

The Semantic Web Needs More Cognition

Editors: Krzysztof Janowicz, Pennsylvania State University, U.S.A. and Pascal Hitzler, Wright State University, U.S.A.
Solicited reviews: Philipp Cimiano, Universität Bielefeld, Germany and Aldo Gangemi, ISTC-CNR Rome, Italy

Martin Raubal^{a*} and Benjamin Adams^b

^a*Department of Geography, University of California, Santa Barbara, 5713 Ellison Hall, Santa Barbara, CA 93106-4060, U.S.A.*

^b*Department of Computer Science, University of California, Santa Barbara, Engineering I, Room 2104, Santa Barbara, CA 93106-5110, U.S.A.*

Abstract. One of the key deficiencies of the Semantic Web is its lack of cognitive plausibility. We argue that by accounting for people’s reasoning mechanisms and cognitive representations, the usefulness of information coming from the Semantic Web will be enhanced. More specifically, the utilization and integration of conceptual spaces is proposed as a knowledge representation that affords two important human cognitive mechanisms, i.e., semantic similarity and concept combination. Formal conceptual space algebra serves as the basis for the Conceptual Space Markup Language (CSML), which facilitates the engineering of ontologies using a geometric framework. We demonstrate the usefulness of the approach through a concrete example and suggest directions for future work, especially the need for combining geometric representations and reasoning mechanisms with existing Semantic Web structures.

Keywords: Conceptual space, Conceptual Space Markup Language, CSML, cognition, representation, markup language

1. Introduction

In 2004, Peter Gärdenfors argued that “the Semantic Web is not semantic” because it is good for syllogistic reasoning only and there is more to semantics than that [1]. In 2010, we claim here that the Semantic Web is still not semantic in the human sense because it does not sufficiently account for people’s cognition, i.e., human conceptual representations and reasoning mechanisms. This must not be confused with a search for *Strong Artificial Intelligence*, i.e., a Semantic Web whose intellectual ability cannot be distinguished from that of a human being [2]. But eventually, what comes out of the Semantic Web should be useful for people and it is our conviction that the better we integrate and account for people’s reasoning mechanisms and cognitive representations the more useful such information will be.

Consider the example of looking for a warm climate vacation (Figure 1). This search involves sever-

al questions that cannot be handled by the current Semantic Web, such as what is the meaning of ‘warm’ in a particular person’s context of climate and how important is this dimension compared to other dimensions, such as distance and cost? This example makes it clear that the Semantic Web has to be based on a solid foundation of human concept processing, including limited knowledge and uncertainty in order to become truly semantic. In addition, representation and processing of context information, is key. Semantic models of context and contextualizing ontologies must account for human sensors and move into the direction of dynamic processes [3, 20].

More specifically, we argue that knowledge representations underpinning the Semantic Web should afford two important human cognitive tasks: the *efficient calculation of semantic similarity* (in the context of the vacation example: how similar is the result to my ideal warm climate vacation?) and *combinations of concepts* (‘warm’ and ‘climate’). However,

* Corresponding author. E-mail: raubal@geog.ucsb.edu.

the existing logical foundations of the Semantic Web—description logics and rules—presume a set-based classification scheme that does a poor job of facilitating these operations. By adopting a geometric and topological representational framework called *conceptual spaces* to describe semantics at the conceptual level, these operations can be defined in terms of an efficient vector algebra. This opens up the possibility to go beyond the classical concept combination possibilities of conjunction, disjunction, and negation. Conceptual spaces were conceived as a theory for how concepts are learned based on the paradigm of cognitive semantics [4], which emphasizes the role of similarity and prototype effects in categorization [5], and the importance of metaphorical and metonymic reasoning. Combined with natural language processing and existing methods of sentential representation, geometric conceptual representation has the potential to create a much richer and cognitively plausible Semantic Web [21].

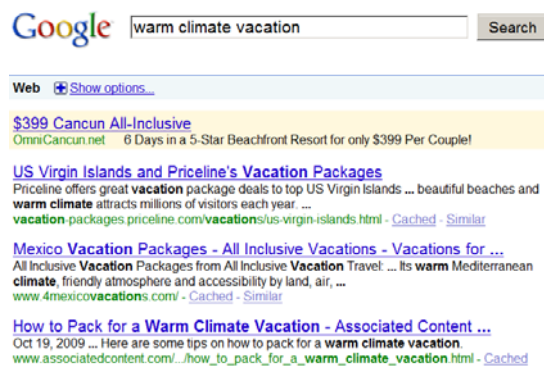


Fig. 1. Search for warm climate vacation.

2. Conceptual Space Algebra

Conceptual spaces were introduced to represent information at the conceptual level [6]. They can be utilized for knowledge representation and sharing, and account for the fact that concepts are dynamic and change over time [7, 8]. A conceptual space is a set of quality dimensions with a geometric / topological structure for one or more domains. Domains are represented by sets of *integral* dimensions, which are separable from all other dimensions. Concepts cover multiple domains and are modeled as n-dimensional regions. Every instance of a category can be represented as a vector in the conceptual space [9]. This allows for expressing the similarity between two in-

stances as a function of the spatial distance between their vectors. The utilization of conceptual space theory within the Semantic Web requires a solid mathematical foundation. Adams and Raubal [10] presented a metric conceptual space algebra, which consists of formal definitions of its components and operations that can be applied to them. Conceptual spaces are defined as multi-leveled structures and a distinction is made between the representation of the *geometric* elements (regions, points) and the *conceptual* elements (concepts, properties, instances). Furthermore, contrast classes—special types of properties, which have meanings that are dependent on the concepts they modify—are specified. Context is defined as a set of salience weights that can be applied to components of any type in the conceptual space, and is therefore a first-order element of a conceptual space. Different algebraic operations, such as metric operations on points and regions, and context-dependent similarity and concept combination query operations can then be applied to the elements of a conceptual space.

In order to facilitate the engineering of ontologies [22] using a geometric framework, languages must be developed to describe the geometric structures. The Conceptual Space Markup Language (CSML) is based on the described algebra and designed for this purpose.

3. Conceptual Space Markup Language

CSML [11] is an XML-based language that allows one to create an ontology of concepts, properties, instances, contrast classes, and contexts as defined in the algebra above. The following shows the *climate* domain described in CSML with two dimensions temperature and precipitation.

```
<csml:Domain csml:ID="Climate">
  <csml:QualityDimension
    csml:ID="Temperature">
    <csml:Scale> interval </csml:Scale>
  </csml:QualityDimension>
  <csml:QualityDimension
    csml:ID="Precipitation">
    <csml:Scale> ratio </csml:Scale>
  </csml:QualityDimension>
</csml:Domain>
```

Different climate properties (e.g., wet, dry, hot, cold, Californian, temperate) are represented as regions within the climate domain. In CSML properties and contrast classes are described as systems of linear inequalities expressed using a variant of MathML.

The following shows how one can represent *warm* as a contrast class in CSML.

```

<csml:ContrastClass csml:ID="Warm"
csml:DomainID="Climate">
  <csml:aVector>
    <cn> 5.4 </cn>
    <cn> 1.0 </cn>
  </csml:aVector>
  <csml:qVector>
    <ci> Temperature </ci>
    <ci> Precipitation </ci>
  </csml:qVector>
  <csml:ccMin> -3.0 </csml:ccMin>
  <csml:ccMax> 4.6 </csml:ccMax>
</csml:ContrastClass>

```

As well, different classes of locations have different climate properties (e.g., California climate) represented as regions in the climate domain bounded by, for example, minimum and maximum average temperatures and precipitation measures.




Fig. 2. Warm contrast class.

For the scenario where the user wants to search for a warm climate vacation, it is straightforward to represent the requisite elements in CSML in a manner that affords semantic search based on context. Specifically, one can frame the goal of this semantic search query by identifying the concepts that are most semantically similar to the user's idealized or prototypical warm vacation location depending on the user's location. Here there are really two different kinds of context at play. First, there is context represented as salience weights on the dimensions for the purpose of similarity measurement. In the example, precipitation might be weighted as highly as temperature because for two locations to have the same climate both precipitation and temperature matter equally. This first kind of context is described using the

`csml:Context` tag in CSML. Second, there is context in terms of which climate property *warm* should modify.

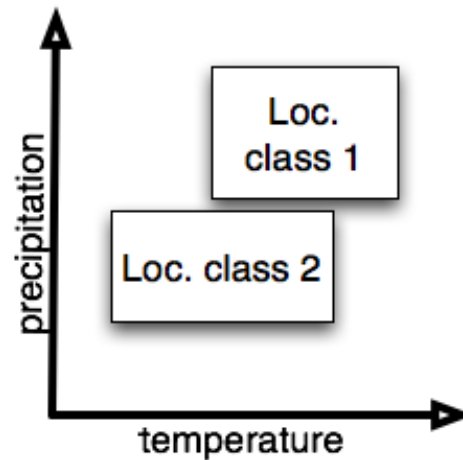


Fig. 3. Climate properties.

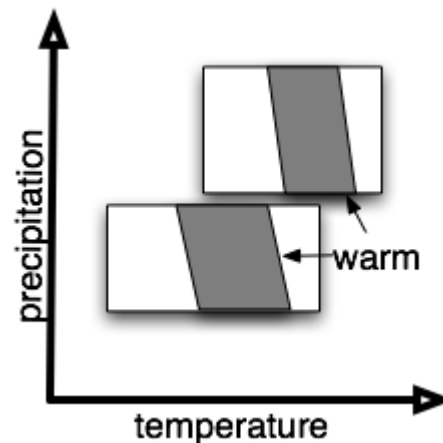


Fig. 4. Combination of warm with climate properties (dark areas).

The concept of a *warm German vacation* entails different semantics for the term *warm* than does *warm California vacation* (or for that matter *warm coffee*, which is actually cold!). As a contrast class, *warm* is represented in a conceptual space as a sub-region of the entire climate domain (Figure 2). The combination of *warm* with another class is not the intersection but rather a geometric affine projection of the *warm* region onto the other class' climate property (Figures 3 and 4). And unlike union and intersection this class constructor is asymmetric.

With this kind of concept combination operation one can very easily reason non-monotonically that what is warm in Sweden is not warm in Europe even if European country is modeled as a generalization or super-class of Sweden. Further, the geometric representation allows one to represent classes in terms of prototypical instances, i.e., as vectors or regions in a conceptual space [12]. This prototype representation is far more natural for representing classes without clear necessary and sufficient features (i.e., classes with degrees of membership determined by similarity to a prototype ideal, such as classes of shapes).

We should note that a semantic search query such as the one above does require a system that can identify which terms are modifying other terms. However, this problem is true for description logics as well and illustrates the need for a natural language processing [13] layer for the Semantic Web.

4. Where To Go From Here

Since it requires identifying the measurable dimensions of a property, the geometric representation might on the face of it seem overly restrictive. However, there is ample evidence that spatial metaphors are used in conceptualizations for many domains of knowledge [14], including any ordinal, interval, and ratio scaled measurements of observable phenomena. In addition, this representation does not necessarily require that the dimensions be identified in the cases when they are modeled as latent variables using techniques such as multidimensional scaling. Further, from the ontology engineer's perspective it makes little sense in many cases to translate the semantics of metric, spatially ordered data into a description logic representation, because 1) it adds unnecessary complexity, since transitivity, disjointness, and other logical characteristics emerge directly from the order topology of the space and 2) it affords the use of linear algebra and computational geometry algorithms as the foundations for many reasoning operations, which can be much more efficient. From a cognitive perspective the latter point aligns with the argument that much similarity measurement happens at the perceptual level without the need for higher-order cognitive representation [15]. Nevertheless, an important future development will be the formalization of mappings between conceptual space representations and OWL based ontologies. This includes the representation of vague information in Fuzzy OWL [16] and compar-

ing the semantic expressiveness between conceptual spaces and Fuzzy OWL.

Description logics and conceptual spaces are two different knowledge representation frameworks with different degrees of semantic expressiveness and thus mappings between the two can result in a loss of information. Generally speaking, numeric datatype properties and object properties can be mapped to dimensions and regions in domains, respectively, but there are exceptions to this rule. In most cases the taxonomic relationships in a conceptual space representation can be 'frozen' into a description logic-based representation, but in doing so it loses expressiveness. For example, categories that are defined using prototypes will entail different memberships depending on context (i.e., dimension weights, which may be set by the user or automatically be assigned through learning from user behavior), so a conceptual space can generate a number of different OWL ontologies dependent on context. In addition, the notion of membership in a category existing on a continuum based on similarity is lost. The representation of regions as sets of linear inequalities might be achievable with the proposed OWL 2 Linear Equations data range extension, but arguably in a very cumbersome manner¹. Since it is likely that ontology engineers will want to retain their ability to use all the existing features of OWL, a hybrid (or dual) representation will be in order. Such a hybrid representation would give ontology engineers the flexibility to define classes based on necessary and sufficient features or prototypes and use set based class constructors as well as more cognitively plausible methods based on contrast classes.

In the semantic web layer cake we conceive of the CSML layer as being a layer that sits on top of XML and beside rules and OWL. CSML can be an earlier stage in the pipeline for building an OWL representation, though it also has a role within the reasoning pipeline, e.g., when doing similarity measurement. Furthermore, reasoning can be done on conceptual spaces without mapping to OWL and this reasoning can exploit the characteristics of the geometric representation as a foundation for more complex sorts of class constructors. The following steps illustrate an example of how the CSML layer can be used to map a set of classes represented by prototypes in a feature space to an OWL ontology (see also Figure 5).

¹<http://www.w3.org/TR/2009/WD-owl2-dr-linear-20090421/>

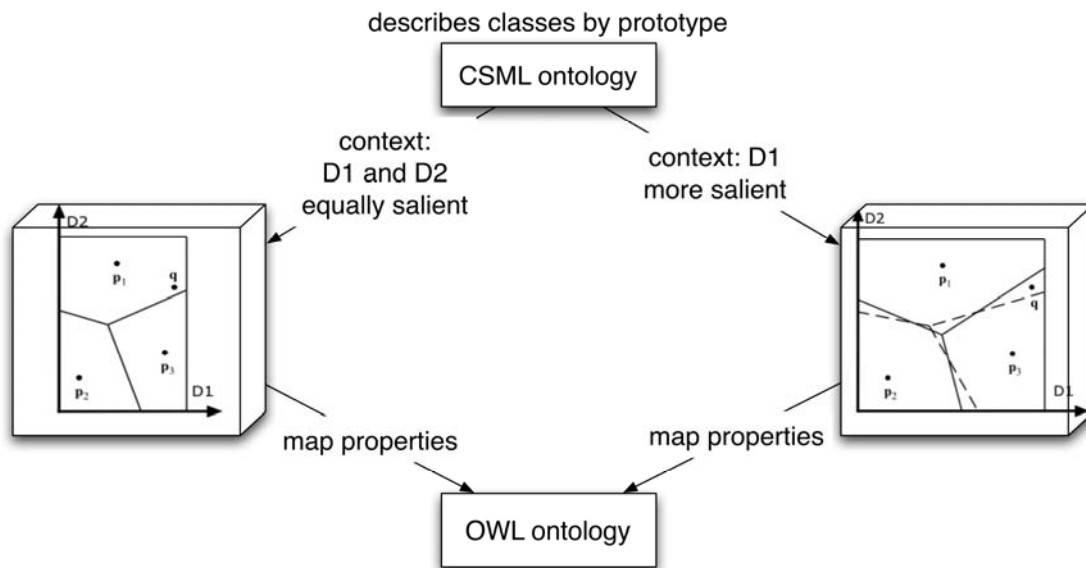


Fig. 5. Example of context changing classification.

1. A machine learning algorithm is used to learn points of central tendency for classes of observations. These points are interpreted as prototypical instances of the classes. Note that most machine learning algorithms of this sort are flat, so in conceptual space parlance it is in fact properties that are being learned not classes (i.e., they are in one domain).
2. The dimensions and points are represented as quality dimensions and instances in CSML.
3. If the classes are disjoint then the Voronoi tessellation of the space based on these prototypes can be used to identify regions representing different properties. By placing different saliency weights on the dimensions the Voronoi tessellation may classify a particular observation differently [17].
4. A mapping is made from CSML to OWL for a given context.

A mapping from a CSML ontology to an OWL ontology is a morphism that reifies the quality dimensions, properties, instances, and concepts as OWL classes, properties, and individuals (see also [18]). In a metric conceptual space quality dimensions are mapped to transitive datatype properties, properties to object properties, instances to individuals, and con-

cepts to classes. Alternately, domains can be mapped to object properties and conceptual space properties to individuals. A logical formalization of this mapping as a function that takes two input parameters, a conceptual space knowledge base and a set of context weights, and outputs a SROIQ(D) description logic knowledge base is a current research objective. Note in particular, that by mapping conceptual space properties to object properties information about the geometric structure is lost, therefore maintaining a link to the CSML representation that generated the OWL ontology can be used for finer-grained similarity reasoning. In addition, conceptual spaces are well-suited for non-monotonic changes based on new observations, e.g., adding new points to the original space can change the points of central tendency and the resulting property regions, which can be mapped to an updated OWL ontology.

This leads directly to the question of how we will be able to generate cognitively plausible ontologies in the future not only from measurement data but from the mass of user-generated data such as Volunteered Geographic Information (VGI) [19]. It will also be necessary to contextualize these ontologies on the fly (see the warm vacation example used here). With conceptual space representations this may be done by putting weights on the dimensions and modifying the

classifications. If we know about people's prototypical concepts for different domains, how can we construct ontologies from there?

The future will show whether what is out there can be integrated with conceptual space theory and whether such combination and integration of ideas will eventually pave the way to a truly cognitively plausible Semantic Web, a Semantic Web that is useful for its users.

Acknowledgments

This work is supported by a UCSB faculty research grant. The two reviewers' suggestions helped improve the content of the paper.

References

1. Gärdenfors, P., *How to Make the Semantic Web More Semantic*, in *Formal Ontology In Information Systems*, A. Varzi and L. Lieu, Editors. 2004, IOS Press: Amsterdam. p. 17-34.
2. Searle, J., *Minds, brains, and programs*. Behavioral and Brain Sciences, 1980. 3(3): p. 417-457.
3. Keßler, C., M. Raubal, and C. Wosniok, *Semantic Rules for Context-Aware Geographical Information Retrieval*, in *Smart Sensing and Context - 4th European Conference, EuroSSC 2009, Guildford, UK, September 16-18, 2009*, P. Barnaghi, et al., Editors. 2009, Springer: Berlin, Heidelberg. p. 77-92.
4. Janowicz, K., M. Raubal, A. Schwering, and W. Kuhn, *Semantic Similarity Measurement and Geospatial Applications (Editorial)*. Transactions in GIS, 2008. 12(6): p. 651-659.
5. Rosch, E., *Principles of Categorization*, in *Cognition and Categorization*, E. Rosch and B. Lloyd, Editors. 1978, Lawrence Erlbaum Associates: Hillsdale, New Jersey. p. 27-48.
6. Gärdenfors, P., *Conceptual Spaces - The Geometry of Thought*. 2000, Cambridge, MA: Bradford Books, MIT Press.
7. Barsalou, L., W. Yeh, B. Luka, K. Olseth, K. Mix, and L. Wu, *Concepts and meaning*, in *Papers from the parasession on conceptual representations*, K. Beals, et al., Editors. 1993, University of Chicago: Chicago Linguistics Society. p. 23-61.
8. Raubal, M., *Representing concepts in time*, in *Spatial Cognition VI - Learning, Reasoning, and Talking about Space. Proceedings of the International Conference Spatial Cognition 2008, Freiburg, Germany*, C. Freksa, et al., Editors. 2008, Springer: Berlin. p. 328-343.
9. Raubal, M., *Formalizing Conceptual Spaces*, in *Formal Ontology in Information Systems, Proceedings of the Third International Conference (FOIS 2004)*, A. Varzi and L. Vieu, Editors. 2004, IOS Press: Amsterdam, NL. p. 153-164.
10. Adams, B. and M. Raubal, *A Metric Conceptual Space Algebra*, in *Spatial Information Theory - 9th International Conference, COSIT 2009, Aber Wrac'h, France, September 2009*, K. Stewart Hornsby, et al., Editors. 2009, Springer: Berlin. p. 51-68.
11. Adams, B. and M. Raubal, *The Conceptual Space Markup Language (CSML): Towards the Cognitive Semantic Web*, in *Third IEEE International Conference on Semantic Computing (ICSC 2009), 14-16 September, Berkeley, California*. 2009, IEEE Computer Society. p. 253-260.
12. Schwering, A. and M. Raubal, *Measuring Semantic Similarity between Geospatial Conceptual Regions*, in *Geo-Spatial Semantics - First International Conference, GeoS 2005, Mexico City, Mexico, November 2005*, A. Rodriguez, et al., Editors. 2005, Springer: Berlin. p. 90-106.
13. Jurafsky, D. and J. Martin, *Speech and Language Processing - An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. 2nd ed. 2009, Upper Saddle River, New Jersey: Pearson - Prentice Hall.
14. Johnson, M., *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason*. 1987, Chicago: The University of Chicago Press.
15. Goldstone, R. and L. Barsalou, *Reuniting perception and conception*. Cognition, 1998. 65: p. 231-262.
16. Bobillo, F. and U. Straccia, *An OWL Ontology for Fuzzy OWL 2*, in *Foundations of Intelligent Systems, 18th International Symposium, ISMIS 2009, Prague, Czech Republic, September 14-17, 2009, Proceedings*, J. Rauch, et al., Editors. 2009, Springer: Berlin. p. 151-160.
17. Gärdenfors, P. and M. Williams. *Reasoning about Categories in Conceptual Spaces*. in *Fourteenth International Joint Conference of Artificial Intelligence*. 2001. p. 385-392.
18. Janowicz, K., B. Adams, and M. Raubal, *Semantic Referencing - Determining Context Weights for Similarity Measurement*, in *Geographic Information Science - Sixth International Conference, GIScience 2010, Zürich, Switzerland, Sep. 14-17 2010, Proceedings*, S. Fabrikant, et al., Editors. 2010, Springer: Berlin.
19. Goodchild, M., *Citizens as sensors: the world of volunteered geography*. GeoJournal, 2007. 69: p. 211-221.
20. Corcho, O. and R. García-Castro, *Five challenges for the Semantic Sensor Web*. Semantic Web - Interoperability, Usability, Applicability, 2010 (in this issue).
21. Gangemi, A. and V. Presutti, *Towards a Pattern Science for the Semantic Web*. Semantic Web - Interoperability, Usability, Applicability, 2010 (in this issue).
22. Guizzardi, G., *Theoretical Foundations and Engineering Tools for Building Ontologies as Reference Conceptual Models*. Semantic Web - Interoperability, Usability, Applicability, 2010 (in this issue).