The Stimulus-Sensor-Observation Ontology Design Pattern and its Integration into the Semantic Sensor Network Ontology

Krzysztof Janowicz*, Michael Compton†

*GeoVISTA Center, Pennsylvania State University, USA
jano@psu.edu

†CSIRO ICT Centre, Canberra, Australia
michael.compton@csiro.au

Abstract. This paper presents an overview of ongoing work to develop a generic ontology design pattern for observation-based data on the Semantic Web. The core classes and relationships forming the pattern are discussed in detail and are aligned to the DOLCE foundational ontology to improve semantic interoperability and clarify the underlying ontological commitments. The pattern also forms the top-level of the the Semantic Sensor Network ontology developed by the W3C Semantic Sensor Network Incubator Group. The integration of both ontologies is discussed and directions of further work are pointed out.

1 Introduction and Motivation

Observations are at the core of empirical science. Sensors transform stimuli from the physical world into observations and thereby allow to reason about the observed properties of particular features of interest [1]. The notion and usage of sensors, however, has drastically changed over the last decades. These days, sensors are omnipresent and part of everyday digital equipment. While the size as well as costs for the production and deployment of sensors are decreasing, their processing power and efficiency is still increase. In the near future, self-organizing networks of sensors will be deployed as smart dust to monitor all kinds of events ranging from traffic jams to floods and other natural or artificial disasters. Additionally, with the advent of the Social Web and Volunteered Geographic Information (VGI) citizens provide an additional and valuable source of observations [2]. Consequently, sensor data is available in near-real time for arbitrary spatial, temporal, and thematic resolutions.

The core motivation of the Sensor Web Enablement (SWE) initiative of the Open Geospatial Consortium is to make data about sensors and their observations available on the Web. Besides storage and access, the initiative also focuses on the tasking of sensors. Hence, SWE supports the full workflow from the selection of specific sensors, their configuration and deployment, over the collection of observations, up to changing their location and spatial distribution in a Web-centric framework. This so-called Sensor Web includes services such as the Sensor Observation Service (SOS), the Sensor Planning Service (SPS), and the upcoming Sensor Event Service (SES), as well as data models and encodings such as Observations and Measurements (O&M) or the Sensor Model Language (SensorML). Judging from the number of available implementations for the aforementioned services, the amount of available sensor data encoded in O&M and SensorML, as well as the number of SWE-based projects, the initiative has been successful so far. Nevertheless, these services and the stored data are only available as parts of Spatial Data Infrastructures (SDI) and, hence, only available to OGC-compliant client applications. Moreover, the retrieval of sensors and especially observation data and, hence, their re-usability have only been partially addressed so far [3]. Essentially, retrieval still boils down to keyword-based catalogs, code-lists, and ambiguous plain text definitions for observable properties. The missing semantic matching capabilities are among the main obstacles towards a plug & play-like sensor infrastructure [4].

In contrast, the Semantic Web offers a technology stack for information retrieval and reasoning beyond keywords as well as knowledge representation languages for conceptual reference models, i.e., ontologies, used to restrict domain vocabularies towards their intended interpretations [5,6]. The Semantic Sensor Web [7,8] combines the idea of a Web-based access and processing of sensor data from the Sensor Web with the extended knowledge representations, reasoning, and organization capabilities of the Semantic Web. This fusion, for instance, allows to dynamically populate an ontology of weather phenomena, such as hurricanes, based on a set of declarative rules and observations from weather stations stored in a semantically-enabled SOS [9]. Similarly, Keßler et al. [10] demonstrate the interplay of observation data with Semantic Web rule engines to develop a surf-spot recommendation service. While the Semantic Sensor Web addresses the retrieval problem of classical SWE services, the sensor data may still be stored in silo-like databases and, therefore, not directly dereferenceable using HTTP and URIs. Additionally, the interpretation of the data as well as the observed properties and features is determined by the application.

To address these challenges, various researchers [11,12,13,14,15] have proposed to blend the Semantic Sensor Web with Linked Data principles [16]. Sensor data should be identified using URIs, looked up by dereferencing these URIs over HTTP, encoded in machine-readable formats such as RDF, and enriched by links to other data. So far, research on Linked Sensor Data has addressed the problem of defining meaningful URIs [14,12], RESTful access and filters [13,14], the integration within Web mashups [11], and the management of provenance information [17]. Among others, the integration of Semantic Web technologies and Linked Data with classical OGC-based Spatial Data Infrastructures has been recently addressed by Schade and Cox [18], Janowicz et al. [19], as well as Maué and Ortmann [20].

While the presented approaches towards a Semantic Sensor Web and Linked Sensor Data differ in their underlying assumptions, goals, application areas, as well as the used technologies, they all build up on the idea of annotating sensor data using ontologies. Consequently, the development of a sensor ontology has been a major research topic. The developed ontologies range from sensor-centric approaches with a strong relationship to SensorML [21], over rather observation-centric ontologies based on O&M [22,23,24], up to work focusing on stimuli, observed properties, or processes [25,26,27]. A survey of existing ontologies, their differences and commonalities has been recently published by Compton and colleagues [28].

The importance of an ontology for sensors and observations is not restricted to enabling the Semantic Sensor Web but reaches beyond. Gangemi and Presutti point out that with the advent of Linked Data the Semantic Web can become an empirical science in its own rights and argue for knowledge patterns as units of meaning for the Web [29]. Based on the previous argumentation on the role of observations for empirical science, a promising approach for the future may be to ground [30] the Semantic Web in the Sensor Web [31]. For instance, based on work of Quine [32], Gibson [33], and others, Scheider et al. [34,35], have shown how to define geographic categories in terms of observation primitives. In analogy to geodetic datums that define the relation between coordinate systems and the shape of the earth, they proposed to use semantic datums [36] to anchor a semantic space in physical observations.

To address the aforementioned challenges, this paper discusses the development of a Stimulus-Sensor-Observation (SSO) ontology design pattern [37]. While the pattern is intended to act as a generic and reusable component for all kinds of observation-related ontologies for Linked Data and the Semantic Web, it also forms the top-level of the Semantic Sensor Network (SSN) ontology developed by the W3C Semantic Sensor Network Incubator Group (W3C SSN-XG)[1]. The pattern, in turn, is aligned to DOLCE Ultra Light (DUL)[2] as a foundational ontology. The relation between these ontologies is best thought of as layers or modules. The first pattern represents the initial conceptualization as a lightweight, minimalistic, and flexible ontology with a minimum of ontological commitments. While this pattern can already be used as vocabulary for some use cases, other application areas require a more rigid conceptualization to support semantic interoperability. Therefore, we introduce a realization of the pattern based on the classes and relations provided by DOLCE Ultra Light. This ontology can be either directly used, e.g., for Linked Sensor Data, or integrated into more complex ontologies as a common ground for alignment, matching, translation, or interoperability in general. For this purpose, we demonstrate how the pattern was integrated as top-level of the SSN ontology. Note that for this reason, new classes and relations are introduced based on subsumption and equivalence. For instance, the first pattern uses the generic *involves* relation,

---

[1] A W3C incubator group is a 1 – 2 year collaborative activity on a new technology, intended to produce a technical outcome, foster collaboration and interest in the topic, and potentially form the basis of further W3C activities. The wiki of the group is available at `http://www.w3.org/2005/Incubator/ssn/wiki/Main_Page` and contains links to the ontologies, use cases, meeting minutes, and related literature.

[2] see `http://www.loa-cnr.it/ontologies/DUL.owl`.

while the DOLCE-aligned version distinguishes between events and objects and, hence, uses *DUL:includesEvent* and *DUL:includesObject*, respectively.

The remainder of this paper is organized as follows. First, we introduce the core concepts and relations of the Stimulus-Sensor-Observation ontology design pattern. Next, we discuss the alignment of this pattern to the DOLCE foundational ontology and present potential extensions. Afterwards, we discuss how the pattern is integrated into the Semantic Sensor Network ontology. We conclude the paper by summarizing the proposed approach and point towards further work. The introduced ontologies are available online at the wiki of the W3C SSN-XG.

2 The Stimulus-Sensor-Observation Pattern

In the following, we describe the set of classes and relations that jointly form the Stimulus-Sensor-Observation ontology design pattern; see figure 1. The pattern is developed following the principle of minimal ontological commitments to make it reusable for a variety of application areas. A selection of uses cases focused on the retrieval of sensor metadata and observations can be found at: `http://www.w3.org/2005/Incubator/ssn/wiki/index.php?title=Favorite_Use_Cases_List`. The names of classes and relations have been selected based on a review of existing sensor and observation ontologies [28]. Examples are given for each class to demonstrate its application, interaction with other classes, and to ease interpretation.

Stimuli

Stimuli are detectable changes in the environment, i.e., in the physical world. They are the starting point of each measurement as they act as triggers for sensors. Stimuli can either be directly or indirectly related to observable properties and, therefore, to features of interest. They can also be actively produced by a sensor to perform observations. The same types of stimulus can trigger different kinds of sensors and be used to reason about different properties. Nevertheless, a stimulus may only be usable as proxy for a specific region of an observed property. Examples for stimuli include the expansion of liquids or sound waves emitted by a sonar. The expansion of mercury can be used to draw conclusions about the temperature of a surface that is in close contact. While the expansion is unspecific with respect to the kind of surface, e.g., water versus skin, the usage as stimulus is limited by its melting and boiling points. Moreover, mercury is not restricted to thermometers but e.g., also used in nanometers. Note, that the stimulus is the expansion of mercury, not mercury as such.

Sensors

Sensors are physical objects that perform observations, i.e., they transform an incoming stimulus into another, often digital, representation. Sensors are not

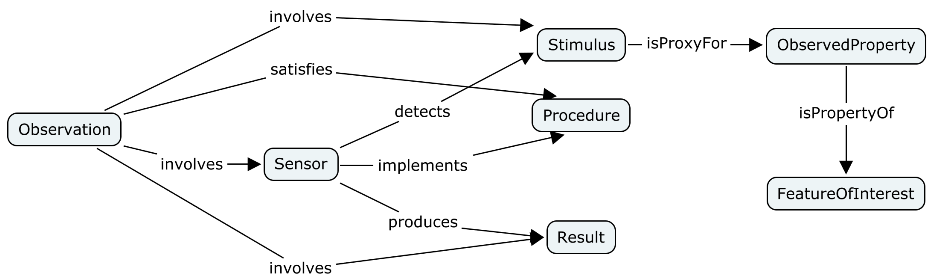

Fig. 1. A conceptual map showing the core concepts and relations (inverse relations are not depicted) forming the Stimulus-Sensor-Observation ontology design pattern. Sensors observe properties of features of interest by detecting stimuli, i.e., changes in the physical world, (directly or indirectly) related to these properties and transforming them to another representation as results. Sensors implement a procedure that describes the transformation of stimuli to results. Observations are the context that bring sensors and stimuli together. They are described by procedures that determine how a certain observation has to be carried out.

restricted to technical devices but may also include humans [2]. A clear distinction needs to be drawn between sensors as objects and the process of sensing. In accordance with the literature on thematic roles [38], we assume that objects are sensors while they perform sensing, i.e., while they are deployed. Furthermore, we also distinguish between the sensor and a procedure, i.e., a description, which defines how a sensor should be realized and deployed to measure a certain observable property. Similarly, to the capabilities of particular stimuli, sensors can only operate in certain conditions. These characteristics are modeled as observable properties of the sensors and includes their survival range or accuracy of measurement under defined external conditions. Finally, sensors can be combined to sensor systems and networks. Examples of sensors range from thermometers and anemometers over telescopes to the human sensory system. Many sensors need to keep track of time and location to produce meaningful results and, hence, are combined with further sensors to sensor systems such as weather stations.

Observations

In the literature, observations are either modeled as database records that store the sensor output and additional metadata or as events in the physical world. We decided to take yet another approach for the design pattern to unify both views. Observations act as the nexus between incoming stimuli, the sensor, and the output of the sensor, i.e., a symbol representing a region in a dimensional space. Therefore, we regard observations as social, not physical, objects. Observations can also fix other parameters such as time and location. These can be specified as parts of an observation procedure. The same sensor can be positioned in different ways and, hence, collect data about different properties. For instance, in contrast to the soil temperature, surface air temperature is measured 2m above ground. Finally, many sensors perform additional processing steps or produce single results based on a series of incoming stimuli. Therefore, observations are rather contexts for the interpretation of the incoming stimuli than physical events.

Observed Properties

Properties are qualities that can be observed via stimuli by a certain type of sensors. They inhere in features of interest and do not exist independently. While this does not imply that they do not exist without observations, our domain is restricted to those observations for which sensors can be implemented based on certain procedures and stimuli. To minimize the amount of ontological commitments related to the existence of entities in the physical world, observed properties are the only connection between stimuli, sensors, and observations on the one hand, and features of interests on the other hand. Examples include, temperature, density, duration, or location and can only be defined with respect to a feature of interest such as a flood.

Features of Interest

Entities are reifications. They are created in an act of cognition and social convention [39,40] and are constructed by relating observable properties [35]. The decision of how to carve out fields of sensory input to form such features is arbitrary to a certain degree and, therefore, has to be fixed by the observation; see above. For instance, soil temperature is measured to determine an optimum time for seeding. The feature of interest is a particular 3-dimensional body of soil restricted by the size of the parcel and a depth relevant for the growth of crop. Without changing the sensor, type of stimulus, and even the observation procedure, the same setting can be used to measure the soil temperature of another feature – e.g., a part of the same parcel to study effects of erosion.

Procedure

Procedures describe how a sensor should be constructed, i.e., how a certain type of stimuli is transformed to a digital representation. Consequently, sensors can be thought of as implementations of procedures where different procedures can be used to derive information about the same type of observed property. For instance, there are many different ways to measure temperature. Procedures can describe so-called contract-based methods, e.g., based on the expansion of mercury, or contact-free methods such as used to construct thermal imaging sensors. Besides procedures for the construction of sensors, procedures also fix observation parameters, e.g., how and where a sensor has to be positioned to measure wind speed and direction. Simplifying, one can think of procedures as cooking recipes.

Result

The result is a symbol representing a value as outcome of the observation. Results can act as stimuli for other sensors and can range from counts and Booleans, to images, or binary data in general. Example includes the number 23 together with degree Celsius as unit of measure.

Extensions to the Pattern

Various extensions to the pattern are possible and can be introduced to fit particular application areas. For instance, work on observation-based symbol grounding requires a semantic datum class. This could be introduces as a subclass of Procedure. Other ontologies developed on top of the pattern may add classes such as Deployment and relate it to the Observation as well as the Sensor class. Similarly, while we discussed how properties of sensors can be specified, we do not introduce such properties to keep the pattern flexible and generic. More details on the integration of sensor-centric classes are described in section 4.

3 Alignment to DOLCE

The presented design pattern is intended as a minimal vocabulary for the Semantic Sensor Web and Linked Data in general. While it introduces the key classes and their relations such as stimuli, sensors, and observations, the terms used for their specification remain undefined. For instance, the pattern does not explain what processes, qualities, or social objects are. The advantage of such an underspecified ontology pattern is its flexibility and the possibility to integrate and reuse it in various applications by introducing subclasses. The downside, however, is a lack of explicit ontological commitments and, therefore, reduced semantic interoperability. Two Linked Data providers can annotate their data using the SSO pattern and still have radically different conceptualizations with respect to the nature of the involved classes. For instance, one provider could consider observations to be events in the real world, while another provider may think of them as database records.

While ontologies cannot fix meaning, their task is to restrict the interpretation of the introduced classes and relations towards the intended model [5,6]. To assist knowledge engineers and users in interpreting the SSO pattern, we align it to the DOLCE Ultra Light foundational ontology. For instance, while the SSO ontology design pattern draws on the functional approach presented by Kuhn [25] their differences only become visible based on their incompatible alignments to DOLCE. Kuhn, for example, defines observations as perdurants, while the pattern defines them as social objects. Probst [22], in contrast, does not introduce the notion of a stimulus but starts at the instrument level. Moreover, he distinguishes between the quale as continuous value of the observed property and the discrete result as its measured approximation. We explicitly avoid going beyond the stimulus (as proxy for the observed property) within the pattern.

In the following we highlight the major aspects of the DOLCE alignment; see figure 2. Each SSO class is defined as subclasses of an existing DOLCE class and related to other SSO and DUL classes using DUL properties. New types of relations are only introduced when the domain or range have to be changed, in all other cases the relations from DUL are reused. The resulting extension to DUL preserves all ontological commitments defined in section 2.

Stimuli

The SSO class Stimulus can be either defined as a subclass of DUL:Event or its immediate subclasses Action and Process. In contrast to processes, actions require at least one agent as participant and, therefore, would be too restrictive for the design pattern. The classifications of events in DUL is work in progress. For instance, there is nothing said about how processes differ from other kinds of events. Therefore, the pattern defines a Stimulus as a subclass of DUL:Event. As a consequence, stimuli need at least one DUL:Object as participant. Such objects include the mercury participating in the event of expansion or a thermometer involved in the detection of a stimulus.

Sensors

Sensors are defined as subsclasses of physical objects (DUL:PhysicalObject). Therefore, they have to participate in at least one DUL:Event such as their deployment. This is comparable to the ontological distinction between a human and a human's life. Sensors are related to procedures and observations using the DUL:isDescribedBy and DUL:isObjectIncludedIn relations, respectively.

Observations

The class Observation is specified as a subclass of DUL:Situation, which in turn is a subclass of DUL:SocialObject. The required relation to stimuli, sensors, and results can be modeled using the DUL:includesEvent, DUL:includesObject, and DUL:isSetting relationship, respectively. Observation procedures can be integrated by DUL:satisfies. The decision to model observations as situations is also conform with the observation pattern developed by Blomqvist[3].

Observed Properties

ObservedProperty is defined as a subclass of DUL:Quality. Types of properties such as temperature or pressure should be added as subclasses of ObservedProperty instead of individuals. A new relation called SSO:isPropoertyOf is defined as a sub-relation of DUL:isQualityOf to relate a property to a feature of interest.

[3] see `http://ontologydesignpatterns.org/wiki/Submissions:Observation`.

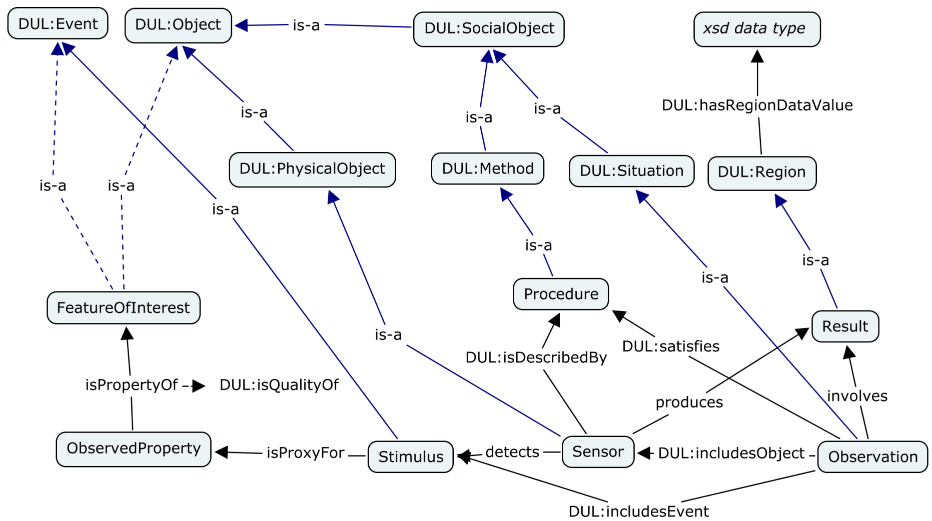

Fig. 2. A conceptual map depicting the core parts of the alignment of the SSO ontology design pattern with the DOLCE Ultra Light foundational ontology.

Features of Interest

Features of interest can be events or objects. We deliberately exclude the DUL classes Quality and Abstract as potential features of interest to avoid complex philosophical questions such as whether there are qualities of qualities. Instead we argue that the need to introduce properties for qualities is an artifact of reification or confuses qualities with features or observations. For instance, accuracy is not a property of a temperature but the property of a sensor or an observation procedure.

Procedure

Procedure is defined as a subclass of DUL:Method which in turn is a subclasses of DUL:Description. Consequently, procedures are expressed by some DUL:InformationObject such as a manual or scientific paper.

Result

The SSO Result class is modeled as a subclass of DUL:Region. A concrete data value is introduced using the data property DUL:hasRegionDataValue in conjunction with some xsd data type.

4 Integrating the Pattern with the SSN Ontology

The SSO pattern and alignment as presented thus far gives a basis for understanding and describing the relationship between sensors and their observations.

However, such a pattern is too high level to be immediately useful in many of the domains mentioned in the introduction. It is once the pattern is set in a broader context of sensors, their capabilities, operating conditions, and all the metadata that sets them in their context that the pattern realizes its potential as a language for sensors and observations.

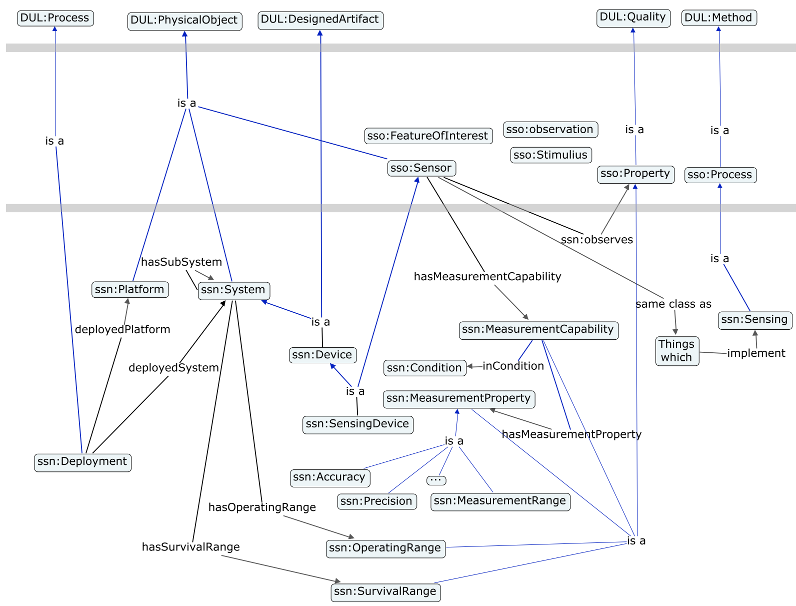

Fig. 3. A partial view on the integration of the DUL-aligned ontology design pattern with the Semantic Sensor Network ontology developed by the W3C SSN-XG.

The SSN ontology provides such a broader perspective. It is intended to cover sensing concepts and little else that would be present in multiple domains or that is not a key notion for describing sensors. The ontology supports different views: (1) it can be centered around sensors, their capabilities and constraints, the procedures they execute and the observations they can make; (2) around observations, i.e., what was observed and how, as well as on (3) features and properties and how to observe them. The SSN ontology builds directly on the SSO classes introduced before and adds additional relations and classes (e.g., Sensing) for measurement capabilities, operating conditions, and deployments; see figure 3. In some cases, the DUL relations and classes used in the pattern alignment are too generic for an intuitive understanding in the domain of sensing and sensors. In this cases, the SSN ontology introduces sub-relations or classes with different names, domains, and ranges. For example, SSN:observedBy (as a sub-relation of DUL:includesObject) is used to relate SSN:Observation to SSN:Sensor.

Sensors

Sensors, in the SSN ontology, as well as making observations by implementing sensing procedures, have capabilities (SSN:MeasurmentCapability) that describe the accuracy, range, resolution and other properties (e.g., SSN:Accuracy and SSN:Resolution) of the sensor. As these are observable characteristics of a sensor, they are modeled as properties (SSO:Property). Similarly, the operating and survival conditions (SSN:OperatingCondition and SSN:SurvivalCondition) are properties of a sensor. Indeed accuracy, battery levels, or further properties can be observed by other sensors. Therefore, one can model a sensor, or parts of it, as a feature of interest observed by another sensor.

Observations

The observation model of the SSN ontology, is similar to the O&M model, with observations of properties of particular features of interest being made at particular times and yielding a result of a value for the property. The SSN model describes the sensor (similar to procedure in O&M) that made the observation, thus linking the observation-centric and sensor-centric view points. The model allows for further refinements describing how the sensor was used. The main difference between the O&M and SSN-XG models is the distinction of classifying them as events and social objects, respectively. While the distinction was discussed above, both views are compatible from a data-centric perspective, i.e., they allow to assign values for observed properties by sensors.

Stimuli, Features of Interest, Observed Properties, and Deployments

Stimuli, features of interest, observed properties, and the modeling of observation results are taken directly from the SSO pattern and do not need to be subclassed in the SSN ontology. Adding types of properties, for instance, is out of scope for the W3C SSN-XG but up to ontologies such as SWEET.[4] Features of interest can be added as subclasses and imported from domain ontologies such as the Ordnance Survey hydrology ontology.[5]

The remaining parts of the SSN-XG ontology are various physical objects such as platforms (SSN:Platform), systems (SSN:Systems), devices (SSN:Device), and the deployment (SSN:Deployment) process that all serve to further define sensors and place them in the context of other systems. Deployments represent the ongoing processes related to deploying a sensor in the field (and are thus subclasses of DUL:Process). A deployment relates a sensor, a platform to which it was deployed, and a time frame, and may involve parts such as installation and maintenance.

[4] see `http://sweet.jpl.nasa.gov/ontology/`.

[5] see `http://www.ordnancesurvey.co.uk/oswebsite/ontology/Hydrology/v2.0/Hydrology.owl`.

5 Conclusions and Outlook

The presented work introduces a generic Stimulus-Sensor-Observation ontology design pattern intended as a building block for work on the Semantic Sensor Web and Linked Sensor Data. The key classes and relations are discussed in detail and illustrated by examples. To support semantic interoperability and restrict the interpretation of the introduced classes towards their intended meaning, the design pattern is aligned with the DOLCE foundational ontology. The integration of the pattern (as well as DUL) into the Semantic Sensor Network ontology is discussed.

By providing a richer axiomatization, the DUL alignment should also improve semantic search, pattern mining, similarity and analogy-based reasoning, and other services. At the same time, the added complexity makes using and populating the ontology more difficult. Open-world reasoning ensures that many required parts of the ontology need not to be filled. For instance, each individual sensor of the aligned SSN ontology requires an event to participate in, e.g., a deployment, as well as a stimulus, and, similarly, each feature of interest either requires the instantiation of a participating object or an event in which the feature participates; however, these need not actually be instantiated as they are implied. Nevertheless, the DOLCE alignment may still be over-engineered for some applications and especially Linked Data and, therefore, restrict the usage of the ontologies. For this reason, the W3C SSN-XG intends to deliver the ontologies together with a script to free them from the DUL alignment. To ensure that first-time users are still aware of the ontological commitments, the ontologies are stored with the alignment by default.

Finally, the presented ontologies and the alignment are still work in progress and will be part of the final report of the W3C SSN-XG. Among other aspects, the relation between sensors and results needs further work. For instance, DUL:Region may be replaced by DUL:Parameter in the future. Further work will also focus on documentations and use cases to demonstrate how to integrate the ontologies or develop extensions. These use cases and additional user feedback will be used for further refinement. Developing modular ontologies at different abstraction levels and based on common agreement raises many other challenges for collaborative ontology engineering and keeping these ontologies in sync.

Acknowledgments

The presented work is part of the W3C Semantic Sensor Network Incubator Group and the 52°*North* semantics community. We are thankful to our colleagues from both initiatives, as well as to Werner Kuhn and Simon Scheider for their input to improve the quality of the paper and the developed ontologies. This work was in part supported by CSIRO's Water for a Healthy Country Flagship.

References

1. Quine, W.V.O.: From Stimulus to Science. Harvard University Press (1995)
2. Goodchild, M.: Citizens as Sensors: the World of Volunteered Geography. GeoJournal **69**(4) (2007) 211–221
3. Jirka, S., Bröring, A., Stasch, C.: Discovery mechanisms for the sensor web. Sensors **9**(4) (2009) 2661–2681
4. Bröring, A., Janowicz, K., Stasch, C., Kuhn, W.: Semantic challenges for sensor plug and play. In Carswell, J., Fotheringham, S., McArdle, G., eds.: Web & Wireless Geographical Information Systems (W2GIS 2009), 7 & 8 December 2009, Maynooth, Ireland. Number 5886 in LNCS, Springer (2009) 72–86
5. Guarino, N.: Formal Ontology and Information Systems. In Guarino, N., ed.: International Conference on Formal Ontology in Information Systems (FOIS1998). IOS Press, Trento, Italy (1998) 3–15
6. Kuhn, W.: Semantic Engineering. In Navratil, G., ed.: Research Trends in Geographic Information Science, Springer (2009; forthcoming)
7. Sheth, A., Henson, C., Sahoo, S.: Semantic Sensor Web. IEEE Internet Computing (2008) 78–83
8. Corcho, O., Garcia-Castro, R.: Five challenges for the semantic sensor web. Semantic Web - Interoperability, Usability, Applicability (2010; to appear)
9. Henson, C.A., Pschorr, J.K., Sheth, A.P., Thirunarayan, K.: SemSOS: Semantic sensor observation service. In: International Symposium on Collaborative Technologies and Systems (CTS 2009). (2009)
10. Keßler, C., Raubal, M., Wosniok, C.: Semantic rules for context-aware geographical information retrieval. In Barnaghi, P., ed.: European Conference on Smart Sensing and Context, (EuroSSC 2009). Volume 5741 of LNCS., Springer (2009) 77–92
11. Phuoc, D.L., Hauswirth, M.: Linked open data in sensor data mashups. In Kerry Taylor, Arun Ayyagari, D.D.R., ed.: Proceedings of the 2nd International Workshop on Semantic Sensor Networks (SSN09). Volume Vol-522., CEUR (2009) 1–16
12. Sequeda, J., Corcho, O.: Linked stream data: A position paper. In K., T., Ayyagari, A., De Roure, D., eds.: Proceedings of the 2nd International Workshop on Semantic Sensor Networks (SSN09). Volume Vol-522., CEUR-WS (2009) 148–157
13. Page, K., De Roure, D., Martinez, K., Sadler, J., Kit, O.: Linked sensor data: Restfully serving rdf and gml. In K., T., Ayyagari, A., De Roure, D., eds.: Proceedings of the 2nd International Workshop on Semantic Sensor Networks (SSN09). Volume Vol-522., CEUR (2009) 49–63
14. Janowicz, K., Bröring, A., Stasch, C., Everding, E.: Towards meaningful uris for linked sensor data. In Devaraju, A., Llaves, A., Maué, P., Keßler, C., eds.: Towards Digital Earth: Search, Discover and Share Geospatial Data. Workshop at Future Internet Symposium. (2010; forthcoming)
15. Patni, H., Henson, C., Sheth, A.: Linked sensor data. In: 2010 International Symposium on Collaborative Technologies and Systems, IEEE (2010) 362–370
16. Bizer, C., Heath, T., Berners-Lee, T.: Linked data – the story so far. Journal on Semantic Web and Information Systems (IJSWIS); Special Issue on Linked Data **5**(3) (2009) 1–22
17. Patni, P., Sahoo, S., Henson, C., Sheth, A.: Provenance aware linked sensor data. In Kärger, P., Olmedilla, D., Passant, A., Polleres, A., eds.: Proceedings of the Second Workshop on Trust and Privacy on the Social and Semantic Web. (2010)

18. Schade, S., Cox, S.: Linked data in sdi or how gml is not about trees. In: Proceedings of the 13th AGILE International Conference on Geographic Information Science - Geospatial Thinking. (2010)
19. Janowicz, K., Schade, S., Bröring, A., Keßler, C., Maue, P., Stasch, C.: Semantic enablement for spatial data infrastructures. Transactions in GIS **14**(2) (2010) 111 – 129
20. Maué, P., Ortmann, J.: Getting across information communities: Embedding semantics in the sdi for the amazon. Earth Science Informatics **2**(4) (2009) 217 – 233
21. Russomanno, D., Kothari, C., O. Thomas, O.: Building a sensor ontology: A practical approach leveraging iso and ogc models. In: The 2005 International Conference on Artificial Intelligence (IC-AI 2005). (2005) 637–643
22. Probst, F.: Ontological Analysis of Observations and Measurements. In Raubal, M., Miller, H., Frank, A., Goodchild, M., eds.: Geographic Information Science, 4th International Conference, GIScience 2006. Number 4197 in LNCS, Springer (2006) 304–320
23. OOSTethys: http://mmisw.org/ont/mmi/20090519t125341/general ; last accessed 09/2010.
24. Neuhaus, H., Compton, M.: The Semantic Sensor Network Ontology: A Generic Language to Describe Sensor Assets. In: AGILE 2009 Pre-Conference Workshop Challenges in Geospatial Data Harmonisation, 02 June 2009, Hannover, Germany. (2009)
25. Kuhn, W.: A functional ontology of observation and measurement. In Janowicz, K., Raubal, M., Levashkin, S., eds.: GeoSpatial Semantics. Volume LNCS5892., Springer-Verlag Berlin Heidelberg (2009) 26–43
26. Stasch, C., Janowicz, K., Broering, A., Reis, I., Kuhn, W.: A Stimulus-Centric Algebraic Approach to Sensors and Observations. In: 3rd International Conference on Geosensor Networks. Number 5659 in LNCS, Springer (2009) 169–179
27. Devaraju, A., Neuhaus, H., Janowicz, K., Compton, M.: Combining Process and Sensor Ontologies to Support Geo-Sensor Data Retrieval. In: 6th International Conference on Geographic Information Science (GIScience 2010), Zurich, CH 14-17th September, 2010. (2010; forthcoming)
28. Compton, M., Henson, C., Neuhaus, H., Lefort, L., Sheth, A.: A survey of the semantic specification of sensors. In K., T., Ayyagari, A., De Roure, D., eds.: Proceedings of the 2nd International Workshop on Semantic Sensor Networks (SSN09). Volume Vol-522., CEUR (2009) 17–32
29. Gangemi, A., Presutti, V.: Towards a pattern science for the semantic web. Semantic Web - Interoperability, Usability, Applicability (2010; to appear)
30. Harnad, S.: The symbol grounding problem. Phys. D **42**(1–3) (1990) 335–346
31. Janowicz, K.: The role of space and time for knowledge organization on the semantic web. Semantic Web - Interoperability, Usability, Applicability (2010; to appear)
32. Quine, W.V.O.: Pursuit of Truth; revised edition. Cambridge, MA: Harvard University Press. (1992)
33. Gibson, J.: The theory of affordances. In Shaw, R., Bransford, J., eds.: Perceiving, Acting, and Knowing - Toward an Ecological Psychology. Lawrence Erlbaum Ass., Hillsdale, New Jersey (1977) 67–82
34. Scheider, S., Janowicz, K., Kuhn, W.: Grounding geographic categories in the meaningful environment. In Hornsby, K., Claramunt, C., Denis, M., Ligozat, G., eds.: Conference on Spatial Information Theory (COSIT 2009). Volume 5756 of LNCS., Springer (2009) 69–87

35. Scheider, S., Probst, F., Janowicz, K.: Constructing bodies and their qualities from observations. In: Proceeding of the 2010 Conference on Formal Ontology in Information Systems (FOIS 2010), IOS Press (2010) 131–144
36. Kuhn, W.: Semantic reference systems (guest editorial). International Journal of Geographical Information Science **17**(5) (2003) 405–409
37. Gangemi, A.: Ontology design patterns for semantic web content. In Gil, Y., Motta, E., Benjamins, V., Musen, M., eds.: The Semantic Web – ISWC 2005. Volume 3729 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg (2005) 262–276
38. Sowa, J.F.: Knowledge Representation: Logical, Philosophical and Computational Foundations. Brooks Cole Publishing Co., Pacific Grove, CA, USA (2000)
39. Mark, D.M.: Toward a theoretical framework for geographic entity types. In: Spatial Information Theory: A Theoretical Basis for GIS, International Conference COSIT '93. (1993) 270–283
40. Lehar, S.: The World in Your Head: A Gestalt View of the Mechanism of Conscious Experience. Lawrence Erlbaum (2003)