based on data[1]. In 2007, the World Bank (Hoornweg et al., 2007) recognized that "there are thousands of different sets of city (or urban) indicators and hundreds of agencies compiling and reviewing them. Most cities already have some degree of performance measurement in place. However, these indicators are usually not standardized, consistent or comparable (over time or across cities), nor do they have sufficient endorsement to be used as ongoing benchmarks." In other words, there does not exist an agreed upon set of metrics that are consistently applied across cities, thereby rendering existing comparisons of cities meaningless.

[1] "Data driven decision making is one of the reasons New York City is the safest big city in America," said Mayor Bloomberg. "Just as data helps us reduce crime, prevent fire fatalities and keep incarceration levels low, we believe understanding data can help us work with judges and criminal justice agencies to further improve the effectiveness and efficiency of our criminal justice system." Press Release, New York City, PR-012-13, 7 January 2013.

In response to this challenge, the international standard ISO 37120 “Sustainable development of communities — Indicators for city services and quality of life” was published in May, 2014. The ISO standard contains 100 indicators and has been adopted by over 250 cities across the world as a basis of reporting, comparing and understanding how well they perform.

The primary problem with the ISO standard is that definitions are provided in English, and not in a more formal, computational language. The lack of formal definitions allows the reader to impose their own interpretation based on their understanding of the language and the environment in which they reside (i.e., how their own city may define some terms).

Consider the definition of a Student/teacher ratio as provided by the World Bank (2008, p. 18):

> “Student/teacher ratio”, where the numerator is “Number of Students”, and the denominator is “Number of Teachers”.

There are many problems with this brief definition. Does “student” refers to full time students, or part time students? Are they regular students or special needs students? Do they include kindergarten students or not? It is difficult to compare an indicator for a single city across time if the definition of student changes. For example, today the educational system includes students with special needs, but 30 years ago they may not have been enrolled. Without a more precise definition of terms, it makes it difficult to compare an indicator across cities where each city interprets what a student is differently, or against a city itself when definitions change.

Obviously, the definition and documentation of indicators can be expanded, as has been done in ISO 37120, in recognition of the limited value of prior indicator definitions. Following is the definition of student teacher ratio provided by the standard:

> "The student/teacher ratio shall be expressed as the number of enrolled primary school students (numerator) divided by the number of full-time equivalent primary school classroom teachers (denominator). The result shall be expressed as the number of students per teacher. Private educational facilities shall not be included in the student/teacher ratio. One part-time student enrolment shall be counted as one full-time enrolment; in other words a student who attends school for half a day should be counted as a full-time enrolment. If a city reports full-time equivalent (FTE) enrolment (where two half day students equal one full student enrolment), this shall be noted. The number of classroom teachers and other instructional staff (e.g. teachers’ aides, guidance counselors) shall not include administrators or other non-teaching staff. Kindergarten or preschool teachers and staff shall not be included. The number of teachers shall be counted in fifth time increments, for example, a teacher working one day per week should be counted as 0.2 teachers, and a teacher working three days per week should be counted as 0.6 teachers.”

The revised definition of student teacher ratio clearly addresses some of the issues raised earlier. Never the less, there will always be a disconnect between the actual value of a city’s indicator and the data sources and processes used to measure it; while the indicator’s value is recorded in a machine-readable form (e.g., database or semantic web), the sources and measurement processes are buried in datasets and documents that are inaccessible or not machine readable. In the end, all we are left with is a record of indicator values without an understanding of what they actually measure and how they were measured. We have to rely on the good will of the people who reported to the data to adhere to the definitions.

Our goal is to formalize the definition of city indicators using Ontologies as implemented in the Semantic Web. By doing so we will:

- enable the representation of precise definitions of indicators thereby reducing the ambiguity of their interpretation and making them computationally accessible;
- represent the "data behind the data" or supporting data, enabling drilling down; and
- achieve interoperability, namely the ability to access, understand, merge and use measurement data available from datasets spread across the Semantic Web by providing a standard representation using OWL 2.

This paper describes a set of Ontologies that form a foundation for all indicators in the ISO standard. For each category of indicators, such as Education or Health, additional domain specific Ontologies are needed, and are defined in separate papers, e.g., Fox (2014).

In Section 2 we provide background on city indicators and related research. Section 3 defines the methodology we use to construct GCI Foundation Ontology. We then present in section 4 the core ontologies for representing place, measurement, populations, and statistics, and meta-information such as provenance, validity and trust. Section 5 introduces a core extension that focuses on measurements of populations that form the basis of almost every indicator. Section 6 provides an example of a student teacher ratio implemented using the ontology. Section 7 then evaluates the ontology[2].

2. City Indicators

In this section we review earlier work related to ontologies and City Indicators.

The rapid growth of Asian cities led the Asian Development Bank to launch a city indicator project in 1999. The objectives of the project were to "to establish a policy-oriented urban indicators database for research, policy formulation, monitoring of the development impact of interventions in the urban sector, comparison of performance between cities, and improving the efficiency of urban service delivery." (Westfall and de Villa, 2001 p. x). The result of the project provides the motivation and detailed definition of indicators. It also anticipates an important role for the World Wide Web in the representation and interconnection of indicator data.

The Organization for Economic Co-operation and Development (OECD: www.oecd.org) "provides a forum in which governments can work together to share experiences and seek solutions to common problems." At the core of their work is a large number of indicators spanning topics such as health, education, environment and trade. The indicators are documented in detail, in English, and the results are published as spreadsheets. Definitions of

[2] The Global City Indicator Foundation ontology can be found in: http://ontology.eil.utoronto.ca/GCI/Foundation/GCI-Foundation.owl.

the indicators using Semantic Web ontologies are not available. On the other hand, some OECD datasets have been the object of research in how to automatically transform statistical databases into linked data (Hausenblas et al., 2009; Capadisli et al., 2013).

In light of previous efforts to define city indicators, Hoornweg et al. (2007), identified the following aspects a good "indicator must possess to be accurate, timely and relevant for policy purposes:

- **Objective**: clear, well defined, precise and unambiguous, simple to understand.
- **Relevant**: directly related to the objectives.
- **Measurable and replicable:** easily quantifiable, systematically observable.
- **Auditable**: valid, subject to third-party verification, quality controlled data (legitimacy across users).
- **Statistically** representative at the city level.
- **Comparable/ Standardized** longitudinally (over time) and transversally (across cities).
- **Flexible**: can accommodate continuous improvements to what is measured and how. Have a formal mechanism for all cities and interested parties to comment on.
- **Potentially Predictive**: extrapolation over time and to other cities that share common environments.
- **Effective**: tool in decision making as well as in the planning for and management of the local system.
- **Economical**: easy to obtain/inexpensive to collect. Use of existing data.
- **Interrelated**: indicators should be constructed in an interconnected fashion (social, environmental and economics).
- **Consistent and sustainable over time:** frequently presented and independent of external capacity and funding support."

In response to the above, in 2010 the Global City Indicators Facility (GCIF, 2013)[3] was established to work with cities to identify a common set of indicators with standardized definitions and methodologies that can be consistently applied globally. The outcome of this effort is the international standard ISO 37120 "Sustainable development of communities — Indicators for city services and quality of life"[4].

In our review of the literature, there has been surprisingly little work done on creating ontologies to represent city indicators. There are only two related efforts:

As part of IBM's Smart Cities initiative, they have developed a comprehensive Ontology for representing various types of city knowledge, including city organization and services, flow of events and messages, and key performance indicators (Uceda-Sosa et al., 2011). OWL definitions of the classes and properties are provided. The ability to represent the definitions of indicators and precise semantics of the supporting data used to derive an indicator is missing. Axiomatization is limited and so its use of foundational ontologies.

Ghahremanloo et al. (2012) examined the indicators employed by the cities of Melbourne and Vancouver in order to identify a taxonomy of indicators, leaving the definitions of the indicators as strings. The focus of the work is primarily on methodology. The capability to

[3] http://www.cityindicators.org/
[4] First edition published 2014-05-15.

represent the definitions of indicators and precise semantics of the supporting data used to derive an indicator is not addressed.

In summary, the domain of city indicators represents a new area for the application of Ontologies. None of the existing efforts provide an Ontology for representing indicator definitions nor the supporting data and meta-data from which they are derived.

3. Methodology

The methodology we employ is based on Grüninger and Fox (1995). The first step is to identify scenarios for which the Ontology is to be used.

Cities want to perform better given their limited resources and funds. So imagine a world where cities openly publish their performance metrics and the supporting data from which they are derived. In this world, cities have software that can access this data over the Semantic Web and analyse it to determine why their performance differs from other cities. Based on this analysis they are able to modify their processes to perform better. The first step along this path is to be able to measure their performance, which they have for many years. But if they wish to learn from other cities they have to use the same metrics. That is why global cities are adopting ISO 37120. The second step is to diagnose the root cause of differences in performance. There are two types of diagnosis referred to in the literature: longitudinal analysis (i.e., how and why a city's indicators change over time) and transversal analysis (i.e., how and why cities differ from each other). If only an indicator's value is published, then all a city can determine is that their indicator value is the same or different. It cannot determine why it is different. In order to diagnose the root cause, a city would have to publish all of the supporting data that was used to derive an indicator's value. If we assume, in the spirit of Open Government and Open Data (Fox, 2013a), that cities openly publish their indicators and supporting data using Semantic Web standards, what ontologies should they use?

The second step of the methodology is to identify informal Competency Questions that the Ontology should enable the answering of. The competency questions act as a set of requirements that the Ontology must satisfy. In the next two sections, we begin each subsection with a set of competency questions that the ontology should satisfy.

The third step of the methodology is to formally identify the components of the Ontology. This is done using the OWL language. In this paper we depict some of the OWL axioms using diagrams. The complete set of axioms for the Ontology is provided in OWL files listed in the Appendix.

The fourth step is to formalize the competency questions using the ontology. We have done this using SPARQL as the query language. An example can be found in Fox (2014). It provides the complete Ontology for the ISO 37120 Education indicators, along with the formalization of competency questions as SPARQL queries.

4. Core Ontologies for Global City Indicators

In this section we select and extend ontologies that will form the foundation for representing ISO 37120 Global City Indicators (GCI) definitions and their published instances. By foundation we mean ontologies that generally apply to the representation of all indicator definitions and published instances. We illustrate the construction of the foundation ontology using the Student/Teacher Ratio (STR) indicator from ISO 37120. A number of issues arise in representing its definition. These issues will be addressed as we build the foundation ontology "one brick at a time". Each subsection begins by identifying a subset of competency questions that will motivate the needs for the ontology.

4.1. Placename Ontology

- What is the city being measured?
- What area does it cover?
- What places does it contain?

The STR is computed over a geographic area. In the case of GCIs, it would be a city. Hence, a requirement of the GCI ontology is the ability to identify the geographic area over which the indicator has been calculated. That is, to associate a "placename" with a geographic area. Such placenames could conceivably be applied to areas larger than a city, such as a region, state or country, or smaller than a city, such as a neighborhood or postal code. For example, a reference to Toronto should cover the city of Toronto but a reference to the Greater Toronto Area should cover the larger area encompassing neighbouring cities. But it must be clear which each refers to. A second requirement is that when two indicators are supposed to be computed over the same geographic area, they are in fact the same area. This means that an area has to have a unique identifier.

There are a number of ontologies that represent geographic and place information. Schema.org[5] provides classes of placenames such as `sc:City`, `sc:Country`, and `sc:State`. It also provides classes for `sc:GeoCoordinates` (i.e., elevation, latitude, and longitude) and `sc:GeoShape` denoted by a polygon or circle. The Linkedgeodata.org ontology[6] extends what can have a placename by providing classes for `gd:neighborhood`, `gd:building`, `gd:bridge`, `gd:hospital`, `gd:airport`, `gd:prison`, etc.

The GeoNames project provides over ten million placenames spanning the world. It provides an International Resource Identifier (IRI) for every placename so that they can be uniquely referred to. The GeoNames' placenames are instantiations of the Geonames Ontology[7] that integrates a number of ontologies, including Schema.org and Linkedgeodata.org, to provide a broad set of classes that span almost every conceivable type of place. For example, the unique IRI for the city of Toronto is: `http://www.geonames.org/6167865`. It is asserted to have a `geo:parentCountry` of `geo:6251999` which is the unique IRI for Canada.

---

[5] The Schema.org ontology is available at: http://schema.org/. We will use the prefix "sc:" to identify classes and properties from the ontology.
[6] The Linkedgeodata.org Ontology is available at: http://www.linkedgeodata.org/ontology/. We will use the prefix "gd:" to identify classes and properties from the ontology.
[7] The Geonames Ontology is available at: http://www.geonames.org/ontology/ontology_v3.1.rdf#. We will use the prefix "geo:" to identify classes and properties from the ontology.

4.2. Measurement Ontology

- What is the quantity being measured?
- What are its units? Meters? Yards? Kilometers?
- What is the measurement scale? Nominal, Ordinal, Interval, Ratio?
- How is the indicator derived?

A city indicator is a measure of some property of a city. At the core of an indicator lies a number. The question is what does that number represent? Measurement ontologies provide the basic concepts that underlie numbers. They divide measurement into a Quantity such as length (*the what*) and a Unit of Measure such as meters (*the how*). A Unit of Measure has a scale classified as nominal, ordinal, interval or ratio, and whether the number is the composition of dimensions such as velocity being composed of speed and direction, and whether it has a starting point such as absolute zero on the Kelvin scale.

In the case of the STR, the purpose of a measurement ontology is to provide the underlying semantics of the number. The importance of grounding an indicator in a measurement ontology is to assure that the indicator values are comparable, not that they are measuring the same thing (which is dealt with later), but the actual measures are of the same type, e.g., ratio of student and teacher population counts, or that the counts of the student and teacher populations are of the same scale (i.e., thousands vs millions).

Upper level ontologies such as SUMO (Niles and Pease, 2001) and CYC (Matuszek et al., 2006) provide classes for representing quantities, but the OM ontology[8] (Rijgersberg et al., 2011) provides a more rigorous ontology based on measurement theory. In addition to covering ratio scales, it covers nominal, ordinal and interval. QUDT[9] is an alternative to OM. We chose OM over QUDT for the reasons expressed in Rijersberg et al. (2013). The core classes of OM include:

- om:Quantity: What is being measured, such as a length or diameter.
- om:Unit_of_measure: It is the units and type of the quantity, such as a meter or yard.
- om:Measure: It combines the number with both the unit of measure and what is being measured (i.e., Quantity).

4.3. Statistics Ontology

- What defines the members of a population?
- Over what area is the population being drawn from?
- What is its unit of measure?
- How is the population size estimated?

The STR indicator is based on a measure of the number of students and teachers within a population designated by a city (Placename). One can view both as a statistical measurement in the sense that there is a population that we want to perform a measurement

[8] The OM ontology can be found at: http://www.wurvoc.org/vocabularies/om-1.8/. We will use the prefix “om:” to identify classes and properties from the ontology. Definitions and examples are taken directly from the ontology where quoted.

[9] http://www.qudt.org.

of, namely a city population, and we are counting the number of members that satisfy a description of a Student and a Teacher, respectively. While the STR requires a count of the population, other indicators may require statistical measures of mean, deviation, etc. of other characteristics of the population.

Anticipating the broader requirements of the Global City Ontology, we have adopted the GovStat[10] general statistics ontology (Pattuelli, 2009). In GovStat, the core class is the `gs:Population` to be measured. A `gs:Population` is linked to a parameter (e.g., mean, standard deviation) by the `gs:is_described_by` property, and the parameter is a sub class of `gs:Parameter`. In statistics it is almost always the case that only a portion of the population is measured. This portion is represented by the class `gs:Sample`, and the parameter being measured is represented as a subclass of `gs:Statistic`.

What is missing at this point is a definition of the population that we are measuring or from which a sample is to be taken. For the STR indicator the `gs:Population` must identify the area in which the population resides, i.e., the city, and what characterizes a member of the population, namely the characteristics of a Student or Teacher. For example, the characteristics of a Teacher could be:

- Fulltime, defined as teaching 30 or more hours per week, and
- Teaches at the primary or secondary level, where primary spans grades 1 thru 8 and secondary spans 9 thru 12.

We have extended the GovStat ontology as follows:

- Added a property to `gs:Population`, `gs:located_in`, that identifies the area that the Population is drawn from.
- Added a property to `gs:Population`, `gs:defined_by`, that identifies the class that all members of the Population are subsumed by.

A related Ontology is the Data Cube Vocabulary (Cyganiak et al., 2014). Its primary purpose is to represent the actual data points that comprise a Population or a Sample thereof (i.e., multi-dimensional data). It is also designed to support OLAP operations. It supports the association of attribute measures with datasets and their slices. It can be used to extend GovStat in representing the underlying datapoints.

4.4. Provenance Ontology

- Who created the actual value of the GCI?
- When was it created?
- What process was used to create it?
- What datasets is it based on?
- Has this GCI been revised?

An important aspect of an indicator is its provenance, namely where did it come from and how was it derived. Over the last decade, concerns around information validity, provenance and trust have grown. With the web now containing trillions of documents authored by millions of

[10] The GovStat Ontology is not available online, but a version with the GCI extensions can be found at: http://ontology.eil.utoronto.ca/govstat#. We will use the prefix “gs:” to identify classes and properties from the ontology.

people, the need to know whether the content is valid, where the content came from and whether to trust its creator, has taken on an increasing importance.

Much of the research into provenance has grown out of workflow management where the focus has been the evolution of a document as it proceeds through a sequence of edits, perhaps by different people and/or systems. Tracking the various versions created, who did what and when has been the primary concern. This research has culminated in the proposed Semantic Web ontology: PROV[11] (Belhajjame et al., 2012), which is based on the work of Hartig & Zhao (2010) and Moreau et al. (2010). We have chosen to use PROV as it satisfies the competency questions above and it is an emerging standard.

At the heart of the PROV ontology are three classes:

- **`pr:Entity:`** represents any artifact for which we want to specify its provenance. In our case it would be an indicator or the data from which the indicator was directly or indirectly derived.
- **`pr:Activity`**: the action (or sequence of actions) that creates or transforms an entity. In our case it may be a computation performed over some data set such as census data.
- **`pr:Agent`**: the person, organization, or system that performs or plays some role in the activity that transforms an entity. In our case it may be a software application that mines a data set or a person who reviews a data set.

Along with these classes are defined a set of properties that define the causal relationship among entities and activities, including:

- **`pr:wasGeneratedBy`**: It links an `pr:Entity` (domain) to a `pr:Activity` (range), identifying the activity that generated the entity.
- **`pr:used`**: It links an `pr:Activity` (domain) to an `pr:Entity` (range), identifying the entities used by an activity to produce a new entity.
- **`pr:wasDerivedFrom`**: Links two `pr:Entity`'s where domain entity was derived from the range entity (without indicating the method of derivation).
- **`pr:generatedAtTime`**: It links a `pr:Entity` (domain) to a `pr:time` (range), identifying the time the entity was generated.

4.5. Time Ontology

- When was this GCI constructed?
- What is the time of the measure before, during or after an event?
- Was the teacher population sizing done during the same time that the student population sizing was done?

Fundamental to the concept of provenance is the time at which measurements are taken, computed or derived. Questions may arise regarding the temporal relationship among indicators and among their supporting data. Not just at what time something occurred, but whether something occurred before, after or during some external event. For example, was "Total Employment" of New Orleans determined before or after Hurricane Katrina? Or did

[11] The PROV Ontology can be found at: http://www.w3.org/ns/prov#. We will use the prefix "pr:" to identify classes and properties from the ontology.

Katrina take place during the interval that the indicator was determined? To answer these questions, we need a much richer notion of time that supports reasoning about time points, time intervals and the relationships amongst them. Many time ontologies have been developed. We have chosen OWL-Time[12] for its simplicity and ability to represent time as a point or interval. OWL-Time is based on the work of Allen & Ferguson (1997) and described in Hobbs & Pan (2006).

4.6. Validity Ontology

- Over what period of time should the GCI be considered valid?
- Is the GCI believed to be an accurate measure by its creator?

An ongoing issue with the web is whether information/data found on a page is correct (true) or incorrect (false). Whether the creator of the information deliberately makes false statements, or unknowingly copies false information from another site, there is no way to discern what is correct from incorrect. The same holds with city indicators. Data and analyses that are believed to be true at the time they are gathered or computed, may be found over time to be incorrect. Or it may not be clear whether the information is true or not, especially if the indicator is based on a sampling of a population. But one can assign a degree of validity to the information. In addition, in the case where data is derived from other data, and the latter is no longer valid at some point of time, then the former becomes invalid for that same point of time. For example, if the STR is derived from a count of students and teachers, and if the student count is valid only within an interval of time such as the year in which it is gathered, then outside of that interval, both the student count and its dependent STR's validity are unknown.

Fox & Huang (2005) provide an ontology, called the Knowledge Provenance Ontology[13] (KP), for representing the validity (certainty) of a proposition, and axioms for propagating validity within a dependency network Huang & Fox (2004a; 2004b). It assigns to a "proposition" a validity between [0,1] or "unknown." This validity may be dynamic in that it changes over time. An example of the latter is a population count that is representative of a population only at a point of time or for an interval of time. The time interval during which the proposition's validity is known is called the "effective" time interval.

4.7. Dynamic Placenames

- Has the city's boundary changed during the time between two measures of an indicator?

Consider the unique placename for the City of Toronto. If we wish to do a longitudinal analysis of an indicator for Toronto, we run into a problem. The geographic definition of Toronto changed in 1998 after its amalgamation with five adjacent municipalities. Yet in the Geonames ontology there is a single Toronto; there is no representation for how placenames

[12] The OWL-Time Ontology can be found at: http://www.w3.org/2006/time. We will use the prefix "ot:" to identify classes and properties from the ontology.
[13] The Knowledge Provenance Ontology can be found at: http://ontology.eil.utoronto.ca/kp#. We will use the prefix "kp:" to identify classes and properties from the ontology.

evolve over time. Kauppinen and Hyvönen (2007) have addressed this problem. They propose an ontology based on Spatial Temporal Regions. A placename has associated with it a spatial region, defined by a polygon, and a time interval over which the placename and the region do not change. In the Global City Ontology we will refer to placenames whose spatial regions can change over time as Dynamic Placenames. Rather than adopt Kauppinen and Hyvönen's terminology directly, we adapt their ideas by reusing the provenance, time and validity ontologies to represent how place names change over time and the cause of their change.

4.8. Trust Ontology

- Do you trust the creator of the GCI?
- Do you trust the process used to create the GCI?
- How does your trust affect the validity of the GCI?

How do we represent the degree of trust we have in the creator of indicator values and the data from which they are derived. This representation of trust differs from degree of validity as trust refers not to the degree of certainty in the data but our trust in the agent/organization that produced the data. The obvious example is how to represent the trust we have in an organization that has a history of "cooking the numbers." The consequence of not having trust in the producer of data is that the validity one assigns to data or indicator will be reduced by this lack of trust.

Huang & Fox (2006) and Huang (2008) provide an ontology of trust[14]. The ontology views trust as occurring between two agents, where $agent_1$ has or has not trust in $agent_2$. Trust arises out of direct experience or the experience of others whom you may trust. Trust is also context dependent. For example, $agent_1$ may trust $agent_2$ in providing information on topics relevant to their expertise, such as a meteorologist characterizing the climate of a city, but lacks trust in $agent_2$ outside of their field of expertise. Finally, they identify two types of trust: 1) *trust in belief*, where $agent_1$ believes what $agent_2$ believes, and 2) *trust in performance*, where $agent_1$ believes that $agent_2$ will perform an activity properly. The Trust Ontology also addresses how the validity of an indicator or data changes by taking the original degree of validity, asserted by the creator ($agent_2$), and modifying it by the degree of trust the "user" ($agent_1$) has in the creator. This resultant validity is dependent on $agent_1$ and $agent_2$.

4.9. Composite Indicators

ISO 37120 does not take a position on how its indicators are combined to make quality of life choices. Never the less, the GCI foundation ontology is able to support compensatory (e.g., weighted average of indicators) and non-compensatory choice models by providing appropriate representations that decision models can use. For compensatory models, the OM Ontology provides classes and properties for multiplying an indicator value by a weight (i.e., om:Unit_multiplication) where the terms are defined by properties (i.e., om:term_1 and om:term_2) which can be used to link a weight to an indicator. For non-compensatory models of choice we need to represent indicators as an ordered sequence where benefits of indicators later in the sequence cannot over balance shortfalls of indicators earlier in the

[14] The Trust Ontology can be found at: http://ontology.eil.utoronto.ca/trust#. We will use the prefix "tr:" to identify classes and properties from the ontology.

sequence. The OM Ontology contains the class om:Ordinal_scale. Elements of the scale are defined by om:Ordered_measurement_scale_category, and their ordering in the scale are defined by the property om:rank.

5. Measuring Populations

Many of the indicators in ISO 37120 are a ratio of population measures. For example:

- Number of internet connections per 100,000 population (17.1)
- Annual number of public transport trips per capita (18.3)
- Annual number of tress planted per 100,000 population (19.2)
- Percentage of city population served by wastewater collection (20.1)

In order to represent these indicators, we need to combine and extend the measurement, statistics and placename ontologies. This section describes the extensions.

The STR is the ratio of the number of students to the number of teachers (within a city). Both students and teachers represent sets, i.e., the set of all students within a city (Placename) and the set of teachers within the same city (Placename). We need to represent the cardinality of these sets.

Figure 1 depicts the new Unit of Measure classes required to represent the number of students and teachers. We start by defining a unit of measure: `gci:Cardinality_unit`. Just as the meter is the unit of measure for length, a `gci:Cardinality_unit` is the unit of measure for the size of a set. The `gci:Cardinality_unit` is a ratio scale: `gci:Cardinality_scale`, which is a subclass of `om:Ratio_scale` and is has a zero element (namely zero). We specialize the `gci:Cardinality_unit` to the class `gci:Population_cardinality_unit` which is the unit of measure for the cardinality of a set defined by a Population and associate the symbol “pc” with it. For example, 1100pc represents a population cardinality (or size) of 1100. We can take full advantage of prefix notations available in OM to scale the numbers by defining units of measures: `gci:kilopc`, `gci:megapc` and `gci:gigapc` which are multiples of `gci:Population_cardinality_unit`. 1.1 kilopc represents 1100 pc.

We can now introduce the unit of measure for measuring a population ratio such as STR. `gci:Population_ratio_unit` is defined to be a subclass of `om:Unit_division`. It has two properties:

- `om:numerator` whose range is restricted to being a `gci:Population_cardinality_unit`.
- `om:denominator` whose range is restricted to being a `gci:Population_cardinality_unit`.

In other words, a population ratio is derived from two population cardinalities.

The above, provides the unit of measures for populations (pc) and population ratios (pc/pc) (the *how*). We now have to define *what* we are measuring which is referred to as a Quantity in the OM ontology. First, we need to define the `om:Quantity` for the size of the teacher and student populations from which the STR is derived. In Figure 2 we introduce `gci:Population_size` as a subclass of `om:Quantity`. Its `om:unit_of_measure` is the `gci:Population_cardinality_unit`.

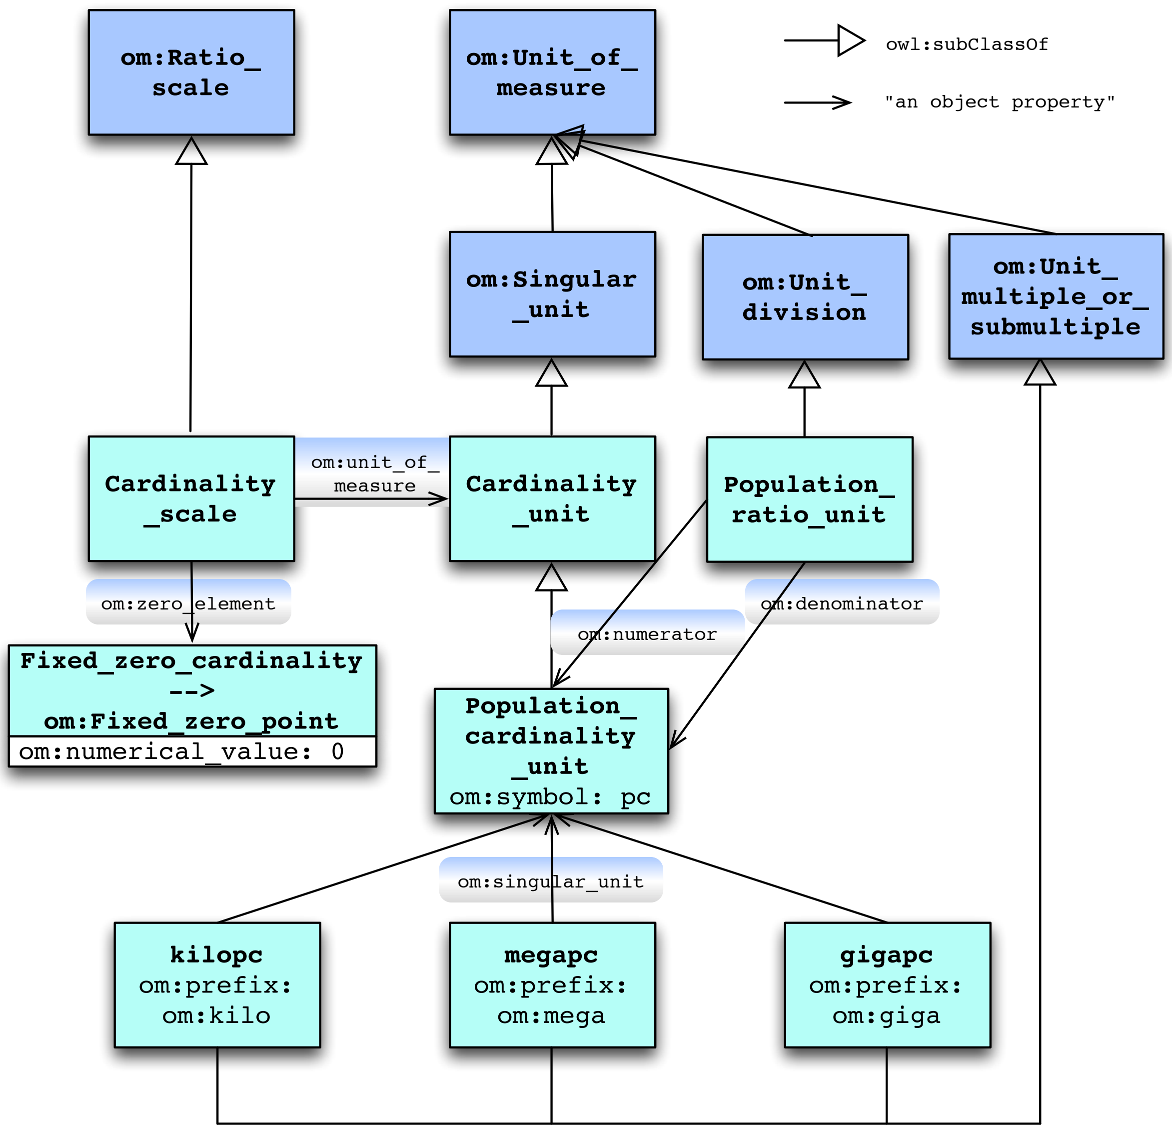

Figure 1: Population Cardinality and Ratio Units

What we now need to define is the population that we are measuring or from which a sample is to be taken. For the STR indicator the `gs:Population` must identify the area in which the population resides, i.e., the city, and what characterizes a member of the population, namely the characteristics of a Student or Teacher. For example, the characteristics of a Teacher could be:

- Fulltime, defined as teaching 30 or more hours per week, and
- Teaches at the primary or secondary level, where primary spans grades 1 thru 8 and secondary spans 9 thru 12.

As depicted in Figure 2, we have extended the GovStat ontology as follows:

- Added a property to `gs:Population`, `gs:located_in`, that identifies the area that the Population is drawn from.
- Added a property to `gs:Population`, `gs:defined_by`, that identifies the class that all members of the Population are subsumed by.

In order to complete definition of `gci:Population_size` pictured above, we need a further constraint. The property of (`gs:is_property_of`) the `gci:Population` must be a `gs:Count` parameter.

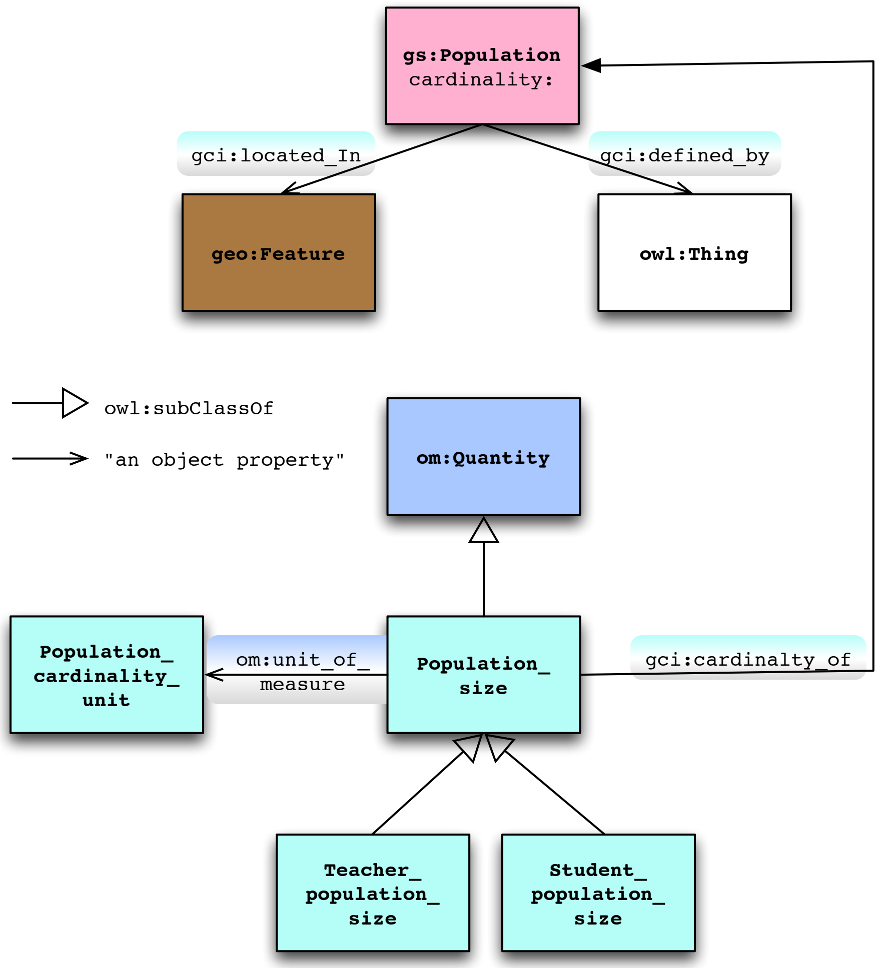

Figure 2: Population Class

We now have the requisite infrastructure to define population ratio GCIs (Figure 3). First we define the class of GCIs, `gci:Global_city_indicator`, as a subclass of `om:Quantity`. All GCIs will be a subclass of `gci:Global_city_indicator`. `gci:Education_GCI` is introduced as a subclass of `gci:Global_city_indicator` with a property that it is a `gci:for_city_service` `gci:Education_city_service`. Simply, this denotes that this indicator is for the education city service.

The actual value for a city's STR will be an instance of the measure `Student_teacher_ratio_measure`, which is linked to `Student_teacher_ratio_GCI` `and` is a subclass of `gci:Education_GCI`. It has the following properties:

- **`om:unit_of_measure`**, whose range is the `gci:Population_ratio_unit`. This signifies that the quantity is a ratio with a numerator and denominator that are restricted to being `gci:Population_cardinality_unit's`.
- **`gci:numerator & gci:denominator,`** whose ranges are `gci:Student_population_size` and `gci:Teacher_population_size` classes

respectively, which satisfy the `gci:Population_ratio_unit` numerator and denominator constraints.

- **`gci:for_city`**, whose range is a `geo:Feature` that uniquely identifies the city for which this is an indicator.
- **`gci:teacher_def & gci:student_def`**, whose range are a subclass of `Teacher` and `Student` respectively. These define the properties of the teachers and students that we are measuring. For example, all full time students in grades 1 through 12.

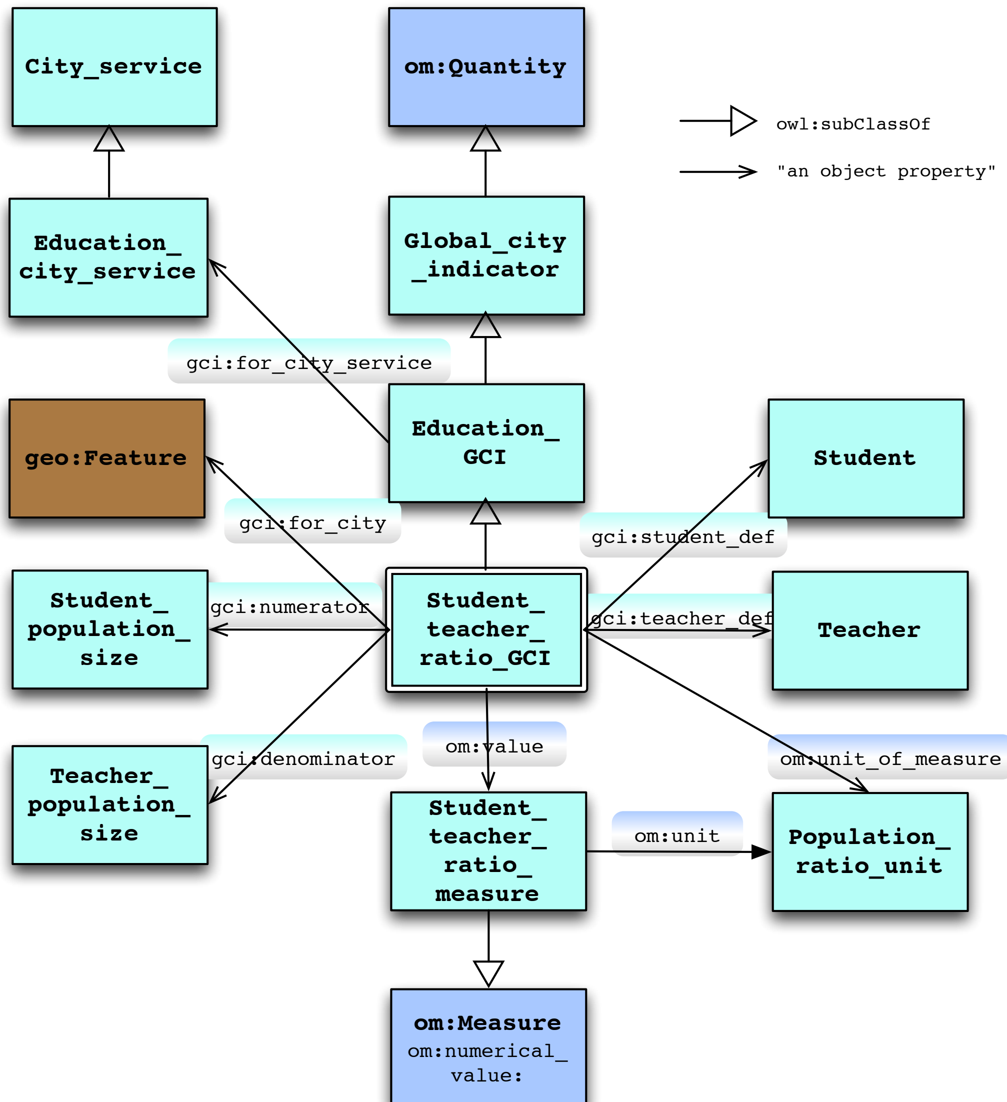

Figure 3: STR Definition

6. Example

In this section we show how a specific instance of the STR is represented in the ontology. The instances are represented as bifurcated rectangles where the top part identifies the name of the instance followed by the class it is an instance of (<instance name> -->> <class name>). The bottom part contains data properties and object properties for which, for brevity, we do not depict using links.

Let's assume that we want to create a STR for the city of Toronto, what we have to do is (Figure 4):

- Create an instance (`ex:TO_str`) of quantity `gci:Student_teacher_ratio_GCI`. This will be starting point of the value for the city of Toronto's student teacher ratio indicator. Set the `gci:for_city` object property to the URI of the Toronto placename.
- Create an instance of measure `gci:Student_teacher_ratio_measure` (`ex:TO_str_m`), fill the `om:numerical_value` property with the actual ratio (40 in this example). Link `ex:TO_str` to `ex:TO_str_m` using the `om:value` object property.

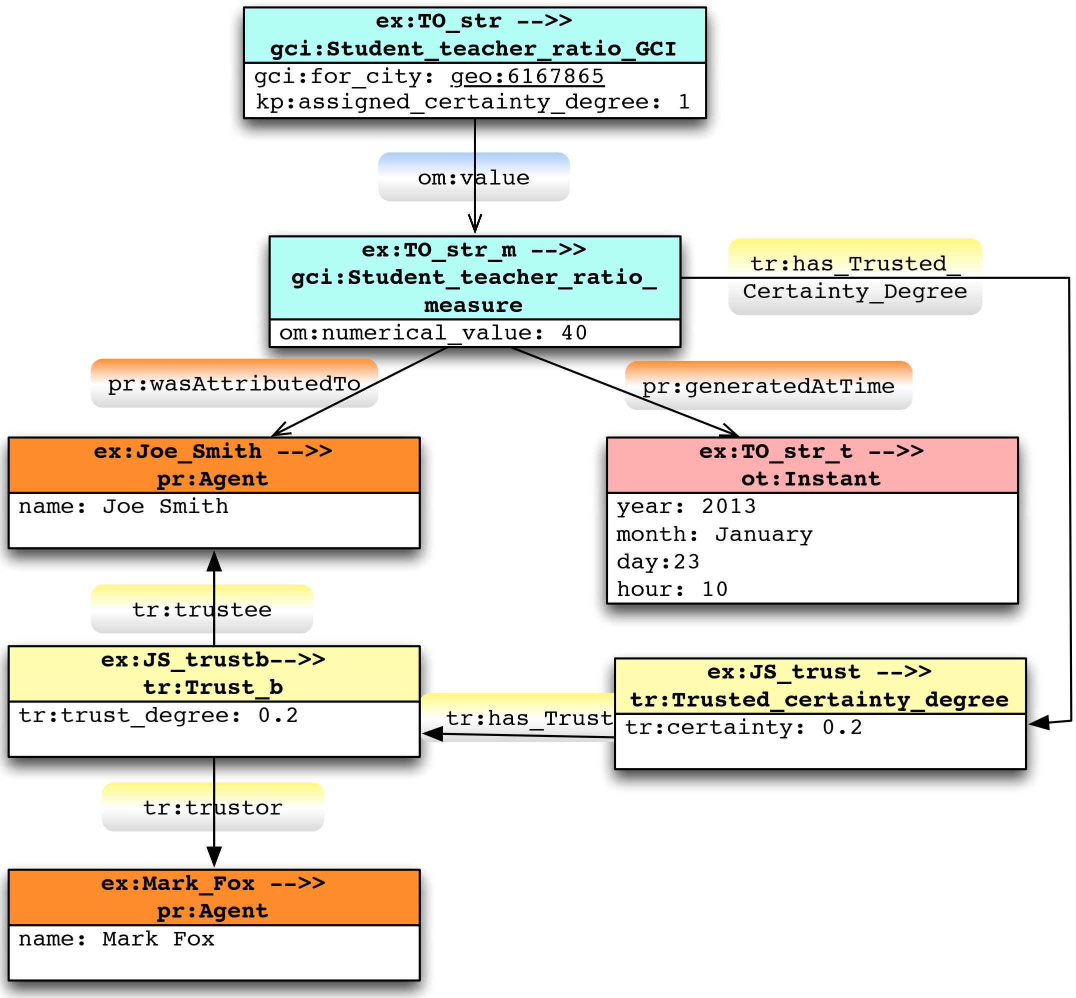

Figure 4: STR Measure plus Meta Information

With this representation, based on the classes these are instances of, we know that `ex:TO_str` represents a Student-Teacher ratio for the Education city service for the city of Toronto. The unit of measure is the ratio of two Population sizes, where the populations are defined by the Student and Teacher classes, respectively, which we have yet to define.

Next we will add two types of provenance:

- The date/time this ratio was created (23 January 2013 at 10am), and
- The Agent who created it (Joe Smith),

by instantiating their respective classes.

We then add the validity of the indicator. Since `gci:Student_teacher_ratio_GCI` is a subclass of `kp:KP_prop`, it inherits the `kp:assigned_certainty_degree` data property which we set to 1, i.e., the creator believes the value of the indicator is true.

Based on the above, we know that the source of the indicator was Joe Smith, which he created on January 23rd, 2013 at 10am.

Lastly, we add the degree we trust Joe Smith by doing the following:

- Adding the trustor, Mark Fox, and the degree to which he trusts Joe Smith by instantiating the trust in belief class (`tr:Trust_b`).
- Adding the trusted certainty degree to the Student Teacher ratio.

With this additional information, we know that although Joe Smith believes the STR to be correct, Mark Fox assigns a trusted certainty degree of 0.2 because he does not trust what Joe Smith believes.

The example to this point shows how a the single number that represents the student teacher ratio is represented using the OM ontology, along with the meta information representing provenance, validity and trust. If the supporting data is available, we can extend this example to represent what was used to derive the STR. Figure 5 depicts how the denominator of the STR is represented. `ex:TO_tpopsize` is an instance of `gci:Teacher_population_size` with the value specified by the `om:value` property linking it to a measure `ex:TO_tps_m`, which in turn specifies the teacher population size to be 1 `kilopcs`. The teacher population size is the cardinality of the population `ex:TO_tpop` linked to by `gci:cardinality`. This population is located in Toronto, and its membership is defined by the class of teachers who teach fulltime in primary school (this class definition has been abbreviated).

The numerator, i.e., the number of students, is defined similarly. Where they differ is in the definition of population membership, namely the definition of students and teachers. In order to precisely define a student and a teacher, we need an Education Ontology that covers the main points raised in the ISO 37120 definition, such as, primary vs. secondary school, administrative vs. teaching staff, school age students, cohorts, etc. An Education Ontology along with the complete definitions of the ISO37120 educational indicators can be found in Fox (2014).

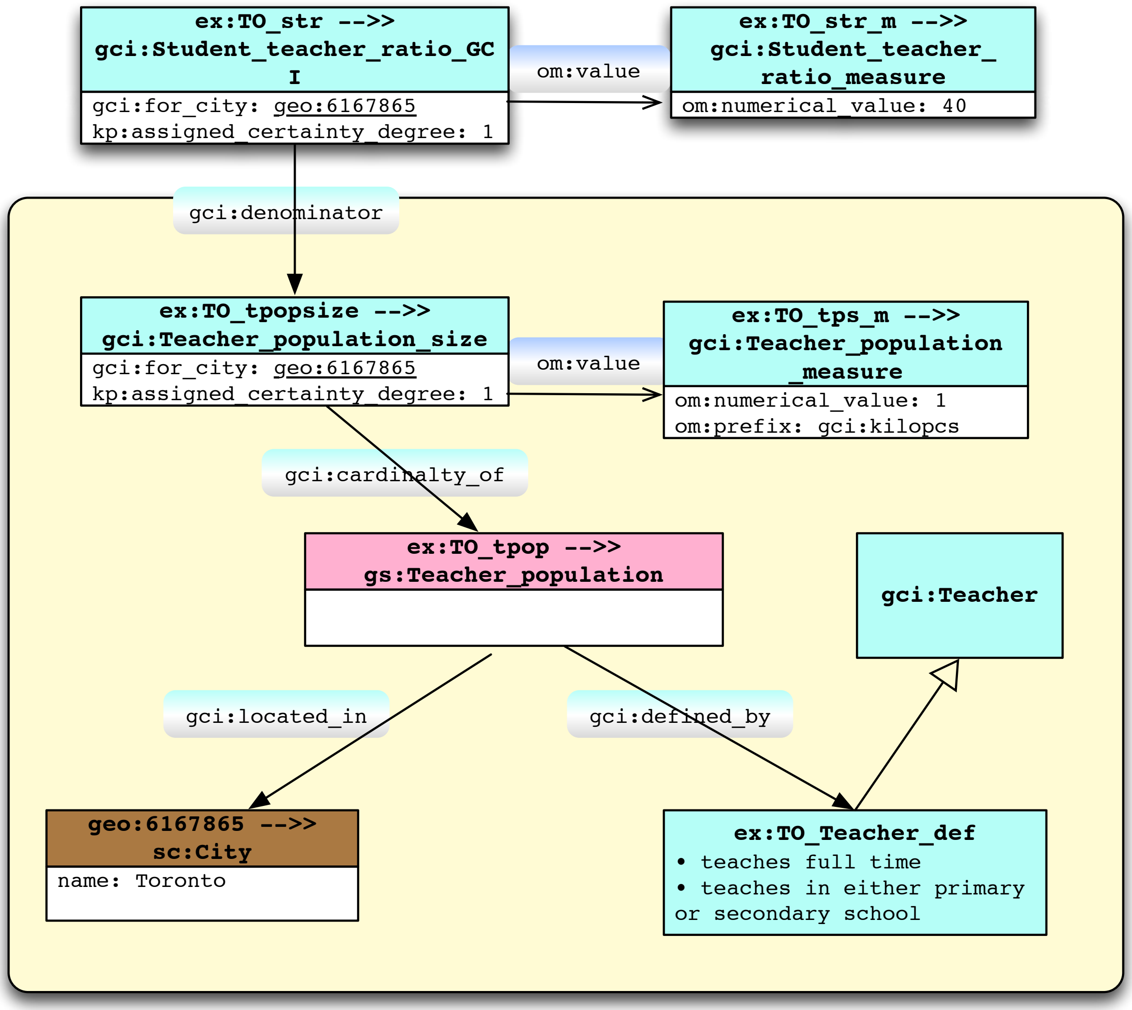

Figure 5: Decomposition of STR

7. Evaluation

We approach the evaluation of Global City Indicator Foundation ontology from five perspectives:

1. Is the ontology Competent? In Grüninger & Fox (1995), the requirements of an ontology are defined by a set of competency questions. These questions define how the ontology is to be used by applications. In order for an ontology to be competent with respect to a set of questions, it must be able to correctly deduce answers assuming the model has been instantiated correctly.
2. Is the ontology General? Are the classes general enough to represent other indicators and can it be easily extended where necessary?
3. Do we satisfy the aspects of a good indicator as identified in Hoornweg et al. (2007)?
4. Has the ontology been adopted by cities?

7.1. Competency

The competence of an ontology is defined by a set of questions the ontology must be able to answer. These questions fall into four categories:

1. Questions that require a simple retrieval of the value of a property. For example, the city of a particular indicator.

2. Questions that require the following of one or more links (properties) in the network.
3. Questions that focus on the consistency of the representation.
4. Questions that require some type of computation. For example, longitudinal or transversal analysis.

Regarding categories 1 and 2, Fox (2014) defines the competency questions for the STR indicator plus the remaining Education six indicators defined in ISO 37120. Each competency question is implemented in SPARQL. Regarding category 3, Fox (2013b) defines a set of consistency competency questions which have been implemented and tested. Regarding category 4, our future work will explore the types of analysis questions that need to be answered and any further extensions to the ontologies required.

7.2. Generality

A major goal of the development of the GCI Ontology is to make it as general as possible so that it can be reused across the all ISO 37120 indicators. Seven ontology modules were used in the STR example: Placename, Measurement, Statistics, Provenance, Time, Validity and Trust. Some of these modules were externally developed and used without change or extensions, such as Time and Provenance, some were extended, such as Measurement.

An analysis of the complete set of ISO 37120 indicators has shown that the foundation ontologies described in this paper apply across all categories of indicators (Fox, 2013b). But additional domain ontologies will have to be developed, specific to each indicator category, in order to represent the remaining indicators. 17 new modules have been identified:

- Census
- Economy
- Education
- Energy
- Environment
- Geography
- Fire
- Health
- Municipal Finance
- Municipal Governance
- Recreation
- Safety
- Shelter
- Waste
- Telecommunication
- Transportation
- Urban Ecology

7.3. Aspects of a Good Indicator

In Section 2, we discussed the aspects of a good city indicator defined by Hoornweg et al. (2007). In this section we revisit these aspects from the perspective of what and how the GCI Ontology achieves these aspects.

- **Objective**: *clear, well defined, precise and unambiguous, simple to understand.*

  The ontology provides a clear, precise representation of an indicator that is grounded in more foundational ontologies such as measurement theory, statistics, etc., This reduces, if not removes in most cases, ambiguity in the interpretation of the indicator.

- **Measurable and replicable**: *easily quantifiable, systematically observable.*

  This aspect is not addressed by the ontology.

- **Auditable**: *valid, subject to third-party verification, quality controlled data (legitimacy across users).*

  With the inclusion of provenance, validity and trust information in the ontology, the ability to audit the information is greatly enhanced. Add to it the more detailed information on the populations from which the data is drawn from, the quality of the data can be further verified.

- **Statistically representative** *at the city level.*

  While this aspect is not addressed by the ontology, the detailed representation of the place and populations sampled enables the audit function determine whether the information is statistically representative.

- **Comparable/ Standardized** *longitudinally (over time) and transversally (across cities).*

  The incorporation of dynamic placenames, measurement, time, statistics, and provenance makes it possible to perform longitudinal and transversal analysis and to verify that the data being compared is consistent with each other.

- **Flexible:** *can accommodate continuous improvements to what is measured and how. Have a formal mechanism for all cities and interested parties to comment on.*

  The ontology can be easily extended to include other measures as demonstrated by the generality of the underlying modules (i.e., placenames, provenance, measurement, etc.).

- **Interrelated**: *indicators should be constructed in an interconnected fashion (social, environmental and economics).*

  The Semantic Web's network representation is fundamentally an integrated representation, and enables the integration of indicators and the information they are based on.

- **Consistent and sustainable over time**: *frequently presented and independent of external capacity and funding support.*

  An important aspect of publishing indicators and their supporting data on the Semantic Web is the universal access it provides and its availability over time.

7.4. Adoption

The major challenge we face in getting cities to adopt and publish indicator data using any ontology is that most open data and information systems city staff are unaware of the Semantic Web and the role that ontologies play. Consequently, adoption requires a lengthy education process to get to the point that city staff understand its importance.

In the fall of 2014,the City of Toronto published its ISO 37120 indicators (Toronto, 2014). The publication is a PDF file whose content does not conform to semantic web standards. At the beginning of 2015, the Open Data Group of the City of Toronto agreed to publish the ISO

37120 data based on the GCI Foundation Ontology described in this paper. At the time of writing, the Education indicators have been completed. A pre-published version of these indicators is available at http://ontology.eil.utoronto.ca/ISO37120_6_2013_TO.owl. Once the remaining indicators have been completed they will be published on the city's open data portal.

The City Protocols project[15], led by the City of Barcelona, has adopted the GCI Foundation Ontology for its representation of city indicators, which include the ISO 37120 indicators.

The Research Data Alliance's[16] Quality of Urban Life Interest Group[17] has adopted the GCI Foundation Ontology for its representation of city indicators.

One of the takeaways from this project is that indicator *supporting* data is extremely sparse if not unavailable. This has led to the development of an Open Data Completeness Model (Fox & Pettit, 2015) that measures both the degree to which the data supporting an indicator is opening published, and the format and ontologies that are used to publish it. Another takeaways is the need for tools for instantiating indicators and supporting data. Tools like Protégé are not designed to instantiate complex patterns provided by the indicator definitions. An interactive pattern instantiator, as alluded to in Djedidi & Aufaure (2010), is needed. But ultimately the ontologies have to incorporated into a city's enterprise software; the acquisition, representation and publishing of indicator data has to be part of the software systems used to run a city.

8. Conclusions

Industrial Engineering and Management Science both share the view that you cannot manage what you do not measure. Enhancing the quality and efficiency of the operations and services of a city depends upon the ability to measure them. The development of city metrics faces many challenges. The first challenge is the selection and definition of the metrics. The second challenge is the adoption and use of these metrics by a large number of cities. These first two challenges have been the focus of the Global City Indicator Facility for the last five years and has resulted in the creation of ISO 37120 and the adoption of the standard by over 250 cities worldwide.

The third challenge is to represent the indicator definitions, city indicator values and their supporting data so that they can be published, linked, merged, and analyzed based on the principles of the Semantic Web. This work addresses this third challenge. In doing so it makes two unique contributions:

- It provides an ontology for the precise representation of the definitions of ISO 37120 indicators, and a city's specific indicator values and supporting data, so that they can be published on the semantic web, and
- It demonstrates the integration and extension of several core ontologies on a significant and important problem.

---

[15] http://cityprotocol.org
[16] http://rd-alliance.org/
[17] https://rd-alliance.org/groups/urban-quality-life-indicators.html

There are three directions that our current research is heading. The first direction is to complete the Global City Indicator Ontology to span the entire set of ISO37120 Indicators. For example, Fox (2014) defines an Education Ontology and combines it with the Foundation Ontology described herein to define ISO 37120 education indicators. The second dimension is to automate the consistency analysis of indicator data. Initial work has been completed and is summarized in Fox (2013b). The third direction is to automate the longitudinal and transversal analyses of city data. An overview of the goals and scope of the project can be found in Fox (2015).

9. Acknowledgements

We would like to thank Hajo Rijgersberg for his feedback on the use of the OM Ontology, and Patricia McCarney for her feedback on Global City Indicators. This research is supported in part by the Natural Sciences and Engineering Research Council of Canada.

10. References

Allen, J. F., and Ferguson, G., (1997), "Actions and events in interval temporal logic." In *Spatial and Temporal Reasoning*. O. Stock, ed., Kluwer, Dordrecht, Netherlands, pp. 205-245.

Belhajjame, K., Deus, H., Garijo, D., Klyne, G., Missier, P., Soiland-Reyes, S., and Zednik,S., (2012), "PROV Model Primer", http://www.w3.org/TR/prov-primer.

Capadisli, S., Auer, S., Ngonga Ngomo, A-C., (2013), "Linked SDMX Data: Path to high fidelity Statistical Linked Data for OECD, BFS, FAO and ECB", Submitted to *Semantic Web*, IOS Press.

Cyganiak, R., Reynolds, D., and Tennison, J., (2014), "The RDF Data Cube Vocabulary", http://www.w3.org/TR/vocab-data-cube/

Djedidi, R., and Aufaure, M-A., (2010), "ONTO-EVO^{A}L an Ontology Evolution Apporach Guided by Pattern Modeling and Quality Evaluation", *Proceedings of the Foundation of Information and Knowledge Systems: 6th Intl Symposium,* S. Link and H. Prade (Eds.), Springer-Verlag Berlin, pp. 286-305.

Fox, M.S., (2013a), "City Data: Big, Open and Linked", *Municipal Interfaces*, November, pp. 19-25. A longer version appeared as a Working Paper, Enterprise Integration Laboratory, University of Toronto. http://eil.utoronto.ca/smartcities/papers/City-Data-v5.pdf

Fox, M.S., (2013b), "A Foundation Ontology for Global City Indicators", Working Paper 3, Global Cities Institute, University of Toronto. http://eil.utoronto.ca/smartcities/papers/GCI-Foundation-Ontology.pdf

Fox, M.S., (2014), "An Education Ontology for Global City Indicators", Working Paper, Enterprise Integration Laboratory, University of Toronto, http://www.eil.utoronto.ca/smartcities/papers/fox-eil14.pdf

Fox, M.S., (2015), "The PolisGnosis Project: Representing and Analysing City Indicators", Working Paper, Enterprise Integration Laboratory, University of Toronto. http://eil.utoronto.ca/smartcities/papers/PolisGnosis-11mar2015.pdf

Fox, M.S., and Huang, J., (2005a), "Knowledge Provenance in Enterprise Information", *International Journal of Production Research,* Vol. 43, No. 20., pp. 4471-4492. http://www.eil.utoronto.ca/km/papers/fox-ijpr05.pdf

Fox, M.S., and Huang, J., (2005b), "An Ontology for Static Knowledge Provenance", In *Knowledge Sharing in the Integrated Enterprise,* IFIP, Vol. 183/2005, pp. 203-213.

Global City Indicators Facility (2013). "Global Cities Indicator Facility Briefing Document", University of Toronto. http://www.cityindicators.org/Deliverables/GCIF - Overview Report - Fall 2013_1-7-2014-1051

Ghahremanloo, L., Thom, J.A., Magee, L., (2012), "An Ontology Derived from Heterogeneous Sustainability Indicator Set Documents", *Proceedings of the Seventeenth Australasian Document Computing Symposium*, pp. 72-79.

Grüninger, M., and Fox, M. S., (1995), "Methodology for the Design and Evaluation of Ontologies." *Proceedings of the Workshop on Basic Ontological Issues in Knowledge Sharing*, IJCAI-95, Montreal, Canada.

Hartig, O., and Zhao, J., (2010), "Publishing and Consuming Provenance Metadata on the Web of Linked Data", *Proceedings of the Third International Provenance and Annotation Workshop.*

Hausenblas, M., Halb, W., Raimond, Y., Feigenbaum, L., & Ayers, D. (2009). "Scovo: Using statistics on the web of data." In *The Semantic Web: Research and Applications* (pp. 708-722). Springer Berlin Heidelberg.

Hobbs, J.R., and Pan, F., (2006), "Time Ontology in OWL", http:www.w3.org/TR/owl-time/.

Hoornweg, D., Nunez, F., Freire, M., Palugyai, N., Herrera, E.W., and Villaveces, M., (2007), "City Indicators: Now to Nanjing", World Bank Policy Research Working Paper 4114.

Huang, J., (2008). "Knowledge Provenance: An Approach to Modeling and Maintaining The Evolution and Validity of Knowledge", PhD Thesis, Dept. of Mechanical and Industrial Engineering, University of Toronto.

Huang, J., and Fox, M.S., (2004a). "Dynamic Knowledge Provenance", *Proceedings of Business Agents and Semantic Web Workshop*, pp. 372-387, National Research Council of Canada.

Huang, J., and Fox, M.S., (2004b), "Uncertainty in Knowledge Provenance", *Proceedings of the European Semantic Web Symposium*, Springer Lecture Notes in Computer Science.

Huang, J., and Fox, M.S, (2006), "An Ontology of Trust – Formal Semantics and Transitivity," *Proceedings of the International Conference on Electronic Commerce,* pp. 259-270. http://www.eil.utoronto.ca/km/papers/huang-ec06.pdf

Kauppinen, T., and Hyvönen, E., (2007), "Modeling and Reasoning About Changes in Ontology Time Series", In *Ontologies: A Handbook of Principles, Concepts, and Applications to Information Systems*, Integrated Series in Information Systems, Vol. 14, pp. 319-338, Springer.

Matuszek, C., et al., (2006), "An Introduction to the Syntax and Content of Cyc", In *Proceedings of the 2006 AAAI Spring Symposium on Formalizing and Compiling Background Knowledge and Its Applications to Knowledge Representation and Question Answering*.

McCarney, P. L., (2011) "Cities and climate change: The challenges for governance" Coordinating Lead Author with H. Blanco, J. Carmin and M. Colley, Ch. 9 in *Climate Change and Cities: First Assessment Report of the Urban Climate Change Research Network*, Cambridge University Press.

Moreau et al., (2010), "The Open Provenance Model Core Specification (v1.1)", *Future* Generation Computer Systems. Also see openprovenance.org.

Niles, I. & Pease, A., (2001), "Towards a Standard Upper Ontology", In *Proceedings of the 2nd International Conference on Formal Ontology in Information Systems FOIS-2001*, Ogunquit, Maine. New York: ACM Press.

Pattuelli, M.C., (2003), "The GovStat Ontology: Technical Report". The GovStat Project, Integration Design Laboratory, School of Information and Library Science, University of North Carolina at Chapel Hill, http://ils.unc.edu/govstat/papers/govstatontology.doc.

Rijgersberg, H., Wigham, M., and Top, J.L., (2011), "How Semantics can Improve Engineering Processes: A Case of Units of Measure and Quantities", *Advanced Engineering Informatics*, Vol. 25, pp. 276-287.
Toronto, (2014), "Toronto's 2013 Results Under ISO 37120 Indicators of City Service Delivery and Quality Of Life". http://www1.toronto.ca/City%20Of%20Toronto/City%20Managers%20Office/Toronto%20Progress%20Portal/ISO%2037120/Final%20-%20Summary%20of%20Toronto%27s%20WCCD-%20ISO%2037120%20Results-6%20.pdf
Uceda-Sosa, R., Srivastava, B., and Schloss, B., (2012), "Building a Highly Consumable Semantic Model for Smarter Cities", *In Proceedings of the workshop on AI for an Intelligent Planet,* ACM.
Westfall, M.S., and d Villa, V.A., (2001), *Urban Indicators for Managing Cities*, Asian Development Bank.
World Bank, (2008), "Global City Indicators Program Report: Preliminary Final Report", April 2008.

11. Appendix: Ontologies

The Global City Indicator Foundation ontology can be found in: http://ontology.eil.utoronto.ca/GCI/Foundation/GCI-Founation.owl.

The Global City Indicator Education ontology can be found in: http://ontology.eil.utoronto.ca/GCI/Education/GCI-Education.owl.

The Global City Indicator Innovation ontology can be found in (to appear): http://ontology.eil.utoronto.ca/GCI/Innovation/GCI-Innovation.owl.

URIs for all of the ISO37120 indicators can be found in: http://ontology.eil.utoronto.ca/ISO37120.owl.

Definitions of the ISO37120 education indicators, using the GCI Foundation and Education ontologies can be found in: http://ontology.eil.utoronto.ca/GCI/ISO37120/Education.owl.

Definitions of the ISO37120 innovation indicators, using the GCI Foundation and Innovation ontologies can be found in (to appear): http://ontology.eil.utoronto.ca/GCI/ISO37120/Innovation.owl.