

Linked Environment Data: SCOVO-fying the Environment Specimen Bank

Thomas Bandholtz
innoQ Deutschland GmbH
Halskestr. 17
40880 Ratingen, Germany
+49 178 4049387

Till Schulte-Coerne
innoQ Deutschland GmbH
Halskestr. 17
40880 Ratingen, Germany
+49 1515 7132845

Maria R ther
Federal Environment Agency
Corrensplatz 1
14195 Berlin
+49 30 89031503

thomas.bandholtz@innoq.com till.schulte-coerne@innoq.com maria.ruether@uba.de

ABSTRACT (of submission)

In this paper, we describe the RDF schema and external linkage of the German Environmental Specimen Bank (ESB). The schema is based on SCOVO and SKOS with only little but powerful extensions. The submission raises several strategic SCOVO design issues that are currently subject of discussion in the Linked Data community. The proposed solutions may change in some details until the date of camera-ready delivery.

Categories and Subject Descriptors

E.1 [Data Structures]: Distributed Data Structures

H.5 [Information Interfaces and Presentation]

J.3 [Life and Medical Sciences]: Health

General Terms

Design, Standardization.

Keywords

Triplification, Vocabularies, Linked Government Data, Statistical Data, Environmental Data, Content Negotiation, SPARQL.

1. INTRODUCTION

Currently several projects at the German Federal Environment Agency (UBA) begin with the design and implementation of a public data network that is technologically based on Linked Data. The first ones will be the Environmental Specimen Bank (ESB) and the Semantic Network Service (SNS); the inclusion of the Dioxin Database and the Joint Substance Data Pool of the German Federal Government and the German Federal States (GSBL) is currently under discussion. The undertaking is an international cooperation jointly with the Ecoterm Initiative and the European Environment Agency (EEA), and it is envisioned to include the partners of the International Environmental Specimen Bank Group (IESB).

These projects and partners provide the key instruments in the field of environmental observation that enable the long-term analysis of substance exposure of humans and the environment.

2. RDF Schema for the ESB

The ESB reports the accumulation of pollutants/substances in test subjects at specific places with respect to time.

Basically, this can be expressed with the Statistical Core Vocabulary (SCOVO) [1].

However, the simple SCOVO model needs some tailoring and extension in order to represent the complexity of the ESB model. One aspect is about representing hierarchical dimension parameters by linking SCOVO and SKOS, the other is about tailoring subclasses of `scv:Dimension` according to the ESB model.

2.1 Specializing `scv:Dimension`

Basically, SCOVO is modelled around `scv:Item` which links to several `scv:Dimension` individuals and a `scv:Dataset`, and finally it provides a simple `rdf:value` Literal.

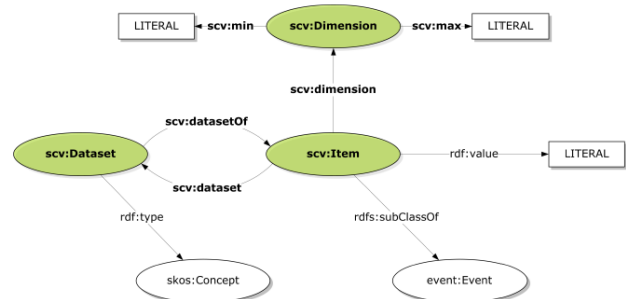


Figure 1. Original SCOVO Model (from [1]).

In the referred SCOVO publication [1], the flight on-time arrival example has three dimensions *OnTime*, *Airport*, and *TimePeriod*. The value is given as a simple number like in the following example:

```
ex:myItem a scv:Item ;
  rdf:value 74 ;
  scv:dataset ex:ontimeFlights ;
  scv:dimension ex:Q12006 ;
  scv:dimension ex:onTimeArrival ;
  scv:dimension ex:AtlantaHartsfield .
```

Example 1. Original SCOVO example (from [1]).

Compared to this, ESB has four major dimensions:

- *SpecimenType*
- *Analyte*
- *Location*

- *TimeReference*

So far, this is structurally similar. The adapted example would look as follows:

```
esb:specimen a scv:Item ;
  rdf:value 2.1672 ;
  scv:dataset esb:chemicalExposure ;
  scv:dimension esb:breamLiver ;
  scv:dimension esb:Copper ;
  scv:dimension esb:lowerRhine.Bimmen ;
  scv:dimension esb:year2008 .
```

Example 2. Adapted ESB example.

2.2 Data Type Considerations

Still something is missing: what does `rdf:value 2.1672` mean? In Example 1, the value 74 means “74 percent of on-time arrivals ...” but this has not been stated explicitly in the data, so the value might be easily misunderstood as meaning a *count* of on-time arrivals.

Tauberer [2] tries to solve this for the 2000 U.S. Census by providing the value in a string such as “122283145776 m^2”, for square meter which he calls “a strange half-number/half-unit literal value” himself. Obviously, this might also be expressed as a typed literal¹ such as

```
"122283145776"^^xsd:double .
```

But what to do if the publisher wants to specify that the value should be interpreted as an `xsd:double` the same time?

In the ESB, we have several complex units of measurement (such as “microgram per litre wet weight”, *and* we have several summary statistics (such as “geometric mean” or “standard deviation”) for most values. The first approach of the ESB project was to extend the `rdf:value` attribute, thereby allowing to provide a complex data structure as each item’s value. However, this would result in reduced flexibility and proprietary processing requirements. The preferred solution is a further extension of the dimension set by adding two `scv:Dimension` subclasses:

- *UnitOfMeasurement* and
- *SummaryStatistic*.

From a data warehouse point of view, those two might not be understood as dimensions but as metadata of different `scv:DataSet` individuals. In our opinion, this design would result in another loss of flexibility in the SCOVO representation.

2.3 Item Filtering vs. Dimension Hierarchies

Furthermore, the ESB supports anamnesis filters (such as “gender”, “smoker”, “amalgam fillers”, etc) for human test subjects. The overall value of a human sampling collective can be split in subgroups by such filters.

Tauberer [2] has a similar structure in the 2000 U.S. Census, and his design (when translated to SCOVO) is a hierarchy of datasets, where each dataset is bound to one of these “filters”. Alternatively, this could be modelled as another hierarchy level of the dimension (see 3 below). However, we find this too

restrictive: If you have a dataset for smokers and another one for females, how then to query the values of female smokers?

For this reason, we added another dimension.

- *AnamnesisFilter*.

In summary, a complete example of an ESB item might look like:

```
esb:specimen a scv:Item ;
  rdf:value 0.1800 ;
  scv:dataset esb:chemicalExposure ;
  scv:dimension esb:24hSamplingUrine ;
  scv:dimension esb:Greifswald ;
  scv:dimension esb:Mercury ;
  scv:dimension esb:year2008 ;
  scv:dimension esb:geometricMean ;
  scv:dimension esb:mgPerLitreWetWeight ;
  scv:dimension esb:female ;
  scv:dimension esb:smoker .
```

Example 3. Extended ESB example.

2.4 Sub-Properties of `scv:dimension`

Example 3 shows concise information but somehow difficult to read by itself, as we have eight `scv:dimension` properties without a local indication of the respective dimension. Following the example from [1], such information is not contained in the `scv:item` itself. Somewhere else in the `esb: namespace` you will find statements like:

```
esb:LocationDimension a owl:Class ;
  rdfs:subClassOf scovo:Dimension .

esb:Greifswald a esb:LocationDimension .
```

Example 4. Dimension subclass and instance.

In order to understand the meaning of each `scv:dimension` property, one has to look up the type definition of the provided individual in the object role. The SCOVO publication [1] contains a SPARQL example for such a use case.

It *might* be more convenient to setup sub-properties of `scv:dimension`, such as `esb:location`.

```
esb:location a owl:ObjectProperty ;
  rdfs:subPropertyOf scv:dimension .
```

Example 5. `scv:dimension` subProperty example.

This would result in a more readable item serialization as given in the following example:

```
esb:specimen a scv:Item ;
  rdf:value 0.1800 ;
  scv:dataset esb:chemicalExposure ;
  scv:specimenType esb:24hSamplingUrine ;
  scv:location esb:Greifswald ;
  scv:analyte esb:Mercury ;
  scv:timeReference esb:year2008 ;
  scv:summaryStat esb:geometricMean ;
  scv:uom esb:mgPerLitreWetWeight ;
  scv:anamnesisFilter esb:female .
  scv:anamnesisFilter esb:smoker .
```

Example 6. More expressive ESB item example.

¹ <http://www.w3.org/TR/rdf-concepts/#dfn-typed-literal>

2.5 scv:Item Subclasses

Spinning this out, one might define subclasses of `scv:Item` and use them as the domain of dedicated properties like those given in example 6. This could bind a certain item type to an intended set of dimension types, including cardinality constraints. This idea is somehow seductive from an object-oriented point of view, but Linked Data is bound to the open world assumption, and Linked Data is only about publishing. One will need such constraints when data gets entered but not when it is simply published or queried. In the Linked Data context, “valid” combinations of dimension subclasses are exactly those which are used in the published `scv:Item` individuals. Such supported combinations can be explored by dedicated SPARQL queries.

(Note at the time of submission: the ESB project is implementing Example 3. Examples 4, 5 and 6 are only provided for discussion of possible future enhancements. This might change before camera-ready delivery of this paper).

2.6 ESB Datasets

The definition `scv:Dataset` is rather vague: “a dataset, representing the container of some data, such as a table holding some data in its cells”.

From a data warehouse point of view, `scv:DataSet` corresponds to the “fact” table. In this table, a specific set of dimension individuals makes up the primary key, and this key is bound to the corresponding measured value. Each `scv:Item` individual corresponds to one row (not to one cell) in such a fact table.

In the discussion about using SCOVO in the Vocabulary of Interlinked Datasets (VOID), there is one remarkable statement by Richard Cyganiak:

“In SCOVO, `scovo:Items` are grouped into `scovo:Datasets`, and there seems to be an implicit assumption that all items in such a dataset share the same dimensions.”²

Strictly speaking and referring to Example 6, Richards does not mean “share the same dimensions” but “share the same set of `scv:dimension` sub-properties”.

Taken this and the design explained before, we have only two of such different `scv:DataSet` individuals in the ESB:

- *chemicalExposure*
- *biometrics*.

Chemical exposure describes the amount of some substance which has been measured in some human or environmental test subjects, while biometrics describes certain different aspects of such test subjects (age, gender, weight, etc.), so there is a different set of analytes in use in the biometric dataset.

This division is somehow inherited from the ESB staff’s culture: they are used to differentiate between those two since more than two decades. From the current RDF point of view: what makes the difference between measuring the average weight vs. the chemical exposure of the same specimen collective? Do we need the `scv:Dataset` construct at all? Couldn’t we provide all statistical data in only one dataset with each item simply identified by a set of dimension references?

² <http://code.google.com/p/void-impl/issues/detail?id=18>

In order to carry on with this issue, we have to take a closer look at the structure of `scv:Dimension` and the dimension use case in the ESB.

3. ESB Dimensions, SCOVO, and SKOS

SCOVO offers several “hooks” [1] for linking statistical data with domain ontologies, two of them are “subclassing the SCOVO dimension class”, and the “built-in support” for `skos:Concept`. We will discuss both approaches in the following sub-chapters.

3.1 SKOS for Dimension Hierarchies

Figure 1, taken from the original publication of SCOVO [1], shows `skos:Concept` as the “type” of `scv:Dataset`, “in order to allow hooking into a categorisation scheme” [1]. First of all, both `skos:Concept` and `scv:Dataset` are classes, so we would propose to make `scv:Dataset` subclass of `skos:Concept`, not a `skos:Concept` individual.

What’s more important: we think that SKOS is helpful for structuring dimensions rather than datasets. In most statistical data collections (just like in the ESB), the dimension individuals, such as locations, are organized in a hierarchy, and there is an individual hierarchy for each dimension type. The ESB Web application provides a hierarchy of “profile” pages for each dimension, so we already have a linked tree of information resources. This tree can simply be rendered in SKOS to provide some simple but valid domain ontology.

```
esb:dimensions a skos:ConceptScheme ;
  skos:hasTopConcept esb:specimenType .

esb:specimenType a skos:Concept ;
  skos:narrower esb:limneticSample .

esb:limneticSample a skos:Concept ;
  skos:narrower esb:bream .

esb:bream a skos:Concept ;
  skos:narrower esb:breamLiver .
```

Example 7. `skos:ConceptScheme` for ESB dimensions.

In the SCOVO scheme, `scv:dimension` has no explicit range, so we can simply use `skos:Concept` individuals in the object role of `scv:dimension` property statements like in the previous examples (see `esb:breamLiver` in Example 2).

3.2 Subclassing `scv:Dimension`

In the SKOS approach of the previous sub-chapter, the ESB specific dimensions appear as top concepts (`skos:hasTopConcept`) of a mono-hierarchical tree.

Now we want to model them as subclasses of `scv:Dimension`.

```
esb:SpecimenType
  rdfs:subClassOf scv:Dimension .
esb:Analyte
  rdfs:subClassOf scv:Dimension .
esb:Location
  rdfs:subClassOf scv:Dimension .
esb:TimeReference
  rdfs:subClassOf scv:Dimension .
```

Example 8. Subclasses for ESB dimensions.

Now we can make the bream liver an instance of `esb:SpecimenType` in two ways: either by a direct assertion, or by

a specific range of `scv:dimension` sub-properties. Referring to examples 5 and 6, we simply add a range to the subproperty, and from this would be inferred that the object belongs to a specific dimension subclass.

```
esb: SpecimenType
  rdfs:subClassOf scv:Dimension .

esb:specimenType a owl:ObjectProperty ;
  rdfs:subPropertyOf scv:dimension ;
  rdfsRange esb:SpecimenType .

esb:specimen a scv:Item ;
  scv:specimenType esb:breamLiver .
```

Example 9. Using `rdfs:Range` in subproperties.

3.3 Bringing Both Approaches Together

So far, this results in subclasses of `scv:Dimension` (such as `esb:SpecimenType`, `esb:Analyte`, etc.) on one side, and a `skos:ConceptScheme` tree with the corresponding top concepts on the other. Which of the `skos:Concept` individuals belongs to what `scv:Dimension` subclass is only inferred by its usage in some `scv:Item` instance. If we want to make this more explicit, we could use a double class assertion (`skos:Concept` *and* the `scv:Dimension` subclass). Alternatively, we make `scv:Dimension` a subclass of `skos:Concepts`, and so the SKOS tree (Example 7) and the dimension subclassing (Example 8) integrates to the following.

```
scv:Dimension
  rdfs:subClassOf skos:Concept .

esb: SpecimenType
  rdfs:subClassOf scv:Dimension .

esb:dimensions a skos:ConceptScheme ;
  skos:hasTopConcept esb:specimenType .

esb:specimenType a esb: SpecimenType ;
  skos:narrower esb:limneticSample .

esb: limneticSample a esb: SpecimenType ;
  skos:narrower esb:bream .

esb: bream a esb: SpecimenType ;
  skos:narrower esb:breamLiver .
```

Example 10. Dimension as subclass of Concept.

In this case, we do not need the subproperties of `scv:dimension` (Example 9), as this information can be taken from the class assertion of the respective object. This brings us back to the simple pattern of Example 1 or 3, but now we have a SKOS tree of dimension individuals at the same time.

4. ESB External Linkage

Hausenblas et al. [1] recommend linking the dimension individuals to some domain ontology which has been published in the Web, or to the more general DBpedia using `owl:sameAs`.

In the case of the ESB, we have a built-in domain ontology which is published together with the statistic data; however, this does not make a big difference: we are linking the ESB domain ontology to several external vocabularies anyway.

This happens in an international context of cooperating governmental authorities, such as the eTerminology Workshop [5]

or the Ecoterm Group³ with members from many European countries and the US. These authorities have started setting up a trusted network of domain ontologies [4] from different environmental facets and multiple languages. Such ontologies will be published in RDF (namely SKOS) and described in the Vocabulary of Interlinked Data (VOID). Some of them have already gone live or will be going live in the next months.

Interlinking of the ESB will focus on these vocabularies and further environmental measurement data to be published in the Web as well. Governmental authorities are quite reserved against non-governmental vocabularies such as DBpedia or Geonames. Publications of the agencies usually are subject to legal obligations and thus sensitive about the provenance of the sources. Currently there is no decision about such linkage.

In the following, we will give a short overview of the intended linkage and the structure and state of each target.

4.1 Semantic Network Service of the Federal Environment Agency

Semantic Network Service (SNS)⁴ is maintained by the agency since 2003. SNS includes a thesaurus (UMTHES), a gazetteer and a chronicle with occasional interlinkage among each other. All three are currently available in the XML Topic Maps⁵ format. A first draft of an RDF vocabulary for SNS has been presented in 2006, but until today only the thesaurus has been migrated into a SKOS-XL representation.

The gazetteer contains the ESB sampling areas (among others). The RDF schema may extend the Geonames ontology⁶ by a domain specific type system and some properties for spatial intersections between individuals which are not organized in a hierarchy (e.g. river crosses city).

The chronicle will be interlinked with the ESB time references. For example, the Elbe flood in 2002 had a considerable influence on the exposure of limnetic samples in this area in the following year. Such influences can be expressed by linking the respective statistical items (in this case not the dimensions) to this event.

The chronicle can be published in the Linked Event Ontology⁷ with a domain specific type system, and a somehow simplified pattern of describing the time references.

4.2 More Vocabularies from the Ecoterm Space

The Ecoterm group members represent more than 20 different vocabularies but we will focus on only three for the beginning.

The European reference vocabulary since years is the GEneral Multilingual Environmental Thesaurus (GEMET)⁸, maintained by the European Environment Agency (EEA). GEMET has been one of the first SKOS use cases in 2004 and is still available in this

³ <http://ecoterm.infointl.com>

⁴ <http://www.semantic-network.de>

⁵ <http://isotopicmaps.org>

⁶ <http://www.geonames.org/ontology>

⁷ <http://linkedevents.org/ontology>

⁸ <http://www.eionet.europa.eu/gemet>

serialization. Since last year it is also published using the Linked Data technical patterns.

UMTHES is already linked with GEMET, so we do not need any direct linkage between ESB and GEMET. GEMET is much smaller than UMTHES (which has been one of its sources) but it is available in 29 languages.

The second vocabulary from the EEA is the EUNIS biodiversity database⁹, with a focus on species. EUNIS has been published in RDF early this year, using several properties from the Darwin Core vocabulary¹⁰.

The third example is the Environmental Applications Reference Thesaurus (EARTH)¹¹ from Italy, which has been published in SKOS and linked with EUNIS as well.

4.3 Environmental Information Systems

All those partner vocabularies mentioned above are reference vocabularies, not data about the state of the environment. ESB provides such data, and we will link it to more data. Currently there is some environmental data published within the Data-gov projects of the US¹² and UK¹³.

Within the realm of the Federal Environment Agency in Germany, there are two such systems which may follow in the near future:

- The Dioxin Database¹⁴ is built on the same model as the ESB and currently accessible through Web Services. Once we have established the ESB in Linked Data, the same patterns can be adopted to this database.
- The Joint Substance Data Pool of the German Federal Government and the German Federal States (GSBL)¹⁵ contains detailed dossiers about substances, some of which occur as analytes in the ESB. The ESB substances have already been linked to these dossiers, but the GSBL is based on a completely different database model, and it will take some extra effort to get this published in RDF.

Finally, there is ongoing work to establish an International Environmental Specimen Bank Group (IESB)¹⁶ with partners from Scandinavian countries, Canada, Japan, South Korea, France, UK, and US.

4.4 owl:sameAs Consideration

Most Linked Data contributors (including [1]) tend to use owl:sameAs for such cross-references. Following the OWL reference, “Such an owl:sameAs statement indicates that two URI references actually refer to the same thing: the individuals have the same “identity”.”¹⁷

⁹ <http://eunis.eea.europa.eu/>

¹⁰ <http://rs.tdwg.org/dwc/>

¹¹ http://uta.iiia.cnr.it/earth_eng.htm

¹² http://data-gov.tw.rpi.edu/wiki/The_Data-gov_Wiki

¹³ <http://data.gov.uk/>

¹⁴ <http://www.pop-dioxindb.de/>

¹⁵ http://www.gsbl.de/eng_home.htm

¹⁶ <http://www.inter-esb.org/>

¹⁷ <http://www.w3.org/TR/owl-ref/#sameAs-def>

This has two implications: (1) make sure about exact semantic identity, and (2) a reasoner will merge both individuals, including their class assignments.

- (1) There are cases where two things appear to be the same but are not. In the ESB, for example, two different species, the bream and the deer, have a skos:narrower, which might be named “liver”. Someone might state that both are the same and link them to some anatomy vocabulary talking about the abstract concept of “liver”. Most certainly, however, a deer liver is not the same as a bream liver. SKOS provides a more subtle set of mapping relations¹⁸ skos:closeMatch, skos:exactMatch, skos:broadMatch, skos:narrowMatch and skos:relatedMatch. Using these, one might express that esb:breamLiver and esb:deerLiver both have a skos:broader relation to anatomy:liver, which would come much closer to the domain knowledge than owl:sameAs.
- (2) This brings up the merging issue. Using owl:sameAs, the esb:breamLiver (a skos:Concept) and anatomy:liver (some other class) would be treated as just one individual which now belongs to both classes. In the respective schemas, there may be collisions between the two class definitions. In the case of the SKOS mapping relations, domain and range of these properties would make both individuals instances of skos:Concept. Such implications may be ignored as long as linked data gets simply browsed, but they will result in semantic collisions as soon as a reasoner or some more specific agent is traversing the Web of Data.

The ESB provides scientific data which can give input for further research and integration with different sources, so we will link it carefully.

5. Exploring the ESB

Feigenbaum [5] describes his approach of exploring the Eurostat LOD contribution (which has not been published in SCOVO) by querying the graph with SPARQL.

This is something that will work with any published RDF data that supports a SPARQL endpoint. ESB will support multiple approaches to explore its structure and data.

5.1 Starting at the ESB Homepage

One starting point is the homepage of the ESB Web application at www.umweltprobenbank.de. This application provides various information resources about the ESB, a tree of interlinked, human readable dimension profiles, and a comfortable HTML/JQuery database interface. Results are displayed in tables or diagrams, and they can be exported in CSV and Excel formats. We are planning to make most of this accessible in RDF in several ways:

1. Each dimension profile will support content negotiation as described in [6];
2. For human visitors with standard browsers and no linked data plug-in there will be direct links to the respective RDF representation;
3. The query dialog will provide the URI of the SPARQL access point and display the SPARQL version of each query.

¹⁸ <http://www.w3.org/TR/skos-reference/#mapping>

Complementing this, there will be a Linked Environment Data introduction page with a description of the project and links to the documented RDF schema, VOID metadata, and to the SPARQL endpoint. The VOID metadata will be linked to the Ecoterm project.

5.2 Coming from an Interlinked Vocabulary

Our intention is to have bi-directional links between the ESB and the partner systems listed in chapter 4. Whenever an agent explores one of these partner systems, she will detect outgoing links pointing to some related ESB dimensions or items. From here the agent can explore the linkage of this resource as described in subchapter 5.3.

Bi-directional links may not be realistic in some cases, as this requires some extra effort from the partner agencies. We can provide a concise list of the cross-references to minimize this effort but we cannot incorporate such links into the partner system ourselves.

If bi-directional links cannot be established, the upcoming Ecoterm platform may help: “a simple umbrella Web page would be created to aid in the management, promotion, and access to this network of linked data” [4]. The first step will be linking to the VOID metadata of all member vocabularies from this umbrella page. If each of these VOID datasets links back to the umbrella page at least, this would offer an indirect path to the interlinkage: from any `ex1:resource` to the `ex1: namespace` to the `ex1: VOID` to the Ecoterm platform to the `ex2 VOID` dataset which describes the linkage from `ex2` to `ex1`. Ecoterm could also host lists of cross-references. Currently this is still a vague idea which needs some clarification and agreement in detail.

5.3 Exploring the Neighbourhood of a Single ESB resource

There are various scenarios which lead to a single ESB resource as the starting point. The previous chapter described bi-directional linkage but you may also find such a link in a human readable document or an e-mail.

In most cases, this link will lead to an ESB dimension or item individual.

In case of a `scv:Dimension` individual, ESB offers different names in German and English, a short description in the domain context, and links to further information resources and related data. As we modelled `scv:Dimension` as a subclass of `skos:Concept`, you may traverse the dimension tree level by level. Using SPARQL, one may retrieve all statistical items that refer to this individual (which will be quite many), or inspect the existing combinations with different dimension subclasses first.

Having an item as the starting point, you may simply move on to explore each of the connected dimension individuals for a better understanding of the meaning of this item. Using SPARQL, you may drop one of the dimensions to get a list of related items, e.g. a time line, or a comparison of different locations at the same time.

These are only some simple examples. Being aware of the ESB schema, SPARQL gives you nearly unlimited options.

6. Conclusions and Ongoing Work

(Preliminary, will be updated in the camera-ready version).

This is the first time a German Authority is publishing measurement data as Linked Data.

The SCOVO model can be applied with some extensions.

ESB needs an extended set of dimension subclasses.

ESB uses explicit statements about units of measurement and summary statistics.

ESB integrates the domain vocabulary by making `scv:Dimension` subclass of `skos:Concept`, and so integrates a dimension tree.

Each dimension individual links to various information resources for extended human reading.

Those information resources link back to the RDF resource.

ESB will be carefully linked to multiple reference vocabularies from the Ecoterm space. Bi-directional links are aspired.

ESB will be linked to other data collections of environmental authorities.

ESB (Germany) initiates a close linkage between different international ESBs.

7. REFERENCES

- [1] Hausenblas, Michael; Halb, Wolfgang; Raimond, Yves; Feigenbaum, Lee; Ayers, Danny. 2009. SCOVO: Using Statistics on the Web of Data. ESWC 2009. <http://sw-app.org/pub/eswc09-inuse-scovo.pdf>.
- [2] Tauberer, Joshua. 2007. The 2000 U.S. Census: 1 Billion RDF Triples. <http://www.rdfabout.com/demo/census/>.
- [3] Bandholtz, Thomas et al.. 2009. Summary of W4 eEnvironment Terminology. e-envi2009. Integrating Environmental Knowledge in Europe. March 25-27, 2009 Prague. <http://www.e-envi2009.org/SummaryTerminologyW4.pdf>
- [4] Ecoterm Group. 2009. Report on the outcome of the Ecoterm V Workshop. UN Food and Agriculture Organization, Rome, Italy on 5-6 October 2009 http://eea.eionet.europa.eu/Public/irc/envirowindows/jad/library?l=/ecoinformatics_indicator/ecoterm_5-6102009
- [5] Feigenbaum, Lee. 2008. Modeling Statistics in RDF - A Survey and Discussion. TechnicaLee Speaking. http://www.thefigtrees.net/lee/blog/2008/03/modeling_statistics_in_rdf_a_s.html
- [6] Bizer, Chris; Cyganiak, Richard; Heath, Tom. 2007. How to Publish Linked Data on the Web. <http://www4.wiwiw.fu-berlin.de/bizer/pub/LinkedDataTutorial>