

Representing Dataset Quality Metadata using Multi-Dimensional Views

Jeremy Debattista
University of Bonn /
Fraunhofer IAIS, Germany
jeremy.debattista@iais-
extern.fraunhofer.de

Christoph Lange
University of Bonn /
Fraunhofer IAIS, Germany
math.semantic.web
@gmail.com

Sören Auer
University of Bonn /
Fraunhofer IAIS, Germany
auer@cs.uni-bonn.de

ABSTRACT

Data quality is commonly defined as *fitness for use*. The problem of identifying quality of data is faced by many data consumers. Data publishers often do not have the means to identify quality problems in their data. To make the task for both stakeholders easier, we have developed the Dataset Quality Ontology (daQ). daQ is a core vocabulary for representing the results of quality benchmarking of a linked dataset. It represents quality metadata as multi-dimensional and statistical observations using the Data Cube vocabulary. Quality metadata are organised as a self-contained graph, which can, e.g., be embedded into linked open datasets. We discuss the design considerations, give examples for extending daQ by custom quality metrics, and present use cases such as analysing data versions, browsing datasets by quality, and link identification. We finally discuss how data cube visualisation tools enable data publishers and consumers to analyse better the quality of their data.

Categories and Subject Descriptors

H.2.7 [Database Administration]: Data dictionary/directory

General Terms

Measurement, Documentation, Standardization

1. INTRODUCTION

There are various definitions for the term ‘Data Quality’. Robert Pirsig defines quality as *the result of care* [22], whilst Juran defines quality as *fitness for use* [19]. Juran’s views on data quality were shared by Phillip Crosby, who defined quality as *conformance to requirements* [9]. A substantial amount of Linked (Open) Datasets have already been published¹. Most of these facts are extracted from heterogeneous sources, including semi-structured data and unstructured data, which causes great variance in quality. Therefore such sources could lead to various problems such as inconsistencies and incompleteness, which could render a dataset to not be

¹Some statistics can be found in <http://lod-cloud.net> and <http://stats.lod2.eu>

fit for a certain tasks. Moreover datasets might also evolve during their life-span, leading to increase or decrease in quality.

Identifying the right quality factors for a dataset is a challenge which is always faced by many data consumers. The main problem stems from the fact that different domains require different quality metrics. Various research work [6, 13, 17] defines a number of quality factors pertinent to linked open datasets. Zaveri et al. [24] provide a systematic review of this and further literature, categorising the different metrics.

To enable the definition and use of different metrics in a standardised manner, we introduce the *Dataset Quality Ontology (daQ)*. daQ is a light-weight core vocabulary for representing the results of quality benchmarking of a linked dataset, again as linked data. In linked *open* data settings this allows for embedding quality metadata into datasets, thus ‘stamping’ them with a number of quality measures, allowing for the expression of concrete, tangible values that represent the quality of the data.

This paper is an extended version of the original presentation of an early version of daQ [11]. The main *new* contribution lies in re-using the W3C Data Cube Vocabulary [10], which allows for representing data quality metrics as statistical observations in multi-dimensional spaces. This approach enables the representation and analysis of quality assessments in a fine-grained manner. In particular, we will show how the data cube representation enables visual quality analysis using standard tools.

To put the reader into the context of this work, we introduce a use case:

Quimp is a startup company providing various life science linked datasets. Their business model is that their customers (the real publishers) provide their data in various formats to the Quimp Portal. Meanwhile Quimp semantically lift the data to a standardised RDF linked data representation using a number of pre-defined vocabularies. This linked data is periodically updated, having new resources added and others becoming obsolete. Quimp will then offer access for these linked datasets to data consumers. Data consumers often complain about the fact that datasets suffer a lot from incorrect and inconsistent facts. Concerned with these complaints, Quimp decided to start computing quality metrics on the semantically-enriched data. This quality metadata is not only computed for one-time analyses, but it is stored, in order to track the long-term evolution of dataset quality across multiple versions. Visualising the quality metrics and their development over time helps Quimp identify what aspects of their data are not up to standard, and therefore ensuring that quality over the differ-

ent versions does not diminish.

The remainder of this paper is structured as follows: in Section 2 we discuss use cases for the daQ ontology, including its new multi-dimensional aspect. Then, in Section 3 and 4 we discuss the vocabulary design and show how data can be explored and used. Finally, in Section 6 we give an overview of similar ontology approaches before giving our final remarks in Section 7.

2. USE CASES

Linked Open Data quality has different stakeholders in a myriad of domains, however, the stakeholders can be cast under either *publishers* or *consumers*. *Publishers* are mainly interested in publishing data that others can reuse. *Consumers*, both human end users and machine agents assisting them, require to use this published data in their applications.

Data consumers, both human and machine, may find it challenging to assess the quality of a dataset, i.e. its fitness for use. Currently, there is no standard way of how data publishers can assess the quality of their linked datasets. Most of these publishers rely on their mutual trust they have with data providers, believing that the data they provide is good for data consumers, leaving the data publishers in the dark of the value and quality of their data. Undoubtedly, quality metadata is also significant to the data consumer who ultimately has to decide what data is fit to their use case. Activities of the recently formed *W3C Data on the Web Best Practices Working Group* (DWBP)²³ include developing a standard vocabulary for assessing and representing metadata about quality. Our daQ, or an adaptation or extension of it, is a candidate.

The following use cases (UC) show how both data publishers and consumers can benefit from having quality metadata about datasets. The UC thus motivate the need for developing a standard like daQ. Our original paper discusses further use cases [11].

2.1 UC1: Analysis of Data Versions

Ideally, data publishers update their published datasets regularly to (a) keep the data fresh and up-to-date; (b) clean data to improve quality; (c) keep up with the data curation lifecycle. However, it is sometimes difficult to identify which aspects of the data are lacking quality standards. Furthermore, it is even more difficult to analyse how data quality changed over time. Assuming that there are tools that automatically analyse data quality and output daQ metadata (such as our implementation mentioned in Section 5), their data cube structure enables a multi-dimensional representation, where the versions of a dataset form one dimension, and the different quality metrics form the other dimension.

2.2 UC2: “Fit” Dataset for Retrieval

Alexander et al. [1] provide the readers with a motivational use case with regard to how the VoID ontology (cf. Section 6) can help with effective data selection. The authors describe how a consumer can find the appropriate dataset by:

- criteria related to content (what is the dataset mainly about);
- its links from and to other datasets;
- the vocabularies used in the dataset.

The daQ ontology gives an extra edge to “appropriateness” by providing the consumer with quality criteria on the candidate datasets.

An objective assessment of data quality enables data consumers to determine if a dataset is fit for a certain use case. Currently, tools

²<https://www.w3.org/2013/dwbp>

³The authors are affiliated with this WG, contributing to the standardisation of quality assessment of LOD.

targeting human data consumers, such as semantic web search engines [18] or Data Web browsers [5, 15, 16], do not focus on dataset quality when presenting search results. With the introduction of the daQ framework, tools that provide faceted browsing facilities, such as the CKAN data portal engine⁴, are enabled to provide more information about a dataset’s quality attributes. Such functionality is attributable to the flexibility of the ontology, enabling various filtering and ranking possibilities of the dataset quality metrics. This would permit human data consumers to better understand the quality attributes of a dataset, and thus choose which is the most fitting to their use case. The recent addition of multi-dimensionality to quality metadata in the daQ model enables data consumers to track and follow quality improvements of data publishers on their datasets over time. This also opens a sundry of opportunities leading to the assessment of data publishers regarding their willingness to enhance the value of the data in terms of quality. To keep the quality metadata of open datasets easily accessible, we recommend that each dataset contains the relevant daQ metadata graph within the dataset itself.

2.3 UC3: Link Identification

Identifying links between existing datasets is one of the main drivers that makes the Linked Open Data cloud more coherent. Tools such as LIMES [21] or Silk [23] support the automatic identification of links according to built-in as well as user-defined criteria. The introduction of quality metadata to datasets will add another criterion for link identification, in that linking algorithms can also take the quality of the target dataset into consideration before linking to it. Linking tools could also consider the needs of a data consumer who might not only require to link to any high quality entity, but possibly even to those datasets which the consumer deems “fit” to her cause. This can be done by ranking and filtering candidate datasets according to criteria such as weights on specific quality metrics defined by the consumer (as described in UC2). Linking resources of proven quality helps to improve the quality of both datasets participating in the link. The generic framework proposed for the daQ ontology ensures that any custom metric defined by third parties can be easily integrated into any tool supporting such quality metadata for linking.

2.4 UC4: Extension of the Five Star Scheme

The popular five star scheme for deploying open data⁵, which we propose to extend by a sixth star for quality, defines a set of widely accepted criteria that serve as a baseline for assessing data reusability. The reusability criteria defined by the five star scheme and by the quality metrics are largely measurable in an objective way. Thanks to such objective criteria, one can assess the reusability of any given dataset without the major effort of, for example, running a custom survey to determine whether its intended target audience finds it reusable. Such a survey may, of course, still help to get an *even better* understanding of quality issues. As a consumer, the benefits of a sixth star is that good quality datasets can be discovered. On the other hand, as a data publisher, the benefits of having the sixth star are that (i) the published data conforms to the established domain quality metrics; and (ii) catalogued and archived datasets (refer to [11]) can be easily discovered when consumers filter by quality aspects.

3. THE DATASET QUALITY ONTOLOGY (DAQ)

⁴<http://ckan.org>

⁵<http://5stardata.info>

The idea behind the Dataset Quality Ontology⁶ (daQ) is to provide a generic core vocabulary, allowing a uniform definition of specific data quality metrics. This metric definition would then allow publishers to represent the metadata resulting from benchmarking the quality of their datasets, or even to attach these metadata to their linked datasets.

3.1 The basic daQ Concepts

We first recapitulate the fundamental concepts, which were originally introduced in more detail in [11]. Quality metadata is intended to be represented as a *Quality Graph*. The latter concept is a subclass of `rdflg:Graph` [7]. This means that the quality metadata is stored and managed in a named graph that is separate from the dataset but can be embedded into the dataset if desired. Named graphs also allow the digital signing of graphs [8], thus ensuring trust in the computed metrics.

The daQ ontology distinguishes between three layers of abstraction, based on the survey work by Zaveri et al. [24]. As shown in Figure 1 Box B, a quality graph comprises a number of different *Categories*, which in turn possess a number of quality *Dimensions*⁷. A quality dimension groups one or more quality *Metrics*.

3.2 Extending daQ for Multi-Dimension Representation and Statistical Evaluation

The Data Cube Vocabulary [10], abbreviated using the prefix `qb:`, allows the representation of statistical observations in multidimensional spaces. In the initial daQ design [11], computing quality metrics of different resources (usually: datasets), or even different revisions of the same resource, resulted in multiple *quality graphs*, consisting of multiple instances of *Metric* classes representing the individual observations. Multidimensional analysis of these observations, e.g. across the revision history of a dataset, would thus have required complex querying. Reusing the standardised Data Cube Vocabulary in daQ allows us to represent quality metadata of a dataset as a collection of *Observation*. This also permits applying the wide range of tools that support data cubes to quality metadata, including the CubeViz visualisation tool⁸.

Figure 1 shows the current state of daQ, where the introduction of data cubes entails some structural changes over the initial version of the ontology from [11]. A *Quality Graph* is a special case of `qb:DataSet`, which allows us to represent a collection of quality observations complying to a defined dimensional structure. daQ defines the structure of such observations by the `qb:DataStructureDefinition` shown in Listing 1.

The `daq:QualityGraph` definition in Listing 2 also specifies one restriction that controls the property `qb:structure` and its value to the mentioned definition, thus ensuring that all *Quality Graph* instances make use of the standard definition. Having a standard definition ensures that all *Quality Graphs* conform to a common data structure definition, thus datasets with attached quality metadata can be compared.

```

daq:QualityGraph
  a rdfs:Class, owl:Class ;
  rdfs:subClassOf rdflg:Graph , qb:DataSet ,
  [ a owl:Restriction ;
    owl:onProperty qb:structure ;
    owl:hasValue daq:dsd ] ;

```

⁶<http://purl.org/eis/vocab/daq>

⁷In this paper we will refer to these as *quality dimensions*, to avoid confusion with data cube dimensions

⁸<http://cubeviz.aksw.org>

```

daq:dsd a qb:DataStructureDefinition ;
  # Dimensions: metrics and what they were
  # computed on
  qb:component [
    qb:dimension daq:metric ;
    qb:order 1 ; ] ;
  qb:component [
    qb:dimension daq:computedOn ;
    qb:order 2 ; ] ;
  # Measures (here: metric values)
  qb:component [ qb:measure daq:value ; ] ;
  # Attribute (here: unit of measurement)
  qb:component [
    qb:attribute sdmx-attribute:unitMeasure
    ;
    qb:componentRequired false ;
    qb:componentAttachment qb:DataSet ; ] .

```

Listing 1: The Data Structure Definition (Turtle Syntax)

```

rdfs:comment "Defines a quality graph
  which will contain all metadata about
  quality metrics on the dataset." ;
rdfs:label "Quality Graph Statistics" .

```

Listing 2: The Quality Graph Definition (Turtle Syntax)

The detailed definitions of all abstract classes in Figure 1 Box B, have been introduced in [11]. The `daq:Metric` class now no longer has an immediate *value* property, but is linked to from an `qb:Observation` by `daq:metric` (defined as a `qb:DimensionProperty`), whose inverse `daq:hasObservation` we define for convenience. The properties `daq:computedOn` and `daq:value` are now defined as `qb:DimensionProperty` and `qb:MeasureProperty` respectively. The former is defined in each observation instance rather than once as a *Quality Graph* property. Each observation also has a `dc:date` property, which holds the timestamp of when it was computed. Each custom metric definition should also include the `daq:expectedDataType` property. This will indicate the observation's value datatype. The optional property `sdmx-attribute:unitMeasure` can be defined on an observation instance, enabling a system (application) to further specify the semantics of the unit of measurement of the value. For example, some of a dataset's accessibility metrics are reasonably measured as durations in seconds, whereas the majority of metrics can be expressed as plain numeric figures in the normalised interval [0, 1].

4. USING THE ONTOLOGY

4.1 Extending daQ

The classes of the core daQ ontology are intended to be extended by specific and custom quality metrics that characterise a dataset's "fitness for use" [19] in a particular domain. We have defined the quality dimensions and metrics described in [24], some of which are being considered to be standard metrics to calculate quality on Linked Open Data sets whilst others are specific to the DIACHRON project, in whose context we are doing this research (refer to Section 5). **Extending** the daQ ontology means adding new quality protocols that inherit the abstract concepts (Category–Dimension–Metric). Custom quality metrics do not need to

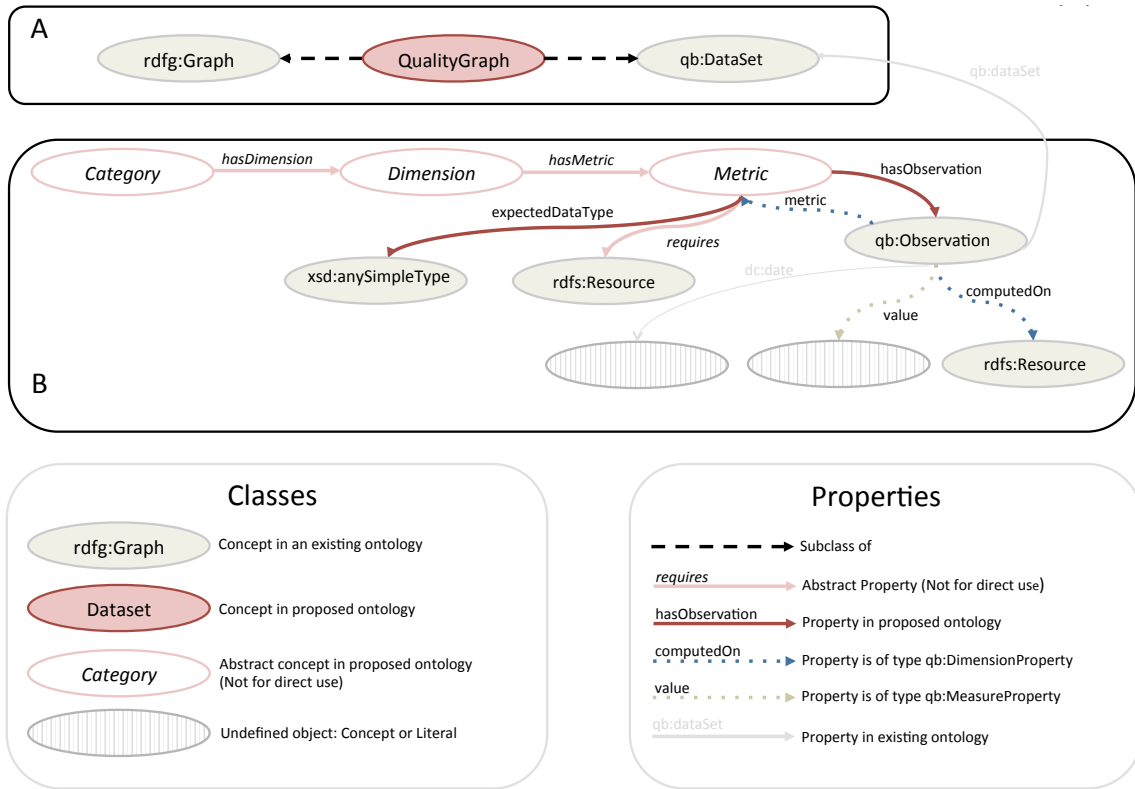


Figure 1: The extended Dataset Quality Ontology (daQ)

be included in the daQ namespace itself; in fact, in accordance with LOD best practices, we recommend extenders to make them in their own namespaces. Figure 2 shows an illustrative example of extending the daQ ontology (TBox) with a more specific quality attribute – the RDF Availability Metric as defined in [24] – and an illustrative instance (ABox) of how it would be represented in a dataset.

The Accessibility concept is defined as an `rdfs:subClassOf` the abstract `daq:Category`. This category has five quality dimensions, one of which is the *Availability* dimension. This is defined as an `rdfs:subClassOf` `daq:Dimension`. Similarly, *RDFAvailabilityMetric* is defined as an `rdfs:subClassOf` `daq:Metric`. The specific properties *hasAvailabilityDimension* and *hasRDFAccessibilityMetric* (sub-properties of `daq:hasDimension` and `daq:hasMetric` respectively) are also defined (Figure 2)

4.2 Publishing daQ Metadata Records

We encourage publishers of linked *open* data to offer daQ metadata as an RDF named graph in their published dataset. Since such a daQ metadata record requires a lot of metrics to be computed, it is not normally intended to be authored manually. Publishing platforms such as CKAN should offer such an on-demand computation to dataset publishers (as sketched in [11]; see section 5 for our own implementation progress). Listing 5 shows an instance of the `daq:QualityGraph` in a dataset. *ex:qualityGraph1* is a named `daq:QualityGraph`. The defined graph is automatically a `qb:DataSet`, and due to the restriction placed on the `daq:QualityGraph` (see Listing 2), the value for the `qb:structure` property is defined as `daq:dsd` (see Listing 1). In the named graph, instances for the `daq:Accessibility`, `daq:Availability`,

`daq:EndPointAvailabilityMetric` and `daq:RDFAvailabilityMetric` are shown. A metric instance may have been used to make a number of observations. Each of these observations specifies the metric value (`daq:value`), the resource the metric was computed on (`daq:computedOn`), when it was computed (`dc:date`), the metric instance (`daq:metric`) and finally to what data cube the observation is defined in (`qb:dataset`).

4.3 Exploring and Visualising the daQ Metadata

CubeViz is a tool for visualising data cubes. Figure 3 depicts four different CubeViz chart visualisations of the same quality metadata⁹.

A *horizontal bar* represents each metric (Figure 3(a)) and shows its value (x-axis) with respect to the dataset (y-axis). Here, the different ‘datasets’ analysed are actually successive revisions of one dataset. This chart provides a clear view of how the value associated to each one of the measured metrics changes as the dataset (in this case) evolves. The horizontal layout is appropriate when the range of metric values is wide, and the number of different datasets is relatively small.

Similar to the horizontal bars chart, the *vertical bar chart* (Figure 3(b)) allows the user to compare the values computed for each of the metrics (y-axis), with respect to the dataset (x-axis). In contrast with its horizontal counterpart, this chart is more appropriate when there are many datasets analysed but the range of metric values is not so wide.

⁹The quality metadata used can be found in https://raw.githubusercontent.com/diachron/quality/master/src/test/resources/cube_qg.trig

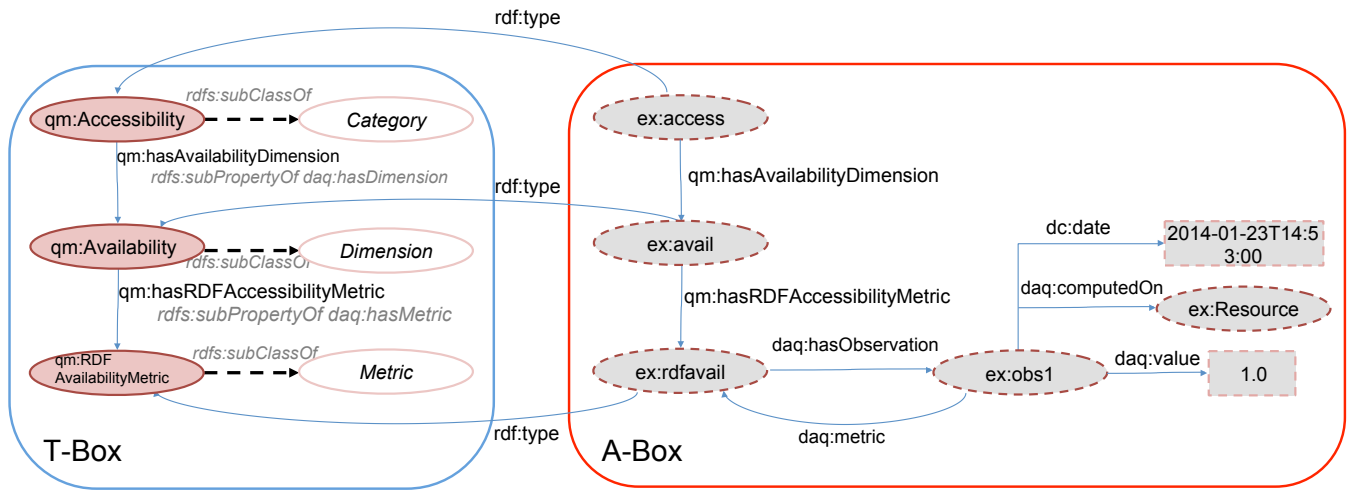


Figure 2: Extending the daQ Ontology – TBox and ABox

In the *radar chart* (Figure 3(c)), the datasets are represented as slices of a circle and the values corresponding to the metrics are depicted as points and lines of a particular colour. This chart provides a clear view of how the values of the metric differ from each other for each particular dataset. Furthermore, it allows one to assess the overall quality of a dataset, by showing whether the values of the metrics are concentrated around sections of the circle regarded as ‘good’ or ‘bad’.

The lines plot (Figure 3(d)), lists the different datasets against the values of the metrics. Here, where ‘different datasets’ are actually different revisions in the evolution of one dataset, this plot provides a comparison of the evolution of the quality of the dataset, with respect to each metric. The lines emphasise the points where the values of the metrics changed noticeably from one version to the next.

4.4 Analysing Observations based on Quality Dimensions

The daQ framework allows the definition of quality metrics in three levels of abstraction: Category–Dimension–Metric. Although the instances have a link between these three levels, we only perform observations on the metric level. Therefore, when visualising and analysing observations, the consumer would only be able to observe the metrics from all quality categories and dimensions, instead by specific quality dimension or category. Thanks to the link between the three levels of abstraction, no manual human intervention is required to analyse a set of metrics *grouped* by a specific quality dimension.

Data Cube slices allow the grouping of observation subsets. Since slices are not intended to represent arbitrary selections in a data cube, but only selections that result from fixing the values for some dimensions, qb:ObservationGroup has to be used instead. Listing 3 shows a SPARQL CONSTRUCT defining an observation group, where all observations in the Accessibility dimension are grouped in a constructed ex:dimObs1 resource.

```

CONSTRUCT {
  ex:dimObs1 a qb:ObservationGroup ;
  qb:observation ?obs .
}
WHERE {

```

```

SELECT DISTINCT ?metricInst ?obs {
  ?dimInst a dqm:Accessibility .
  ?dimInst ?prop ?metricInst .
  ?metricInst daq:hasObservation ?obs .
  ?metricInst a ?metric .

GRAPH <http://www.diachron-fp7.eu/dqm#> {
  ?prop rdfs:subPropertyOf daq:hasMetric .
  ?metric rdfs:subClassOf daq:Metric .
}

```

Listing 3: Creating A Data Cube Observation Group using SPARQL

The resulting construct output is shown in Listing 4.

```

ex:dimObs1 a qb:ObservationGroup ;
qb:observations ex:obs1 , ex:obs2 , ex:obs3
, ex:obs4 .

```

Listing 4: A Data Cube Observation Group.

```

# ... prefixes
# ... dataset triples

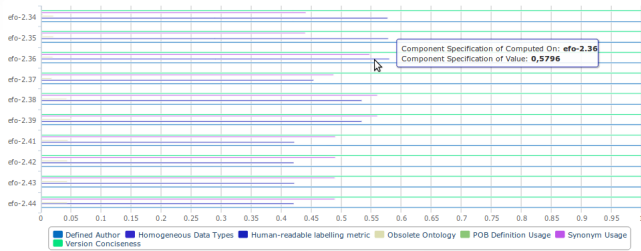
ex:qualityGraph1 a daq:QualityGraph ;
qb:structure daq:dsd .

ex:qualityGraph1 {

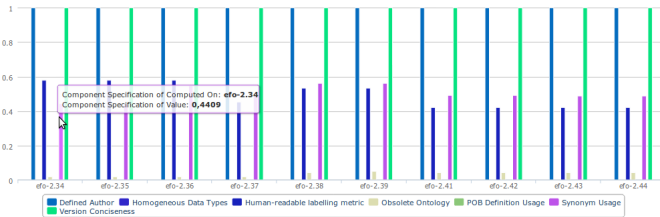
  # ... quality triples
  ex:accessibilityCategory
  a dqm:Accessibility ;
  dqm:hasAvailabilityDimension
  ex:availabilityDimension .

  ex:availabilityDimension
  a dqm:Availability ;
  dqm:hasEndPointAvailabilityMetric
  ex:endPointMetric ;
  dqm:hasRDFAvailabilityMetric
  ex:rdfAvailMetric .
}

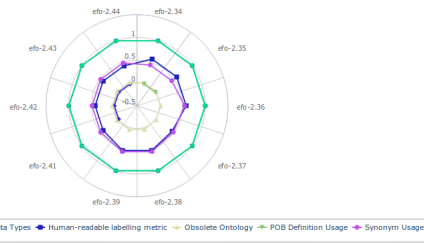
```



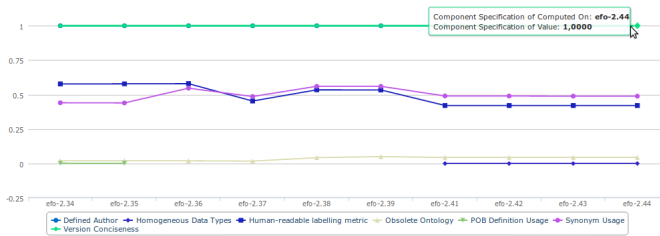
(a) Horizontal Bar Chart



(b) Vertical Bar Chart



(c) Radar Chart



(d) Lines Plot

Figure 3: Visualising Quality Metadata

```

ex:endPointMetric
  a dqm:EndPointAvailabilityMetric ;
  daq:hasObservation ex:obs1, ex:obs2 .

ex:obs1 a qb:Observation ;
  daq:computedOn <efo-2.43> ;
  dc:date "2014-01-23T14:53:00"^^xsd:
    dateTime ;
  daq:value "1.0"^^xsd:double ;
  daq:metric ex:endPointMetric ;
  qb:dataSet ex:qualityGraph1 .

ex:obs2 a qb:Observation ;
  daq:computedOn <efo-2.44> ;
  dc:date "2014-01-25T14:53:00"^^xsd:
    dateTime ;
  daq:value "1.0"^^xsd:double ;
  daq:metric ex:endPointMetric ;
  qb:dataSet ex:qualityGraph1 .

```

```

ex:rdfAvailMetric
  a dqm:RDFAvailabilityMetric ;
  daq:hasObservation ex:obs3, ex:obs4 .

```

```

ex:obs3 a qb:Observation ;
  daq:computedOn <efo-2.43> ;
  dc:date "2014-01-23T14:53:01"^^xsd:
    dateTime ;
  daq:value "1.0"^^xsd:double ;
  daq:metric ex:rdfAvailMetric ;
  qb:dataSet ex:qualityGraph1 .

```

```

ex:obs4 a qb:Observation ;
  daq:computedOn <efo-2.44> ;
  dc:date "2014-01-25T14:53:01"^^xsd:
    dateTime ;
  daq:value "0.0"^^xsd:double ;
  daq:metric ex:rdfAvailMetric ;
  qb:dataSet ex:qualityGraph1 .

```

```

# ... more quality triples
}

```

Listing 5: A Dataset Quality Graph

5. IMPLEMENTATION, EVALUATION, AND THE DIACHRON PROJECT

We are currently implementing the following support for automated quality assessment and analytics: (i) Luzzu, a generic quality assessment framework based on the daQ ontology, implemented in Java using the Jena RDF libraries and offering a web service interface¹⁰, and (ii) so far 38 concrete quality metrics over linked open datasets written by ourselves or by collaborators in the DIACHRON project. They are described in an extension of the daQ ontology¹¹ and implemented on top of the Luzzu framework¹².

The DIACHRON project (“Managing the Evolution and Preservation of the Data Web”¹³) combines several of the use cases mentioned so far in the application domains of open data, enterprise intranets and scientific data. DIACHRON’s central cataloguing and archiving hub is intended to host datasets throughout several stages of their life-cycle (cf. [4] for a general introduction to the linked data life-cycle), mainly evolution, archiving, provenance, annotation, citation and quality assessment. Beyond the implementation currently in progress, we will, together with the project partners, implement a web frontend that pulls dataset metadata from installations of the CKAN data portal system in order to

- allow data publishers to perform quality assessment on datasets, including the generation of quality metadata, visual analytics and support with cleaning datasets that are affected by quality problems, and to
- allow data consumers to filter and rank datasets by multiple quality dimensions.

The daQ ontology is the common language understood by these components.

Further progress with this implementation will enable us to evaluate daQ and its concrete extensions with regard to questions such

¹⁰<https://github.com/EIS-Bonn/Luzzu>

¹¹data quality metrics (DQM; <http://purl.org/eis/vocab/dqm>)

¹²<https://github.com/diachron/quality>

¹³<http://diachron-fp7.eu>

as how much time it takes to compute quality metrics for big datasets, to what extent the ontological foundation facilitates the agile development of end-user frontends, how well our visualisation of multi-dimensional quality metadata enables publishers to improve the quality of their data, and how well quality-based dataset filtering and ranking supports consumers in finding high-quality data.

6. RELATED WORK

To the best of our knowledge, the Data Quality Management (DQM) vocabulary [14] is the only one comparable to our approach. Fürber et al. propose an OWL ontology that primarily represents data requirements, i.e. what quality requirements or rules should be defined for the data. Such rules can be defined by the user herself, and the authors present SPARQL queries that “execute” the definitions of the requirements to compute metrics values. Unlike our daQ model, the DQM defines a number of classes that can be used to represent a data quality rule. Similarly, properties for defining rules and other generic properties such as the rule creator are specified. The daQ model allows for integrating such DQM rule definitions using the *daq:requires* abstract property, but we consider the definition of rules out of daQ’s own scope. Also, the proposed daQ ontology gives the freedom to the user to define and implement any metrics required for a certain application domain.

Our design approach is inspired by the digital.me Context Ontology (DCON¹⁴) [3]. Attard et al. present a structured three-level representation of context elements (Aspect-Element-Attributes). The DCON ontology instances are stored as Named Graphs in a user’s Personal Information Model. The three levels are abstract concepts, which can be extended to represent different context aspects in a concrete ubiquitous computing situation.

Ermilov et. al [12] present a framework calculating comprehensive statistics on Linked Open datasets. This statistical metadata has an underlying ontology based on VoID (see below) and Data Cube. The authors argue that since the VoID ontology was not sufficient to cover the required statistical concepts, the Data Cube extension was required. This extension also presented an opportunity to the authors to represent such statistical data using arbitrary attribute dimensions. Motivated by this work, we model resources and their calculated metrics as Data Cube observations, allowing us to represent quality metadata in a multi-dimensional manner.

The W3C recommends VoID and the Data Catalog Vocabulary (DCAT [20]) for metadata describing datasets. The “Vocabulary of Interlinked Datasets” (VoID) ontology allows the high-level description of a dataset and its links [1, 2]. On the other hand, DCAT describes datasets in data catalogs, which increase discovery, allow easy interoperability between data catalogs and enable digital preservation. With the daQ ontology, we aim to extend what these two ontologies have managed to achieve for datasets in general to the specific aspect of data *quality*: enabling the discovery of a good quality (fit to use) datasets by providing the facility to “stamp” a dataset with quality metadata.

7. CONCLUDING REMARKS

We presented the Dataset Quality Ontology (daQ), a core vocabulary for representing quality benchmarking metadata of linked open datasets, which makes use of the Data Cube Vocabulary. In Section 2 we presented a number of use cases that motivated our idea. These included analysis of data versions, dataset retrieval, automatic link identification based on the quality of data entities, and finally the extension of the five star open data scheme by a star

for quality. The precise definition of these use cases assisted in the development of the daQ ontology, which involved reuse of the Data Cube (Section 3).

The ontology is progressing in a fast pace, and further developments to cover the intended use cases are also in the pipeline. The next iteration phase is to further model the daQ ontology to cover the provenance aspect of quality metadata. The development and extension of new concepts to the daQ ontology should ensure that (i) high standards are kept, and (ii) that the ontology is not bloated out of proportion – i.e. keeping with the main idea of a light-weight quality assessment framework.

Currently, using daQ, we are in the process of implementing (Sections 4 and 5) a number of domain-specific and domain-independent metrics, following a survey of linked data quality metrics [24]. As quality metadata describes the dataset on which quality was calculated, it makes sense to maintain it close to the dataset. In LOD settings we consider it most reasonable to embed the quality metadata into the dataset itself, which is possible thanks to the named graphs approach we chose.

We also demonstrated specific advantages of basing daQ on the Data Cube Vocabulary: quality metadata can be visualised using available Data Cube enabled applications such as CubeViz, and observations can be grouped together automatically using the daQ three level abstract layer.

One of the tools which will support the daQ framework is the DIACHRON platform. This platform will enable consumers to rank and filter datasets by quality. Having tools and platforms supporting the daQ will finally allow us to test and evaluate the ontology thoroughly, to see whether the daQ (and the quality metadata) itself is of a high quality, i.e. fit for use.

8. ACKNOWLEDGMENTS

This work is supported by the European Commission under the Seventh Framework Program FP7 grant 601043 (<http://diachron-fp7.eu>). We would like to thank Santiago Londoño, who carried out the visual analytics of our sample quality metadata.

References

1. Alexander, K. et al. Describing Linked Datasets. On the Design and Usage of void, the “Vocabulary of Interlinked Datasets”. In: *Linked Data on the Web (LDOW)*. (Madrid, Spain, 20th Apr. 2009). Ed. by C. Bizer et al. CEUR Workshop Proceedings 538. Aachen, Apr. 2009. <http://CEUR-WS.org/Vol-538>.
2. Alexander, K. et al. Describing Linked Datasets with the VoID Vocabulary. W3C Interest Group Note. World Wide Web Consortium (W3C), 3rd Mar. 2011. <http://www.w3.org/TR/2011/NOTE-void-20110303/>.
3. Attard, J. et al. Ontology-based Situation Recognition for Context-aware Systems. In: *Proceedings of the 9th International Conference on Semantic Systems (I-SEMANTICS)*. (Graz, Austria, 4th–6th Sept. 2013). Ed. by M. Sabou et al. New York, NY, USA: ACM, 2013, pp. 113–120. <http://dl.acm.org/citation.cfm?id=2506182>.
4. Auer, S. et al. Managing the life-cycle of Linked Data with the LOD2 Stack. In: *The Semantic Web (Part II)*. 11th International Semantic Web Conference (ISWC) (Boston, MA, USA, 11th–15th Nov. 2012). Ed. by P. Cudré-Mauroux et al. LNCS 7650. Springer, 2012. <http://iswc2012.semanticweb.org/sites/default/files/76500001.pdf>.

¹⁴<http://www.semanticdesktop.org/ontologies/dcon/>

5. Berners-Lee, T. et al. Tabulator: Exploring and Analyzing linked data on the Semantic Web. English. In: *Proceedings of the 3rd International Semantic Web User Interaction Workshop (SWUI06)*. Nov. 2006.
6. Bizer, C. Quality-Driven Information Filtering in the Context of Web-Based Information Systems. PhD thesis. FU Berlin, Mar. 2007. http://www.diss.fu-berlin.de/diss/receive/FUDISS_thesis_000000002736.
7. Carroll, J. J. et al. Named Graphs, Provenance and Trust. In: *Proceedings of the 14th WWW conference*. (Chiba, Japan, 10th–14th May 2005). Ed. by A. Ellis, T. Hagino. ACM Press, 2005, pp. 613–622.
8. Carroll, J. J. et al. Semantic Web Publishing using Named Graphs. In: *ISWC Workshop on Trust, Security, and Reputation on the Semantic Web*. Ed. by J. Golbeck et al. Vol. 127. CEUR Workshop Proceedings. CEUR-WS.org, 9th May 2005. <http://dblp.uni-trier.de/db/conf/semweb/iswc2004trust.html#Carroll1BHS04>.
9. Crosby, P. B. Quality is Free. The Art of Making Quality Certain. Mentor book. McGraw-Hill, 1979. http://books.google.ie/books?id=bR_LnQEACAAJ.
10. Cyganiak, R., Reynolds, D., Tennison, J. The RDF Data Cube Vocabulary. W3C Recommendation. World Wide Web Consortium (W3C), 16th Jan. 2014. <http://www.w3.org/TR/2014/REC-vocab-data-cube-20140116/>.
11. Debattista, J., Lange, C., Auer, S. daQ, an Ontology for Dataset Quality Information. In: *Linked Data on the Web (LDOW)*. (Seoul, 8th Apr. 2014). Ed. by C. Bizer et al. 2014. <http://events.linkedata.org/ldow2014/>.
12. Ermilov, I. et al. Linked Open Data Statistics: Collection and Exploitation. In: *Proceedings of the 4th Conference on Knowledge Engineering and Semantic Web*. 2013. http://svn.aksw.org/papers/2013/KESW_LODStats_Demo/public.pdf.
13. Flemming, A. Quality characteristics of linked data publishing datasources. MA thesis. Humboldt-Universität zu Berlin, Institut für Informatik, 2011.
14. Fürber, C., Hepp, M. Towards a Vocabulary for Data Quality Management in Semantic Web Architectures. In: *Proceedings of the 1st International Workshop on Linked Web Data Management (LDWM)*. (Uppsala, Sweden). New York, NY, USA: ACM, 2011, pp. 1–8. <http://doi.acm.org/10.1145/1966901.1966903>.
15. Harth, A. VisiNav: A system for visual search and navigation on web data. In: *Web Semantics: Science, Services and Agents on the World Wide Web 8(4)* (2010). *Semantic Web Challenge 2009 User Interaction in Semantic Web research*, pp. 348–354. <http://www.sciencedirect.com/science/article/pii/S1570826810000600>.
16. Heim, P., Ziegler, J., Lohmann, S. gFacet: A Browser for the Web of Data. In: *Proceedings of the International Workshop on Interacting with Multimedia Content in the Social Semantic Web (IMC-SSW 2008)*. Vol. 417. CEUR Workshop Proceedings. Aachen, 2008, pp. 49–58. <http://CEUR-WS.org/Vol-417>.
17. Hogan, A. et al. An empirical survey of Linked Data conformance. In: *Web Semantics: Science, Services and Agents on the World Wide Web 14(0)* (2012): *Special Issue on Dealing with the Messiness of the Web of Data*, pp. 14–44. <http://www.sciencedirect.com/science/article/pii/S1570826812000352>.
18. Hogan, A. et al. Searching and browsing Linked Data with SWSE: The Semantic Web Search Engine. In: *Web Semantics: Science, Services and Agents on the World Wide Web 9(4)* (2011): *Special issue on Semantic Search*, pp. 365–401. <http://www.sciencedirect.com/science/article/pii/S1570826811000473>.
19. Juran, J. M. *Juran's Quality Control Handbook*. 4th ed. McGraw-Hill, 1974. <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0070331766>.
20. Maali, F., Erickson, J., Archer, P. Data Catalog Vocabulary (DCAT). W3C Recommendation. World Wide Web Consortium (W3C), 16th Jan. 2014. <http://www.w3.org/TR/2014/REC-vocab-dcat-20140116/>.
21. Ngonga Ngomo, A.-C., Auer, S. LIMES – A Time-Efficient Approach for Large-Scale Link Discovery on the Web of Data. In: *Proceedings of the 22nd joint conference on artificial intelligence (IJCAI); Volume Three*. AAAI Press, 2011, pp. 2312–2317.
22. Pirsig, R. *Zen and the Art of Motorcycle Maintenance: An Inquiry Into Values*. Essence (Philosophy). Vintage, 1974. <http://books.google.ie/books?id=M69poeV1UhoC>.
23. Volz, J. et al. Discovering and Maintaining Links on the Web of Data. In: *ISWC*. Vol. 5823. LNCS. Springer, 2009, pp. 650–665.
24. Zaveri, A. et al. Quality Assessment Methodologies for Linked Open Data. In: *Semantic Web Journal* (2014). This article is still under review. <http://www.semantic-web-journal.net/content/quality-assessment-methodologies-linked-data-survey>.