Modeling Collaborative Semantics with a Geographic Recommender

Christoph Schlieder

Laboratory for Semantic Information Technology
University of Bamberg, Germany
christoph.schlieder@wiai.uni-bamberg.de

In the Semantic Web paradigm, geo-ontologies are closely related to geospatial information communities. Each community comes with its own ontology which is modeled, most frequently, within the framework of description logics. The paper questions a central assumption underlying this approach, namely that communities (and ontologies) are defined by crisp semantic boundaries. The idea of a semantic boundary contrasts sharply with the notion of a community of data producers/consumers that characterizes Web 2.0 applications. Well-known examples are GPS-trail libraries for hikers and bikers or image libraries of places of touristic interest. In these applications, conceptualizations are created as folksonomies by voluntary contributors who associate georeferenced objects (e.g. trails, images) with semantic tags. We argue that the resulting folksonomy can not be considered an ontology in the sense of Semantic Web technology and we propose a novel approach for modeling the collaborative semantics of geographic folksonomies. This approach is based on multi-object tagging, that is, the analysis of tags that users assign to composite objects, e.g. a group of photographs.

Subject areas: Spatio-temporal ontologies and reasoning; Definition, extraction, elicitation, and capture of spatial semantics

1 Introduction

Communities of voluntary contributors have been successful at building extensive collections of georeferenced data. Libraries of mountain biking trails recorded with handheld GPS receivers constitute typical examples (Morris et al., 2004). The idea of community-based mapping has been extended to features that, unlike biking trails, belong to the core topographic data set, most notably to road networks[1]. Generally, the rationale behind such low-quality remapping efforts is that of creating a data set which comes with a public domain license. Another type of application that is capable of attracting contributors in great numbers are libraries of georeferenced photographs. For instance, in the United Kingdom and the Republic of Ireland, amateur photogra-

[1] For example: Open Street Map Project (http://www.openstreetmap.org/).

phers collect images of geographic features aiming at documenting each square kilometer of the British Isles with at least one photograph[2].

Most forms of collaborative acquisition of georeferenced data include the acquisition of metadata. Spatio-temporal coverage (Where and when was the image taken, the trail recorded?) and content type (What objects are depicted? What type of trail is it?) constitute indispensable parts of content-related metadata sets. *Social tagging* is probably the simplest way to obtain this metadata from a community. Data producers categorize their contribution by assigning the key words ("tags") that they think best characterize the spatio-temporal coverage and the content type. The resulting tag vocabulary is called a *folksonomy*.

When O'Reilly (2005) described his vision of the Web's future, folksonomies were already widely used and he could present them as a typical tool of the Social Web, or, as he preferred to called it, the Web 2.0. From the point of view of conceptual modeling, the Social Web with its emphasis on collaborative metadata acquisition (Miller, 2005) and the Semantic Web with its emphasis on ontological engineering (Fensel et al., 2003) stress complementary aspects. This is especially true for GI processing where the OGC has always been very clear about the relationship between ontologies and geospatial information communities. The concept of *information community* is given a concise definition in the OGC white paper by Gould and Hecht (2001): "An information community is a group of people who share a common geospatial feature data dictionary (including definitions of feature relationships) and a common metadata schema."

We argue that the most remarkable difference between folksonomies and ontologies is not that of different degrees of formalization but the fact that they address different use cases. Folksonomies handle the modeling of semantics before the emergence of information communities and before crisp semantic boundaries have been established whereas the ontology languages of the Semantic Web (OWL, and to a lesser extent RDF) address the modeling of semantics after that phase. To support that view, we introduce a novel approach for modeling the collaborative semantics of geographic folksonomies.

The paper makes two main contributions: firstly, it bases semantic analysis on multi-object tagging instead of single-object tagging, and secondly, it introduces a new method for analyzing geospatial information communities by exploiting data about user-to-user similarity from a geographic recommender. *Multi-object tagging* is a form of tagging in which not individual items but groups of items are associated with a tag. For instance, a slide show consisting of photographs of tourist sites could constitute a multi-object and would be tagged with a single place concept such as "Southern France". The rest of the paper explores the semantics of multi-object tagging and is organized as follows.

First, we discuss proposals that have been made to formalize single-object tags (section 2). We show that in many cases a judgment about typicality can be recovered from multi-object tags and analyze consequences for the semantics of place concepts (section 3). We connect multi-object tagging to the problem of multi-object recommendation, that is, the task to assist users in selecting a multi-object based on his or

---

[2] Geograph project: http://www.geograph.org.uk/

her previous choices for such composite objects. Using a recommender system for composite images (patchwork postcards) as an example, we illustrate how to obtain data about multi-object tagging. Finally, we argue that for multi-object tagging, a key computational concept of recommending, namely user similarity, should be understood as a measure that indicates the extent to which two users share a common conceptualization. We show that as a consequence, the community of taggers is not going to be divided by crisp semantic boundaries (section 4). The paper concludes with a discussion of related work (section 5).

2 The meaning of tags

The Geograph project that was mentioned in the introduction provides a good example for the standard approach to folksonomies: *single-object tagging*. Information about the tags and their frequency can be retrieved directly from the project's web interface. In the beginning of May 2007, the folksonomy consisted of 2784 categories. It is revealing to sort them in order of decreasing frequency of use. The resulting list starts with the most popular tags, "Church", "Farmland", "Farm", and ends with tags that have been used only once: "Windmill stump", "luminous object in space (Sun)", "Penstock". It is known that folksonomies tend to show an exponential tag frequency distribution (Guy and Tonkin, 2006). In other words, tag frequency follows a power law: 36% of the Geograph tags are used only once, a further 24% appear just 2-5 times in the collection of 422.895 images whereas the most frequent tag "Church" is used 17.360 times. Not surprisingly, the Geograph tag set also shows other properties characteristic of folksonomies. A tag comes in different inflections ("house" vs. "houses"). Synonyms ("mansion" vs. "manor house"), homonyms ("bank" vs. "bank"), and hypernyms ("geological feature" vs. "rock outcrop") appear as tags. And last, but not least, some images are not categorized correctly.

Irrespective of these problems, the Geograph folksonomy may be considered a "shared understanding … in a given subject area" as the frequently quoted definition of ontology by Ushold and Gruninger (1996) requires. We could expect, for instance, the tag set to be disjoint from that of medical folksonomies. However, the understanding that the community of Geograph users has agreed upon is rather limited as it ignores all semantic relations between tags. The ontologies of the Semantic Web, in contrast, are formalized by specifying relations between concepts in form of role restrictions in description logic formulas. This means that while folksonomies can be considered ontologies in the sense of the above defintion, they definitely do not qualify as ontologies in the sense of Semantic Web technology.

Some Web 2.0 applications permit *user tagging* which means that not just the original contributor but also other users may categorize a piece of data. User tagged data items generally receive more than one category tag. We argue that it is user tagging with its potential of semantic conflict which makes folksonomies an interesting alternative to the Semantic Web approach to conceptual modeling. Gruber (2005) pointed out that tagging should not be seen as a binary relation *tagging*(*object*, *tag*). Semantic analysis should instead consider more complex relations all of which have at least three arguments: *tagging*(*object*, *tag*, *user*). In the following, we restrict our attention to Social Web applications capable of recording and evaluating information about user tags corresponding to this ternary tagging relation.

Generally, Social Web applications process tags in order to obtain information about the correct categorization of objects. A typical question asked in that context is: which of the many tags given to *image23.jpg* are the correct ones? We propose to study the converse problem which consists in processing tags to learn about the semantics of the categories. We ask questions such as: which images in the collection have been considered images of historical buildings? This approach amounts to recover from tag data the characterization of a concept by the set of its instances.

However, object categorization and concept characterization are not fully symmetric problems. A single user tag can provide a correct categorization for an object and a few tags may be sufficient to inform us about all categorizations relevant in a given problem domain. This is certainly not true of the *concept characterization problem* because the set of instances is not finite for most relevant concepts. An interesting variant of the concept characterization problem arises if we move from single-object tagging to multi-object tagging.

3 Multi-object tagging

Probably the most frequent type of aggregated objects that users of Social Web applications encounter are collections of (atomic) objects which result from a selection process in the problem domain. The 50 images that have been chosen to illustrate a mobile travel guide to Paris or the 7 roundtrips that a geographic recommender system suggests making during a Weekend stay in Florence constitute *collections of selected objects*. Such collections are not restricted to the digital world. A bestselling book describes 1,000 places to see before you die[3] and the printed patchwork postcard in Fig. 1 features a selection of 8 images that convey a multi-faceted view of the City of Bamberg.

Collections of selected images frequently come with a tagging that informs about the selection criterion. This type of tag can be interpreted as a caption that applies to all images of the collection. Actually, this is exactly the purpose of the place name caption “Bamberg” on the printed patchwork postcard in Fig. 1. Collections of selected images that come with a caption or are tagged otherwise constitute examples of *multi-object tagging*. Formally, we can still use the ternary tagging relation but with a collection as first argument: *tagging*({*object1, … , objectN*}, *tag*, *user*). This generalizes the description of tagging given by Gruber (2005) since single-object tagging can be treated as the special case of a one object collection.

We argue that multi-object tagging shows a semantic characteristic that is missing in single-object tagging, namely that of a selection process which chooses objects that are (1) typical instances of the concept and, at the same time, (2) show the variability of the instances. A patchwork postcard with the caption “London” featuring 8 images of Big Ben would satisfy criterion (1) but not (2). As tagging of multi-images is still

[3] Schultz, P. (2003). 1,000 Places To See Before You Die, Workman Publishing Company.

rare in today's image sharing platforms and data on user-tagged multi-images are not available, we take a look at printed patchwork postcards to support our argument.

Fig. 1. Printed patchwork postcard and its digital counterpart generated by the Tripost system

The sample consists of all patchwork postcards that are currently available of Bamberg, Germany. They show the historical center of the town, a listed UNESCO World Heritage Site. What makes this example particularly interesting is the fact that Bamberg has been listed as architectural ensemble and not because of a single building of outstanding value. Since there are many buildings of equal historical interest, the postcards could show a random sample of the city's built heritage with all historical buildings having about the same probability to appear on a postcard. There would be no typicality effect (criterion 1) in that case. On the other hand, if some historical site appears on almost all patchwork postcards we would assume that the image selection process is guided by some notion of typicality. Similarly, a variability effect (criterion 2) would be present if there are many images appearing on only very few, if any, other postcards. 15 patchwork postcards from Bamberg have been collected with a total of 108 images. However, only 24 different historical sites are depicted. Their frequency distribution is shown in Fig. 2 which reads as follows: half of the sites (12 of 24) appear on 1 to 3 postcards, a quarter (6 of 24) on 4 to 6 postcards, …, and just 1 site appears on 13 to 15 postcards. We find the frequency drop off that is characteristic of the power law, and consequently a first piece of evidence for both, the typicality and the variability of the image selection process.

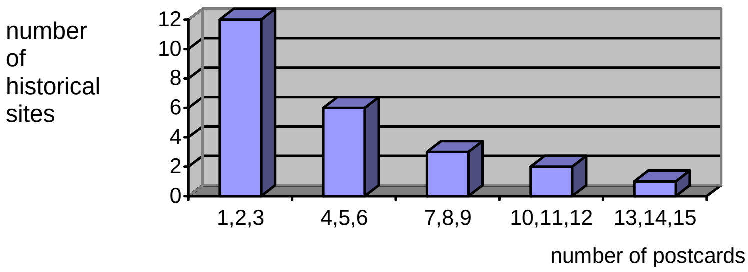

Fig. 2. Typicality of associations between a place name and a geographic object

We have seen that data about multi-object tagging can be useful to analyze the different meanings of a place concept. But from what sources can we obtain such data?

We suggest to collect multi-object taggings as a side-activity of running a recommender service for multi-objects as described in the next section.

4 Collaborative Semantics

Classical recommender systems solve an information filtering task (Burke, 2002). They suggest data sets, typically text documents describing some product or service, that seem likely to be relevant to the user based upon previous choices that this user has made. A widely known recommending approach is the item-to-item similarity approach used by Amazon to recommend books based on books one has bought before (Linden et al., 2003). It is easy to verify that the recommendations are based on item-to-item similarity rather than user-to-user similarity. With a travel guide to Bali in the shopping cart, the recommender suggest buying further travel guides to Bali and not, for that matter, books on Hinduism or surfing. A *geographic recommender system* recommends items from a library of georeferenced objects. In multi-object recommendation, collections of selected items are suggested to the user. A geographic multi-object recommender could make suggestions for a list of cities to visit or a slide show of images illustrating a certain place.

The latter use case is that of the Tripod Geographic Recommender. This recommender is currently implemented as part of the Tripod project which is concerned with the automatic generation of captions for georeferenced images including multi-images[4]. Fig. 3 shows the screenshot of the e-mail postcard application Tripost that takes advantage of the Tripod recommending service to automatically suggest patchwork postcards based on postcards that the user has previously built manually by selecting images from a photo sharing collection.

Data about the tagging of multi-images with place name tags is easily obtained from Tripost submissions to the recommender. Users select a fixed number of images to be shown on the postcard and provide a place name caption. You may use Tripost to generate e-mail postcards of your home town – never one like the other – without that you have to scan the collection for interesting new photographs. Also, you may want Tripost to send you a postcard from a location unknown to you which you specify by a place name. If you submitted, for instance postcards of Bamberg similar to the printed postcard shown in Fig. 1 and asked for a postcard from Florence you will probably get a postcard showing built heritage sites. Only probably, because this depends on what selections other users have made for images of Florence. Obviously, user-to-user similarity is more of a help in this case than item-to-item silimarity.

GroupLens is generally considered the first system that used a collaborative filtering approach for generating recommendations based on user-to-user similarity (Resnick et al., 1994). The authors of that landmark system characterize the basic assumption underlying their approach as “the heuristics that people who agreed in the past, are likely to agree again”. We transfer that idea from the temporal to the spatial domain – GroupLens made recommendations for news based on ratings of old news.

[4]Tripod – Automating caption creation: http://tripod.shef.ac.uk

Our spatial heuristic consists in assuming that people who agreed on the qualities of some places, are likely to agree on the qualities of other places too.

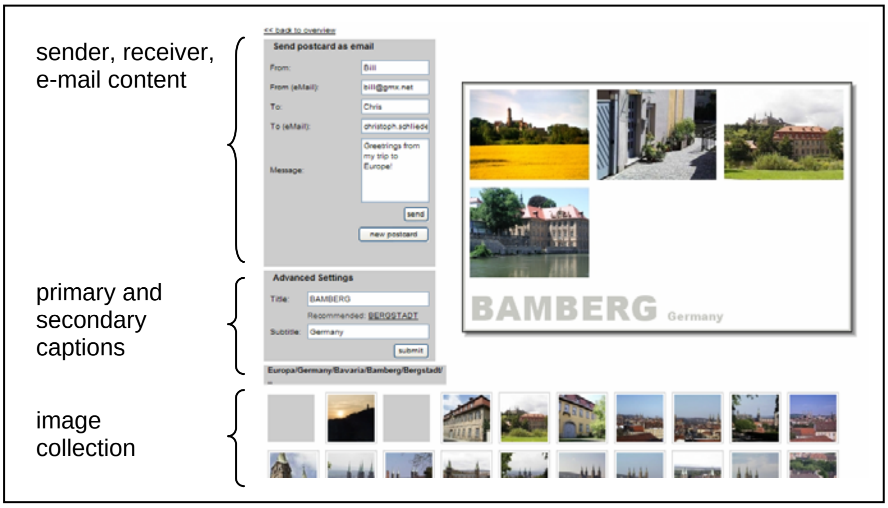

Fig. 3. Screenshot of TRIPOST, a service based on the Tripod Geographic Recommender

We consider a very simple example that helps to explain how a recommender based on user-to-user similarity can help to model the semantics of place concepts. The example data is summarized in Tab. 1. There are 4 place concepts (Antwerpen, Bamberg, Cardiff, Dublin) each of which comes with a collection of exactly 6 images numbered 1 to 6. Furthermore, 6 users (Anna, …, Franz) have submitted patchwork postcards, each showing 3 images, to the recommender with an appropriate place caption. The table entry Antwerpen-Anna, for instance, contains the information that user Anna has chosen image 1, 3, and 6 from the Antwerpen collection to represent that city:

tagging({*imageA1.jpg*,*imageA3.jpg*,*imageA6.jpg*}, *images of Antwerpen*, *Anna*).

The standard data table of a recommender is indexed by users (columns) and objects (rows). Since our goal consists in modeling the semantics of place concepts, we need to rewrite the standard table and index rows by concepts. The objects, that is, the images, now appear as table entries. User-to-user similarity compares users based on their previous selections. In our case, the similarity between two users, say, Bill and Emma, is based on the entries of Tab. 1.

Different similarity measures might be applied such as the variant of the Tversky measure proposed by Rodriguez and Egenhofer (2004) for the comparision of geospatial feature sets. To keep the example as simple as possible, we use the Tanimoto measure for similarity. Let A be the set of images that user a associates with place concept *Place* and B the set of images that user b associates with *Place*. We measure the similarity of the semantics of *Place* specified by A and the semantics of Place specified by B through the expression $\mathrm{sim}(A,B) = |A \cap B| / |A \cup B|$. In other words, the more images two patchwork postcards with the same place name caption share, the more similar is the place name semantics that they convey: sim({2,4,6}, {4,5,6} = 2/3 ≈ 0.66. Again for the sake of simplicity, similarity values are computed for entries that are non-void for both users and averaged: sim(Anna, Bill) = 1/2· (sim({1,2,6},{2,4,6}) + sim({2,4,5},{2,5,6}) = 1/2 · (2/3 + 2/3) ≈ 0.66 (Tab. 2).

	Anna	Bill	Clio	Don	Emma	Franz
Antwerpen	1, 3, 6		2, 4, 5	1, 4, 6		2, 3, 4
Bamberg		2, 3, 5	1, 4, 6	2, 4, 6	1, 4, 5	1, 2, 3
Cardiff	1, 2, 6	2, 4, 6	3, 5, 6	1, 3, 6	4, 5, 6	1, 2, 6
Dublin	2, 4, 5	2, 5, 6	1, 3, 5		1, 2, 3	2, 5, 6

Table 1. Results of multi-object tagging by users of a geographic recommender

	Anna	Bill	Clio	Don	Emma	Franz
Anna		.66	.22	.66	.33	.66
Bill	.66		.22	.33	.44	.77
Clio	.22	.22		.55	.66	.42
Don	.66	.33	.55		.33	.55
Emma	.33	.44	.66	.33		.33
Franz	.66	.77	.42	.55	.33	

Table 2. User-to-user similarity resulting from multi-object tagging

User-to-user similarity measures how similar two users are in their choices of images representing place concepts. We interpret this as a measure for a shared understanding between these users. This interpretation has interesting consequences for the notion of information community. Which are the users that share the understanding of place concepts of user Bill? There is no clear-cut answer because there are degrees of similarity. Most similar to Bill is Franz, then comes Anna, then Emma etc. An obvious way to define the “information community” of a user *u* consists in taking the set of the *k-nearest neighbors of u*. The boundaries of these communities are no longer crisp as *k* may take different values. Two communities that are disjoint when *k* corresponds to 5 % of the users may well overlap when *k* is chosen to correspond to 20 % of the users.

Another central property of information communities, as they are known from the Semantic Web, is lost. For “communities” induced by tagging, user *u* may be a member of the *k*-community of *v*, without that *v* is a member of the *k*-community of *u*. This can be seen from the 3-nearest neighbor communities in our example that are depicted in Fig. 4. Anna is a member of the community of Franz and vice versa, however, Franz is a member of the community of Clio without that Clio is a member of Franz’ community. It is the breaking down of the symmetry of mutual membership that leads us to question whether the notion of information community really makes sense in the context of tagging. The structuring of the users according to similarity seems rather to precede the emergence of information communities in the sense of the Semantic Web.

If this is true then we should think about a novel term to designate groups of users characterized by similarity.

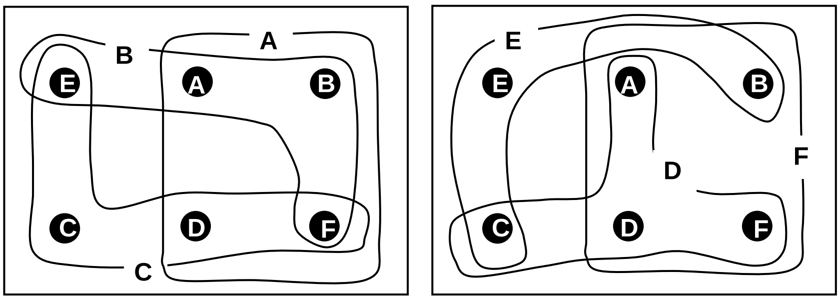

Fig. 4. Communities for *k*=3 of Anna, Bill, and Clio (left); Don, Emma, and Franz (right)

5 Related Work and Discussion

We have presented an approach to collaborative semantics that is based on multi-object tagging and defines “information communities” for folksonomies in terms of user-to-user similarity measurement. We illustrated the approach with an analysis of the semantics of place concepts that are used as tags to describe groups of images.

Our analysis shares the ternary interpretation of tagging with Gruber (2005) who points out that folksonomies and ontologies should not be seen in opposition because “there is increasing value in applying Semantic Web technologies to the Social Web”. However, we generalized the interpretation to include multi-object tagging. To our knowledge, the main idea presented in this paper, namely to analyze geospatial information communities by exploiting data about user-to-user similarity from a geographic recommender, has not been explored before.

Recently, Bishr and Kuhn (2007) have joined Gruber (2005) in rejecting the binary interpretation of tagging by providing arguments from a GI processing perspective. They suggest the use of typed tags (what-tags, where-tags, etc.) to enhance the semantic structure of folksonomies. With a different emphasis, Matyas (2007) has used description logics to implement data quality checks for the collaborative acquisition of geospatial data. While we agree with the main conclusion of these authors – sophisticated Social Web applications will have to make use of the knowledge representation technologies developed for the Semantic Web – we also see that the two approaches to semantics differ radically on the technological, the formal, and probably also on the philosophical level of analysis.

In conceptual modeling, folksonomies are often perceived as the poor man’s ontology. Community-based mapping projects seem to provide evidence for that view. Their spatial data accuracy and the quality of the specification of their feature semantics are often inferior to that of commercial products. However, by adopting this perspective, one would miss the real strength of folksonomies which lies in their ability to handle use cases where information communities have not yet emerged. Web 2.0 applications illustrate that data sharing happens long before it is possible to identify "a group of people who share a common geospatial feature data dictionary (including definitions of feature relationships)", as Gould and Hecht (2001) formulated in their definition of information community. We might therefore think of folksonomies as a way of modeling "pre-community semantics". In that perspective, the collaborative metadata acquisition methods of the Social Web and the ontological engineering methods of the Semantic Web appear truly complementary.

Acknowledgements

The author gratefully acknowledges support by the European Commission who funded parts of this research within the Tripod project under contract number IST-FP6-045335. He also whishes to thank Christian Matyas who implemented the Tri-post service for interesting discussions about the Tripod Geographic Recommender.

References

Bishr, M, & Kuhn, W. (2007). Geospatial Information Bottom-Up: A Matter of Trust and Semantics. *Proc. AGILE Int. Conf. on Geographic Information Science*, Aalborg University.

Burke, R. (2002). Hybrid recommender systems: Survey and experiments, *User Modeling and User-Adapted Interaction*, 12, pp. 331-370

Fensel, D., Wahlster, W., Lieberman, H., & Hendler, J. (2003): *Spinning the Semantic Web: Bringing the World Wide Web to Its Full Potential.* MIT Press, 2003

Gould Hecht (2001). A Framework for Geospatial and Statistical Information Integration. Open GIS Consortium White Paper.

Guy, M., & Tonkin, E. (2006). Folksonomies: Tidying up Tags? In: *D-Lib Magazine*, 12, 1, http://www.dlib.org/dlib/january06/guy/01guy.html.

Gruber, T. (2005). Ontology of Folksonomy: A Mash-up of Apples and Oranges. Keynote at the First Online Conference on Metadata and Semantics Research (MTSR'05). http://tomgruber.org/writing/ontology-of-folksonomy.htm

Linden, G., Smith, B., & York, J. (2003). Amazon.com Recommendations: Item-to-Item Collaborative Filtering, *IEEE Internet Computing*, Feb. 2003, pp. 76-80.

Matyas, S. (2007). Collaborative Spatial Data Acquisition, *Proc. AGILE Int. Conf. on Geographic Information Science*, Aalborg University.

Miller, P. (2005). Web 2.0: building the new library, *Ariadne* 45, published on http://www.ariadne.ac.uk/issue45/miller/intro.html

Morris, S., Morris, A., and Barnard, K. (2004). Digital trail libraries. *Proc. Joint Conference on Digital Libraries*, JCDL '04. ACM Press, pp. 63-71.

O'Reilly, T. (2005). What Is Web 2.0: Design Patterns and Business Models for the Next Generation of Software, published by O'Reilly Media on http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html

Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., & Riedl, J. (1994). GroupLens: An Open Architecture for collaborative Filtering of News. In: *Proc. Conf. Computer-supported cooperative work*, IEEE Press, pp. 175-186.

Rodríguez, A., & Egenhofer, M. (2004). Comparing geospatial entity classes: an asymmetric and context-dependent similarity measure. International Journal of Geographical Information Science, 18(3): p. 229-256.

Uschold, M., & Gruninger, M. (1996). Ontologies: principles, methods, and applications, *Knowledge Engineering Review*, 11, pp. 93-155.