Interoperating Geographic Information Systems, Interop '99, Zurich, Switzerland
A. Vckovski, K. Brassel, and H.-J. Schek (eds.),
Lecture Notes in Computer Science, Vol. 1580, pp. 189-202, Springer-Verlag, March 1999

Assessing Semantic Similarities Among Geospatial Feature Class Definitions*

M. Andrea Rodríguez[12], Max J. Egenhofer[123], and Robert D. Rugg[14]

[1] National Center for Geographic Information and Analysis, University of Maine, Orono, ME 04469-5711, USA
[2] Department of Spatial Information Science and Engineering, University of Maine, Orono, ME 04469-5711, USA
[3] Department of Computer Science, University of Maine, Orono, ME 04469-5752, USA
[4] Department of Urban Studies and Planning, Virginia Commonwealth University, Richmond VA 23284-2008, USA
{andrea, max}@spatial.maine.edu, rugg@vcu.edu

Abstract. The assessment of semantic similarity among objects is a basic requirement for semantic interoperability. This paper presents an innovative approach to semantic similarity assessment by combining the advantages of two different strategies: feature-matching process and semantic distance calculation. The model involves a knowledge base of spatial concepts that consists of semantic relations (is-a and part-whole) and distinguishing features (functions, parts, and attributes). By taking into consideration cognitive properties of similarity assessments, this model expects to represent a cognitively plausible and computationally achievable method for measuring the degree of interoperability.

1. Introduction

Since the first studies on interoperability, progress has been obtained concerning syntactic interoperability, i.e., data types and formats, and structural interoperability, i.e., schematic integration, query languages, and interfaces (Sheth 1998). As current information systems increasingly deal with information and knowledge issues, semantic interoperability becomes a major challenge for the next generation of interoperating information systems.

In information systems, semantics relates the content and representation of information to the entities or concepts in the world (Meersman 1997). The problem of semantic interoperability is the identification of semantically similar objects that belong to different databases and the resolution of their schematic differences (Kashyap and Sheth 1996). Schematic heterogeneity can only exist and, therefore, be solved for semantically similar objects (Bishr 1997). Studies have suggested the use of an ontology (Guarino and Giaretta 1995) as a framework for semantic similarity detection (Bishr 1997, Kashyap and Sheth 1998). One possible approach is to create a knowledge base in terms of a common ontology, upon which it is possible to detect semantic similarities and to define a mapping process between concepts (Lenat and Guha 1990, Kahng and McLeod 1998). On the

* This work was partially supported by the National Imagery and Mapping Agency under grant number NMA202-97-1-1023. Max Egenhofer's research is further funded by grants from the National Science Foundation under NSF grants SBR-8810917, IRI-9309230, IRI-9613646, SBR-9600465, and BDI-9723873; Rome Laboratory under grant number F30602-95-1-0042; the National Aeronautics and Space Administration under grant number COE/97-0015; Bangor Hydro-Electric Co.; and by a Massive Digital Data Systems contract sponsored by the Advanced Research and Development Committee of the Community Management Staff.

other hand, we can expect that in a realistic scenario new concepts will be added to or eliminated from the ontology. There may be different ways to classify a concept based on the specific application and the degree of detail for the concept's definition. Hence, the reuse and integration of existing domain specific ontologies becomes necessary (Kashyap and Sheth 1998, Mena *et al.* 1998).

This paper presents a computational model for similarity assessment among entity classes. We use the term entity classes to describe concepts in the real world and to distinguish their semantics from the semantics of data modeled and represented in a database. The latter case is called data semantics. Naturally, achieving semantic representation of objects in a database implies a good understanding of the semantics of the corresponding concepts in the real world. Consequently, our work considers studies done by cognitive scientists in the area of knowledge and behavior as well as by computer scientists in the domain of artificial intelligence.

The similarity model assumes a common ontology that includes the concepts' distinguishing features and interrelationships. A feature-matching process, together with a semantic distance computation, provides a strategy to create a model that satisfies cognitive properties of similarity assessment. In particular, we capture the idea that similarity assessment is not always a symmetric evaluation, similarity is a result of the commonalities and differences between two concepts, and the relevance of the distinguishing features (functions, parts, and attributes) may differ from one to another. In addition, is-a relations are complemented with part-whole relations to create an ontology that better reflects the interrelationships between concepts.

We focus on the domain of spatial information and we combine two existing sources of information, WordNet (Miller 1995) and the Spatial Data Transfer (USGS 1998), to create a common ontology that is used for the development of a prototype. The scope of this study includes only the evaluation of similarity within this common ontology. The analysis of how to integrate two domain specific ontologies is left for a future work.

The remainder of the paper is organized as follows. Section 2 reviews different approaches to the evaluation of semantic similarity. Section 3 describes the components for the definition of entity types. In Section 4 we present our similarity model and we illustrate its use with an example in Section 5. Finally, conclusions and future work are presented in Section 6.

2. Methods for Comparing Semantics

Most of the models proposed by psychologists are feature-based approaches, which use features that characterize entities or concepts (for example, properties and role). Using set theory, Tverski (1977) defined a similarity measure as a feature-matching process. It produces a similarity value that is not only the result of common features, but also the result of the differences between two entities. A different strategy for feature-based models is to determine a semantic distance between concepts as their Euclidean distance in a semantic, multidimensional space (Rips *et al.* 1973). This approach describes similarity by a monotonic function of the interpoint distance within a multidimensional space, where the axes in this space describe features of concepts. Krumhansl (1978) introduced the distance-density model based on a distance function for similarity assessment that complements the interpoint distance with the spatial density of the space. This model assumes that within dense regions of stimulus range finer discriminations are made than within relatively less dense subregions.

A shared disadvantage of feature-based models is that two entities are seen to be similar if they have common features; however, it may be argued that the extent to which a concept possesses or is associated with a feature may be a matter of a degree (Krumhansl 1978). Consequently, a specific feature can be more important to the meaning of an entity than another. On the other hand, the consideration of common features between entities seems to match the way people assess similarity.

With a different approach, computer scientists have defined similarity measures whose basic strategies make use of the semantic relations between concepts. These semantic relations are typically organized in a semantic network (Collins and Quillian 1969) as the links between nodes denote concepts. The semantic distance results in an intuitive and direct way of evaluating similarity in a hierarchical semantic network. For a semantic network with only is-a relations, Rada *et al.* (1989) pointed out that the semantic relatedness and semantic distance are equivalent and we can use the latter as a measure of the former. They defined conceptual distance as the length of the shortest path between two nodes in the semantic network. This distance function satisfies metric properties—minimality, symmetry, and triangle inequality.

Although the semantic distance models have been supported by a number of experiments and have shown to be well suited for a specific domain, they have the disadvantage of being highly sensitive to the predefined semantic-network architecture. In a realistic scenario, adjacent nodes are not necessarily equal. Irregular density often results in unexpected conceptual distance measures. Most concepts in the middle to high sections of the hierarchical network, being spatially close to each other, would therefore be deemed to be conceptually similar to each other. In order to account for the underlying architecture of the semantic network, Lee *et al.* (1993) argued that the semantic distance model should handle weighted indexing schema and variable edge weights. To determine weights the structural characteristics of the semantic network are typically considered, such as local density network, depth of a node in a hierarchical, type of link, and the strength of an edge link.

Some studies have considered weighted distance in a semantic network. Richardson and Smeaton (1996) used a hierarchical concept graph (HCG) derived from WordNet (Miller 1995) to determine similarity. They defined weights of links in a semantic network by the density of the HCG, estimated as the number of links, and by the link strength, estimated as a function of a node's information content value. Likewise Jiang and Conrath (1997) proposed the use of information content to determine the link strength of an edge. The information content of a node is obtained from the statistical analysis of word frequency occurrences in a corpus. The general idea of the information content is that as the probability of occurrence of a concept in a corpus increases, informativeness decreases such that the more abstract a concept, the lower its information content.

Richardson and Smeaton (1996) and Richardson *et al.* (1994) used a hierarchical network and information theory to propose an information-based model of similarity. Their approach to modeling semantic similarity makes use of the information content as described above, but it does not include distance as a basic strategy for similarity assessment. Conceptual similarity is considered in terms of class similarity. The similarity between two classes is approximated by the information content of the first superclass in the hierarchy that subsumes both classes. In the case of multiple inheritance (Cardelli 1984), similarity can be determined by the best similarity value among all various senses the classes belong to. The information-content model requires less information on the detailed structure of the network. On the other hand, many polysemous words and multi-worded classes will have an exaggerated information content value. The information-content model can generate a coarse result for the comparison of concepts since it does not differentiate the similarity values of any pair of concepts in a sub-hierarchy as long as their "smallest common denominator" is the same (Jiang and Conrath 1997).

Coming from the cognitive-linguistics domain, Miller and Charles (1991) discussed a contextual approach to semantic similarity. They developed a measure for similarity that is defined in terms of the degree of substitutability of words in sentences. For words from the same syntactic category and the same domain, the more often it is possible to substitute one word by another within the same context, the more similar the words are. The problem with this similarity measure is that it is difficult to define a systematic way to calculate it.

Interoperating Geographic Information Systems, Interop '99, Zurich, Switzerland
A. Vckovski, K. Brassel, and H.-J. Schek (eds.),
Lecture Notes in Computer Science, Vol. 1580, pp. 189-202, Springer-Verlag, March 1999

Based on our analysis of current models for semantic similarity, we propose a combination of the features-matching process and the evaluation of semantic distance. We expect that this interpreted model will provide a similarity measure that is not only cognitive plausible, but also computational achievable.

3. Components of Entity Class Definitions

Important components of the entity class definitions are the semantic interrelation among classes. We work on a specific domain, spatial information systems, and we describe the set of entity classes and their semantic relations as an ontology. In artificial intelligence, the term ontology has been used in different ways. Ontology has been defined as a "specification of a conceptualization" (Gruber 1995) and as "logical theory which gives an explicit, partial account of a conceptualization" (Guarino and Giaretta 1995). Thus, an ontology is a kind of knowledge base that has an underlying conceptualization. For our purpose, an ontology will be used as a body of knowledge that defines (1) primitive symbols used in the meaning representation and (2) a rich system of semantic relations interconnecting those symbols.

The most common semantic relation used in an ontology is the is-a relation, also called hypernymic or superordinate relation. This relation goes from a specific to a more general concept such that resembles the generalization mechanism of the object-oriented theory (Dittrich 1986). The is-a relation is a transitive and asymmetric relation that defines a hierarchical structure, where terms inherit all the characteristics of their superordinate terms.

Mereology, the study of part-whole relations (Guarino 1995), plays another important role for ontology. Studies have usually assumed that part-whole relations are transitive such that if *a* is part of *b* and *b* is part of *c*, then *a* is part of *c* as well. Linguists, however, have expressed their concerns about this assumption (Cruse 1979, Iris *et al.* 1988). Explanations to the transitive problem rely on the idea that part-whole relations are not one type of relation, but a family of relations. Winston *et al.* (1987) defined six types of part-whole relations: component-object (e.g., pedal-bike), member-collection (e.g., tree-forest), portion-mass (e.g., slice-cake), stuff-object (e.g., steel-bike), feature-activity (e.g., paying-shopping), and place-area (e.g., oasis-desert). Chaffin and Herrmann (1988) extended the previous classification with a seventh meronymic relation, phase-process (e.g., adolescence-growing up). For this work, we only consider the component-object relation with the properties of asymmetry and (with some reservations) transitivity.

When defining entity classes, the part-whole converse relations do not always hold. For example, we can say that a building complex has buildings, i.e., building complex is the whole for a set of buildings; however, buildings are not always part of a building complex. Thus, we distinguish the two relations, "part-of" and "whole-of," to be able to account for those cases.

Although the general organization of the entity classes is given by their semantic interrelations, we consider that this information is not enough to distinguish one class from another. For example, we can derive that a hospital and an apartment building have a common superclass building; however, that information is insufficient to differentiate a hospital from an apartment building. Considering that entity classes correspond to nouns in linguistic terms, we borrow Miller's (1990) description of nouns and propose to incorporate what he called *distinguishing features* to each class. Distinguishing features are classified into parts, functions, and attributes.

Parts are structural elements of a class, such as roof and floor of a building. We could make a further distinction between "things" that a class must have ("mandatory") or can have ("optional"). Note that parts are related to the relation part-whole previously discussed. While the relation part-whole works at the level of entity-class definitions and forces us to define all the entity classes involved, parts features can have items that are not always defined as entity classes in our model.

Interoperating Geographic Information Systems, Interop '99, Zurich, Switzerland
A. Vckovski, K. Brassel, and H.-J. Schek (eds.),
Lecture Notes in Computer Science, Vol. 1580, pp. 189-202, Springer-Verlag, March 1999

Function features are intended to represent what is done to or with a class. For example, the function of a college is to educate. Thus, function features can be related to other terms such as *affordances* (Gibson 1979) and *behavior* (Khoshafian and Abnous 1990). Attributes correspond to additional characteristics of a class that are not considered by either the set of parts or the set of functions. For example, some of the attributes of a building are age, user type, owner type, and architectural properties. Using a lexicon categorization, parts are given by nouns, functions by verbs, and attributes by nouns whose associated values are given by adjectives or other nouns.

In addition to semantic relations and distinguishing features, two more linguistic concepts are taken into consideration for the definition of entity classes. Entity classes are associated with concepts represented in natural language by words. Natural language understanding distinguishes two problems of the mapping between words and meanings, polysemy and synonymy. Polysemy arises when the same word may have more than one meaning, different *senses*. Synonymy corresponds to the case where two different words have the same meaning (Miller *et al.* 1990). Our class-entity definition incorporates synonyms, such as parking lot and parking area, and different senses of entity classes, such as the case when bank could be an elevation of the seafloor, a sloping margin of a river, an institution, or a building.

4. A Computational Method for Assessing Similarities of Entity Classes

We introduce a computational model that assesses similarity by combining a feature-matching process with a semantic distance measurement. While our model uses the number of common and different features between two entities, it defines the relevance of the different features in terms of the distance in a semantic network.

For each type of distinguishing features (i.e., parts, functions, and attributes) we propose to use a similarity function $S_t(c_1,c_2)$ (Equation 1) that is based on the *ratio model* of a feature-matching process (Tversky 1977). In $S_t(c_1,c_2)$ c_1 and c_2 are two entities classes, t symbolizes the type of features, and C_1 and C_2 are the respective sets of features of type t for c_1 and c_2. The matching process determines the cardinality (#) of the set intersection ($C_1 \cap C_2$) and set difference ($C_1 - C_2$), defined as the set of all elements that belong to C_1 but not to C_2.

$$S_t(c_1,c_2) = \frac{\{C_1 \cap C_2\}_\#}{(\{C_1 \cap C_2\}_\# + \alpha\{C_1 - C_2\}_\# + (1-\alpha)\{C_2 - C_1\}_\#)} \quad (1)$$

This similarity function yields values between 0 and 1. The extreme value 1 represents the case when everything is common between two entity classes, whereas the value 0 occurs when everything is different between two entity classes. The weight α is determined as a function of the distance between the entity classes (c_1 and c_2) and the immediate superclass that subsumes both classes. This corresponds to the least upper bound (l.u.b.) between two entity classes in partially ordered sets (Birkhoff 1967). When one of the concepts is the superclass of the other, the former is also considered the immediate superclass (l.u.b.) between them. The distance of each entity class to the l.u.b. is normalized by the total distance between the two classes, such that we obtain values in the range of 0 to 1. Then, to obtain the final values of α, we define an asymmetric function (Equation 2).

$$\alpha_{(c_1,c_2)} = \begin{cases} \dfrac{d(c_1,l.u.b.)}{d(c_1,c_2)} & d(c_1,l.u.b.) \le d(c_2,l.u.b.) \\ 1 - \dfrac{d(c_1,l.u.b.)}{d(c_1,c_2)} & d(c_1,l.u.b.) > d(c_2,l.u.b.) \end{cases}$$

Interoperating Geographic Information Systems, Interop '99, Zurich, Switzerland
A. Vckovski, K. Brassel, and H.-J. Schek (eds.),
Lecture Notes in Computer Science, Vol. 1580, pp. 189-202, Springer-Verlag, March 1999

The assumption behind the determination of α is that similarity is not necessarily a symmetric relation (Tversky 1977). For example, "a hospital is similar to a building" is a more general agreement than "a building is similar to a hospital." It has been suggested that the perceived distance from the prototype to the variant is greater than the perceived distance from the variant to the prototype, and that the prototype is commonly used as a second argument of the evaluation of similarity (Rosch and Mervis 1975, Krumhansl 1978). Hence, we assume that the non-common features of the concept used as a reference (the second argument) should be more relevant in the evaluation.

An interesting case occurs when comparing a class with its superclass or vice versa. Since subclasses inherit all features of their superclasses, only subclasses may have non-common features. It can be easily seen that when comparing a class with its superclass or vice versa, the weight associated with the non-common features of the first argument is 0 (α) and the weight for the non-common features of the second argument is 1 ($1-\alpha$). By considering the direction of the similarity evaluation, a class will be more similar to its superclass than the same superclass to the class. Currently and with the purpose of calculating the weight α, the part-of relation is treated like the is-as relation. The difference of these two relations lays on the inheritance property of the is-a relation. The effect of the part-of relation can be illustrated when comparing a building with a building complex or vice versa. With our model a stronger similarity is found between the building and the building complex than between the building complex and the building. Note, however, that the similarity between the whole and its parts could also be higher, since there is not an inheritance property for this semantic relations that forces us to have all features of the whole also in its parts.

Synonyms are incorporated into the evaluation of similarity when searching for an entity class at the beginning of the evaluation. In addition, synonyms are also taken into account in the matching process of parts, functions, and attributes. Each term (entity class, part, function, or attribute) is treated in the same way as its synonyms. Words with different semantics or senses (polysemy) are also included in our model. We handle different senses of entity class as independent entity classes with a common name. For parts, functions, and attributes, we first match the senses of the terms and then we evaluate the set-intersection or set-difference operation among the set of features. A term in one sense might have a set of synonyms, therefore, we match terms or their synonyms that belong to the same sense. For example, the verb "to play" has two different senses in our database, play for recreation and play for competition. For any entity class that has the function "to play," the knowledge base also includes the sense of the word such that the system can find the synonyms of "to play" for the respective sense.

The global similarity function $S(c_1,c_2)$ is a weighted sum of the similarity values for parts, functions, and attributes (Equation 3), where ω_p, ω_f, and ω_a are weights of the similarity values for parts, functions, and attributes, respectively. These weights define the importance of parts, functions, and attributes that might vary among different contexts. The weights all together must add up to 1.

$$S(c_1,c_2) = \omega_p \; S_p(c_1,c_2) + \omega_f \; S_f(c_1,c_2) + \omega_a \; S_a(c_1,c_2) \quad ($$

5. An Example

We have implemented a software prototype for the similarity assessment. It used WordNet (Miller 1995) and the Spatial Data Transfer Standard (SDTS) (USGS 1998) to derive a knowledge base. From SDTS we extracted the entity classes to be defined, their partial definition of is-a relations, and the attributes for entity types. By using WordNet we complemented the is-a relations with the part-whole relations and we obtained the structural elements (parts) of entity types. Finally, functions were derived as a combination of the functions or verbs explicitly used in the description of entity classes and a common sense determination.

To illustrate the use of our model for interoperability, consider an urban-planning application that deals with the urban rehabilitation of the downtown of a city. To accomplish the goal, planners have decided to analyze and compare the downtowns of cities of similar sizes that are considered examples of high quality life. In a first step, planners are concerned about the functional components of the downtown, i.e., type of spatial entities, and they have left for *a posteriori* analysis the geometric distribution of those components.

Maps of each downtown are obtained from different spatial databases and at the semantic level we face the problem of comparing the semantic of entity classes. For the time being, we assume that maps are based on a common ontology because they were created by using the same conceptualization. Although the assumption of a unique ontology simplifies the problem of interoperability, the problem of different classification within the same ontology remains possible. For example, what was identified as a sidewalk on one map, it could be identified as a path in another one with a different criteria. This type of problem resembles the abstract level incompatibility discussed by Kashyap and Sheth (1996) when describing the schematic heterogeneities in multidatabases.

Our approach to accomplish the planners' objective is to evaluate the semantic similarity by searching for the best match, entity-to-entity, between two downtown maps. A portion of the knowledge base used for this application, representing an ontology with only is-a relations, is shown in Figure 1. In Figure 1 entities that represent cases of polysemy (i.e., different entity classes with same name but multiple meanings) and cases of entities classes with multiple superclasses (i.e., same entity class with multiple inheritance) are highlighted. Figure 2 shows the complete description of an entity class, i.e., its distinguishing features and its semantic relations.

Since the planners in our example are mostly concerned with the functional components of the downtowns, they may assign a higher weight to the function features. For example, 50% for function features, 25% for part features, and 25% for attribute features. For this application, the direction of the evaluation is determined by the target downtown (the downtown to be redesigned) against which the ideal downtowns are compared.

Figure 3 shows a similarity assessments between a stadium and all other possible entity classes in the knowledge base. Numerically, only four entity classes have similarity values higher than or equal to 0.5: arena (0.78), athletic field (0.62), tennis court (0.6), and construction (0.5). When changing the direction of the evaluation, for example athletic field against all entity classes, it is impossible to notice the asymmetric evaluation of the similarity model. For athletic field, only three entity classes have a similarity value greater than or equal to 0.5: tennis court (0.95), stadium (0.58), and arena (0.57).

Interoperating Geographic Information Systems, Interop '99, Zurich, Switzerland
A. Vckovski, K. Brassel, and H.-J. Schek (eds.),
Lecture Notes in Computer Science, Vol. 1580, pp. 189-202, Springer-Verlag, March 1999

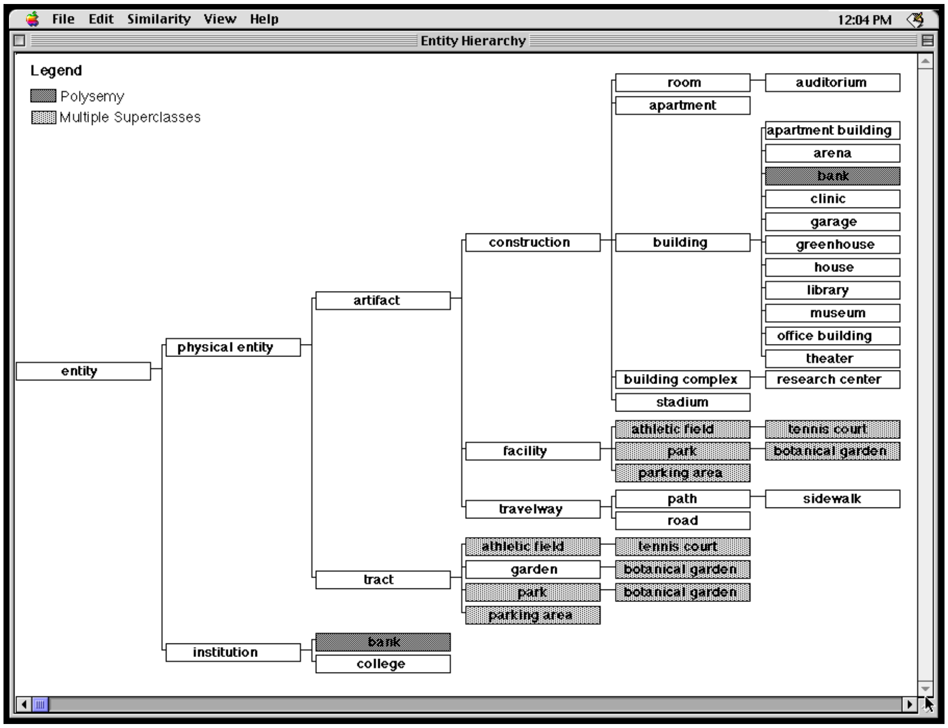

Figure 1: Entity class hierarchy (is-a relations).

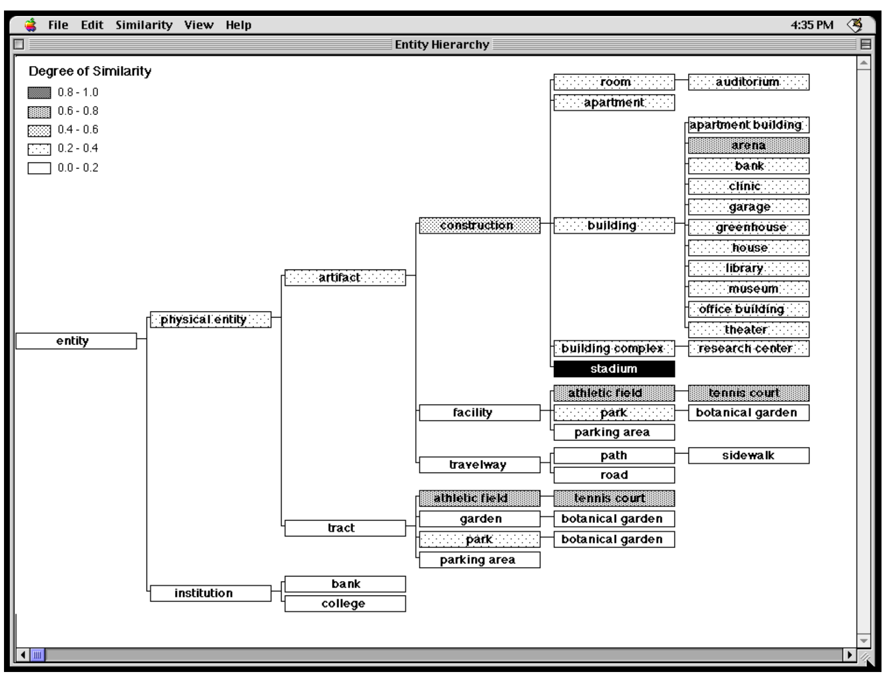

Figure 3: Similarity assessment: stadium against all entity classes.

6. Conclusions and Future Work

Our model of semantic similarity has a strong linguistics basis. It introduces synonyms and different senses in the use of terms. It also provides a first approach to handle part-whole relations in the evaluation of semantic similarity. Furthermore, it defines a semantic-similarity function that is asymmetric for classes that belong to different levels of generalization in the semantic network. Although the model is affected by the definition of parts, functions, and attributes, it reduces the effect of the underlying semantic network when compared with many of the semantic distance models.

As defined by our model, the asymmetric weights for the non-common features of each entity class (α, and $1 - \alpha$) add up to 1. That means that as a total, common and different features have the same weight (i.e., 1). A further refinement can be done to the definition of the weights α and $(1-\alpha)$ if we consider that in the assessment of similarity people may attend to give more importance to the common features (Tverski 1977, Krumhansl 1978).

The global semantic similarity assessment for spatial scenes could also be improved. Our approach evaluates entity-to-entity similarity to obtain a global optimization of the similarity between two scenes. Problems arise when scenes have different numbers of spatial entities. A study of how much non-common entities affect the global similarity assessment will help to obtain a better estimation of the semantic similarity between spatial scenes.

Context has been already suggested to be a relevant issue for semantic similarity (Tversky 1977, Krumhansl 1978) and for interoperability (Kashyap and Sheth 1996, Bishr 1997). We expect to incorporate context, initially through matching a user's intended operations with operations associated with the compared classes, in order to recognize different senses (semantics) of entity classes as well as to be able to define weights that reflect characteristics of a specific application.

Human-subject testing will contribute to testing how closely our model resembles people's similarity judgments. It might also provide us new insights about how important for people are common and non-common features.

Finally, a big challenge for our model is to evaluate similarity across multiple knowledge databases or ontologies. As we assumed ontologies for specific domains and customized by users, we found significant differences in the definition of concepts within a same domain. In order to move forward in the solution of interoperability systems we will need to account for those differences and relax our assumption of a unique ontology by a common ontology that integrates multiple and independent domain specific ontologies.

References

G. Birkhoff (1967). *Lattice Theory*. American Mathematical Society, Providence, RI.

Y. Bishr (1997). *Semantic Aspects of Interoperable GIS*I. Thesis Wageningen Agricultural University and International Institute for Aerospace Survey and Earth Science (ITC), Enshede, the Netherlands.

L. Cardelli (1984). A Semantics of Multiple Inheritance. in: G. Kahn, D. McQueen, and G. Plotkin (Eds.), *Semantics of Data Types*. Lecture Notes in Computer Science 173:51-67, Springer-Verlag, New York, NY.

R. Chaffin and D. Herrmann (1988). The Nature of Semantic Relations: a Comparison of two Approaches. in: M. Evens (Ed.), *Relational Models of the Lexicon: Representing Knowledge in Semantic Network.* pp. 289-334, Cambridge University Press, Cambridge, MA.

A. Collins and M. Quillian (1969). Retrieval Time from Semantic Memory. *Journal of Verbal Learning and Verbal Behavior* 8: 240-247.

K. Dittrich (1986). Object-Oriented Data Base Systems: The Notation and The Issues. in: K. Dittrrich, U. Dayal, and A. Buchmann (Eds.), *International Worskshop in Object-Oriented Database Systems*, Pacific Grove, CA, Washington, D.C. 2-4, IEEE Computer Society Press.

J. Gibson (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston, MA.

T. Gruber (1995). Toward Principles for the Design of Ontologies Used for Knowledge Sharing. *International Journal of Human and Computer Studies* 43(5/6): 907-928.

N. Guarino (1995). Formal Ontology, Conceptual Analysis, and Knowledge Representation. *International Journal of Human and Computer Studies* 43(5/6):625-640.

N. Guarino and P. Giaretta (1995). Ontologies and Knowledge Bases: Towards a Terminological Clarification. in: N. Mars (Ed.) *Toward Very Large Knowledge Bases: Knowledge Building and Knowledge Sharing*, Amsterdam, the Netherlands. pp. 25-32, IOS Press, Amsterdam.

N. Guarino (1997). Semantic Matching: Formal Ontological Distinctions for Information Organization, Extraction, and Integration. in: M. Pazienza (Ed.) *Information Extraction: A Multidisciplinary Approach to Engineering Information Theory*, Frascati, Italy. pp. 139-170, Springer Verlag.

J. Jiang and D. Conrath (1997). Semantic Similarity Based on Corpus Statistics and Lexical Taxonomy. *International Conference on Computational Linguistics (ROCLING X)*. Taiwan, pp. 19-35.

J. Kahng and D. McLeod (1998). Dynamic Classificational Ontologies: Mediation of Information Sharing in Cooperative federated Database Systems. in: M. Papazoglou and G. Schlageter (Eds.) *Cooperative Information Systems: Trends and Directions*. pp. 179-203 , Academic Press, London, UK.

V. Kashyap and A. Sheth (1996). Semantic and Schematic Similarities Between Database Objects: A Context-Based Approach. *The VLDB Journal* 5(4):276-304.

V. Kashyap and A. Sheth (1998). Semantic Heterogeneity in Global Information Systems: The Role of Metada, Context, and Ontologies. in: M. Papazoglou and G. Schlageter (Eds.) *Cooperative Information Systems: Trends and Directions*. pp. 139-178, Academic Press, London, UK.

S. Khoshafian and R. Abnous (1990). *Object Orientation: Concepts, Languages, Databases, User Interfaces*. John Wiley & Sons, New York.

C. Krumhansl (1978). Concerning the Applicability of Geometric Models to Similarity Data: The Interrelationship Between Similarity and Spatial Density. *Psychological Review* 85(5): 445-4 63.

J. Lee, M. Kim, and Y. Lee (1993). Information Retrieval Based on Conceptual Distance in IS-A Hierarchies. *Journal of Documentation* 49(2): 188-207.

D. Lenat and R. Guha (1990). *Building Large Knowledge Based Systems: Representation and Inference in the Cyc Project*. Addison-Wesley Publishing Company, Reading, MA.

R. Meersman (1997). An Essay on The Role and Evolution of Data(base) Semantics. in: R. Meersman and L. Mark (Eds.) *DataBase Application Semantics*. Chapman Hall, London, UK.

E. Mena, V. Kashyap, A. Illarramendi, and A. Sheth (1998). Domain Specific Ontologies for Semantic Information Brokering on the Global Information Infrastructure. *International Conference on Formal Ontology Information Systems*. Available at http://ra.cs.uga.edu/publications/pub_ALL.html.

G. Miller (1990). Nouns in WordNet: A Lexical Inheritance System. *International Journal of Lexicography* 3(4): 245-264.

G. Miller (1995). A Lexical Database for English. *Communications of the ACM* 38(11):39-41.

G. Miller, R. Beckwith, C. Fellbaum, D. Gross, and K. Miller (1990). Introduction to WordNet: An On-Line Lexical Database. *International Journal of Lexicography* 3(4): 235-244.

G. Miller and W. Charles (1991). Contextual Correlates of Semantic Similarity. *Language and Cognitive Processes* 6(1): 1-28.

R. Rada, H. Mili, E. Bicknell, and M. Blettner (1989). Development and Application of a Metric on Semantic Nets. *IEEE Transactions on System, Man, and Cybernetics* 19(1): 17-30.

R. Richardson and A. Smeaton (1996). *An Information Retrieval Approach to Locating Information in Large Scale Federated Database Systems*. Dublin City University, School of

Computer Applications, Working Paper CA-0296, available at http://simpr1.compapp.dcu.ie/CA_Working_Papers/wp96.html#0296.

R. Richardson, A. Smeaton, and J. Murphy (1994). *Using WordNet as a Knowledge Base for Measuring Semantic Similarity Between Words.* Dublin City University, School of Computer Applications, Working Paper CA-1294, available at http://simpr1.compapp.dcu.ie/CA_Working_Papers/wp94.html#1294.html.

L. Rips, J. Shoben, and E. Smith (1973). Semantic Distance and the Verification of Semantic Relations. *Journal of Verbal Learning and Verbal Behavior* 12: 1-20.

E. Rosch and C. Mervis (1975). Family Resemblances: Studies in the Internal Structure of Categories. *Cognitive Psychology* 7: 573-603.

A. Sheth (1998). Changing Focus on in Information Systems: From System, Syntax, Structure to Semantics. in: M. Goodchild, M. Egenhofer, R. Fegeas, and C. Kottman, *Interoperating Geographic Information Systems*. Kluwer Academic Press (in press).

A. Tversky (1977). Features of Similarity. *Psychological Review* 84(4): 327-352.

USGS (1998). *View of the Spatial Data Transfer Standard (SDTS) Document*. Available at http://mcmcweb.er.usgs.gov/sdts/standard.html, Last modification: Friday, 12-Jun-98.

M. Winston, R. Chaffin, and D. Herramann (1987). A Taxonomy of Part-Whole Relations. *Cognitive Science* 11: 417-444.