

Putting Similarity Assessments into Context: Matching Functions with the User's Intended Operations*

M. Andrea Rodríguez and Max J. Egenhofer
National Center for Geographic Information and Analysis
and
Department of Spatial Information Science and Engineering
University of Maine
Orono, ME 04469-5711, USA
Phone: (207) 581-2188
Fax: (207) 581-2206
{andrea, max}@spatial.maine.edu

Abstract

This paper presents a practical application of context for the evaluation of semantic similarity. The work is based on a new model for the assessment of semantic similarity among entity classes that satisfies cognitive properties of similarity and integrates contextual information. The semantic similarity model represents entity classes by their semantic relations (is-a and part-whole) and their distinguishing features (parts, functions, and attributes). Context describes the domain of an application that is determined by the user's intended operations. Contextual information is specified by a set of tuples over operations associated with their respective entity-class arguments. Based on the contextual information, a partial word-sense disambiguation can be achieved and the relevance of distinguishing features for the similarity assessment is calculated in terms of the features' contribution to the characterization of the application domain.

Keywords: semantic similarity, entity classes, functions, part-whole relation, is-a relation, context.

Introduction

Similarity is a judgment process that requires two “things” to be decomposed into aspects or elements in which they are the same and aspects in which they are different. These types of judgments are typically intuitive, subjective, and part of the everyday life such that they usually display no strict mathematical models [1]. In information systems similarity assessment is part of several processes, such as information retrieval [2-4] and data integration [5, 6]. Similarity assessment is particularly important for applications in which no precise definitions underlie the matter of discourse. In such domains, data stored in a database represent particular views of reality and users' queries express only an approximation of what users want to retrieve, which is likely an inexact match to any stored data.

This paper explores the use and effect of context over a model of semantic similarity among entity classes, the matching-distance model [7]. Like feature-based models [1, 8, 9], the matching-distance model defines a similarity function in terms of common and different features of entities (i.e., descriptors and attributes). The matching-distance model, however, defines asymmetric evaluations of semantic similarity that are a product of the weighted contribution of the similarity among different types of distinguishing features (parts, functions, and attributes). The model makes use of a hierarchical structure constructed with is-a and part-whole relations to determine the level of abstraction of entity classes and the model's asymmetric factors.

* 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 grant numbers IRI-9613646, SBR-9600465, and BDI-9723873; the Airforce Research Laboratory under grant number F30602-95-1-0042; 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.

The matching-distance model has a strong basis in linguistics. Since entity classes are identified by words, this model takes into account two linguistic concepts—synonymy and polysemy—that characterize the mapping between words and meanings [10]. Polysemy arises when a word has more than one meaning (i.e., multiple *senses*). Synonymy corresponds to the case when two different words have the same meaning. In linguistics terms, the model represents entity classes as nouns, which are organized into sets of synonyms. Thus, this model allows not only the definition of synonyms, but also the identification of polysemous words.

Context becomes important for a similarity assessment, because it affects the determination of the relevant features [1, 8, 9]. Although a feature-based approach is sensible to the way people assess similarity, 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. Consequently, a specific feature can be more important to the meaning of an entity class than to another. Furthermore, since the matching-distance model allows polysemous words to occur, context helps to distinguish among entity classes identified by the same word (polysemous word).

The remainder of this paper is structured as follows: Section 2 describes the main characteristics of the matching-distance model. Section 3 presents the definition and role of contextual information for the similarity assessment. An example illustrates the use context in Section 4. Section 5 discusses the results found when introducing contextual information into the matching-distance model. Conclusions and future work are presented in Section 6.

2. The Matching-Distance Model

Focusing on the spatial domain, Rodríguez *et al.* [7] defined the matching-distance model for the evaluation of semantic similarity among entity classes. For this work entity classes denote concepts in the real world, rather than entities modeled in a database. These concepts in the real world are cognitive representations that people use to recognize and categorize objects or events [11]. In this sense, this work has a top-down approach by starting from the semantics of entities in the real world instead of the semantics of data stored in a database [12].

The matching-distance model complements the feature-based approach by using semantic relations to organize concepts and a semantic-distance function [13] to determinate the degree of generalization of entity classes. The degree of generalization reflects the prototypical characteristics of entity class definitions, such that more general concepts are located in the top level of the hierarchical structure. Two semantic relations are considered for the entity class definitions: is-a [14] and part-whole [15]. In addition to these semantic relations, the matching-distance model describes entity classes by their distinguishing features. The model differentiates distinguishing features by classifying them into parts, functions, and attributes. Parts are structural elements of a class, such as roof and floor of a building. Function features are intended to represent what is done to or with a class. For example, the function of a college is to educate. Attributes correspond to additional characteristics of a class, such as name, color, and owner. Since the matching-distance model deals with concepts instead of specific real world objects, the model matches type of attributes rather than attributes values. The classification of distinguishing features attempts to facilitate the implementation of the entity class representation as well as to enable the separated manipulation of each type of distinguishing features. The matching-distance model groups synonyms that identify the same entity class. While these synonym sets are considered equivalent for the similarity evaluation, each of the entity classes associated with a polysemous word is handled as an independent definition.

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

$$S(c_1, c_2) = \quad_p S_p(c_1, c_2) + \quad_f S_f(c_1, c_2) + \quad_a S_a(c_1, c_2) \quad (1)$$

Similarity assessment for each type of distinguishing features (i.e., parts, functions, and attributes) is given by Equation 2. In $S_t(c_1, c_2)$ c_1 and c_2 are two entity classes, t symbolizes the type of features, C_1 and C_2 are the respective sets of features of type t for c_1 and c_2 , $\#()$ is the cardinality of a set, \cap is the set intersection, and $-$ the set difference.

$$S_t(c_1, c_2) = \frac{\#(C_1 \cap C_2)}{\#(C_1 \cap C_2) + \#(c_1, c_2)\#(C_1 - C_2) + (1 - \#(c_1, c_2))\#(C_2 - C_1)} \quad (2)$$

The function S_t 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. Here a superclass denotes an entity class that is either *parent* or *whole* of an entity class. The function S_t provides asymmetric values for entity classes that belong to different level of generalization in a hierarchical structure. For instance, the similarity between a museum and a building is greater than the similarity between a building and a museum. The assumption behind the determination of S_t is that the concept used as a reference (the second argument) should be more relevant in the evaluation [8, 16]. In the study of semantic categories, Rosch [17] supported the view that categories are naturally formed and defined in terms of focal points or prototypes. She hypothesized that (1) in sentences such as “a is essentially b,” the focal stimuli (i.e., prototypes) appear in the second position, and (2) the perceived distance from the prototype to the variant is greater than the perceived distance from the variant to the prototype.

3. Context Specification

Context is an important aspect for such diverse areas as natural language processing (NLP), knowledge-based problem solving, database systems, and information retrieval [18-22]. Despite this recognition, the meaning of context in information systems is usually left to the user’s interpretation and its role may vary among different domains [23]. For NLP, context has a sense-disambiguation function [18] so that otherwise ambiguous statements become meaningful and precise. These studies analyze the meaning of words within either a topical context or a local context of a corpus [24]. Knowledge representation involves statements and axioms that hold in certain contexts; therefore, context determines the truth or falsity of a statement as well as its meaning [19]. For knowledge-based problem solving, context is usually defined as the situation or circumstances that surround a reasoning process [20, 25, 26]. Recent studies on data semantics and interoperability have stressed the importance of context to describe data content. In this domain, context is the knowledge needed to reason about another system [27], the intentional descriptions of database objects [21], and the extent of validity for an ontology [28]. For information retrieval, context provides a framework for well-defined queries and consequently, it improves the matching process between a user’s query and the data stored in a database [22].

Following the idea of Naive Physics [29] and Naive Geography [30], we could think in a common sense definition of entity classes. In such a case, we have a comprehensive characterization of entities that contains the entities’ essential properties. We could expect to obtain a good approximation of the similarity assessment among entity classes by considering equally these essential properties as equal. We argue, however, that a similarity assessment is used in an information with a purpose in mind and in this context, some features may be more important than others. Psychologists and cognitive scientists were the first to point out the importance of context for the determination of the relevant features in a similarity assessment [1, 8]. These studies have suggested that the relevance of features is associated with how diagnostic the feature is for a particular set of objects under consideration. The diagnosticity of features refers to how significant a feature is for classifying objects into subclasses [1]. Since a classification process produces groups of entities that share some features, the criteria for classifying a set of entities can be found by looking at these common features. Our work makes use of this notion of diagnosticity and

defines weighted values for the similarity among parts, function, and attributes (p , f , and a of Equation 1) by analyzing the frequency with which distinguishing features characterize entity classes of the domain of discourse. A probability is then translated into a high relevance. We support our approach by arguing that entities associated with a domain of discourse share some features that make them subject to interest for this domain. This way of determining the relevance of distinguishing features is in agreement with the idea that people tend to give more attention to similar than to distinctive features in the evaluation of similarity [1].

Our approach to the determination of the relevance of distinguishing features is based on the determination of a domain of discourse. By the domain of discourse we mean the set of entity classes that are subject of interest for an application. Since a domain of discourse may change among applications, the relevance of distinguishing features changes as well. We describe this situation as the context dependence of the similarity assessment, which is given by the user's intended operations. Our approach is based on the idea that people assess the meaning of a word within the context they can use it [31], which is congruent with the notion of use-based semantics [32]. The user's intended operations may be abstract, high-level intentions (e.g., "analyze" or "compare") or detailed plans (e.g., "purchase a house"). From a linguistics point of view, the user's intended operations are associated with verbs that denote actions. Verbs alone, however, may not be enough to completely describe operations, since they can change the operations' meaning depending on the kinds of noun arguments with which they co-occur [33]. For example, different senses of the verb *play* are *play a role*, *play the flute*, and *play a game*.

We specify contextual information as a set of tuples over operations associated with their respective noun arguments. The nouns correspond to entity classes in the matching-distance model (Equation 3), while the operations refer to verbs that are associated as methods with those classes. In this specification an entity-class argument may be empty if no further explanation is needed for describing the intended operation. Since the context specification uses operations and entity classes, the knowledge base used by the entity class representation of the matching-distance model can be extended to represent the components of the context specification.

$$Context = \langle \langle operation_1, \{entity_1 \dots entity_n\}_1 \rangle \dots \langle operation_m, \{entity_k \dots entity_l\}_m \rangle \rangle \quad (3)$$

For example, if a user wants to analyze some on-line datasets with the purpose of purchasing a cottage, she would describe her intention by $Context = \langle (purchase, \{cottage\}) \rangle$. By using the hierarchical structure of the knowledge base, the operations' arguments can be expressed at different levels of generalization. For example, a user may be looking for sports facilities and in such a case, she can specify $Context = \langle (search, \{sports\ facility\}) \rangle$ or $Context = \langle (search, \{athletic\ field, bowl\ park, tennis\ court, sports\ arena, stadium\}) \rangle$. Another user's intention can be described by using operations without arguments, such as $Context = \langle (play, \{\}) \rangle$. In this case, the operation *play* corresponds to a common function that characterizes the entity classes the user is looking for.

When contextual information is specified, not only relevance of distinguishing features can be determined, but also word-sense ambiguities may be solved. Since the domain of the application is usually a subset of the entire knowledge base, the contextual specification decreases the number of entity classes that possess the same name (i.e., polysemy). Unfortunately, this approach suffers the same disadvantage of the use of topical context for word-sense disambiguation [34], since it may not distinguish polysemous terms that are semantically similar and belong to the same domain of discourse.

The semantic relations among entity classes provide a flexible way to describe context since the specification of one entity class can be used to obtain other entity class that are semantically related. We follow a top-down approach in a hierarchical structure to retrieve all entity classes that belong to the application domain. This approach consists in selecting

- (1) entity classes whose functions correspond to the intended user's operations,

- (2) entity classes that are specified as domain of operations in the context specification, and
- (3) entity classes derived from a recursive search of parts and children of the entity classes found in (1) and (2).

Once entity classes of the application's domain are obtained, the diagnosticity of a type of feature t (P_t) is calculated as the sum of the frequency with which each distinguishing feature of the type t characterizes an entity class in a domain of discourse (Equation 4). In P_t , o_i is the number of occurrence of the feature i in the entity class definitions, n is the number of entity classes, and i is the number of features in a domain of discourse.

$$P_t = \sum_{i=1}^l \frac{o_i}{n} \quad (4)$$

The final weights p , f , and a of Equation 1 are functions of the probability of a type of feature with respect to the probability of the other two types of features (Equations 5a–c).

$$p = \frac{P_p}{(P_p + P_f + P_a)} \quad (5a)$$

$$f = \frac{P_f}{(P_p + P_f + P_a)} \quad (5b)$$

$$a = \frac{P_a}{(P_p + P_f + P_a)} \quad (5c)$$

A special case is when the maximum variability occurs, i.e., each distinguishing feature characterizes only one entity class. In such a case, P_p , P_f , and P_a are zero and the model assigns equal importance to parts, functions, and attributes. The same weights are also obtained when either an application domain has only one entity class or entity classes share all features. When there are no common features among the entity classes, the similarity values are zero, regardless the assignment of weights. Likewise, when features are shared by all entity classes, the similarity values are 1.0, independently of the assignment of weights.

4. Using Contextual Information with the Matching-Distance Model

To illustrate the integration of context into the matching-distance model, we analyze different scenarios for the similarity assessment. A prototype of the matching-distance model has been implemented in C++ and a knowledge base with around 60 entity class definitions was derived from the combination of WordNet [36] and the Spatial Data Transfer Standard [37]. Based on this knowledge base, four scenarios for the similarity assessment were selected:

- (1) Context-0: All distinguishing features are equally important, since context has not been considered in the similarity assessment.

7. Context-1: The user's intention is to retrieve recreational facilities. Figure 1 shows how the user describes her intention by using the operation *play* without arguments.

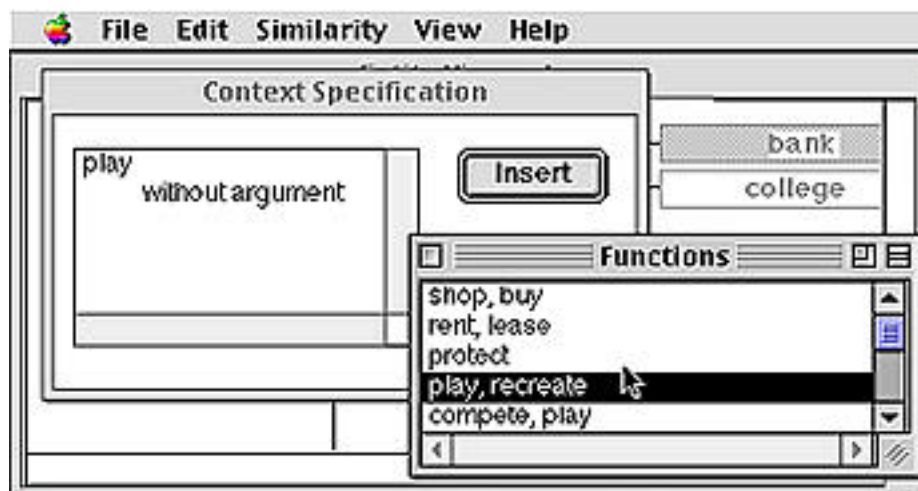


Figure 1: Context-1: Context specification with an operation without argument.

- (3) Context-2: The user's intention is to compare downtowns. The user describes her intention with the operation *compare* and the argument *downtown* (Figure 2).
- (4) Context-3: A user is searching for specific types of buildings and building complexes. This scenario is described by the operation *search* with specific entity classes (i.e., low level in the hierarchical structure) as the operation's argument (Figure 3).

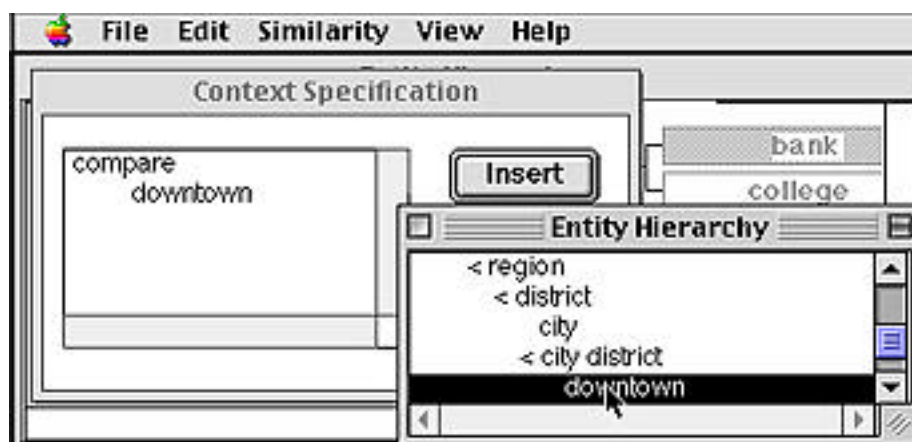


Figure 2: Context-2: Context specification with an operation and a general entity class (top level in the hierarchy) as an argument.

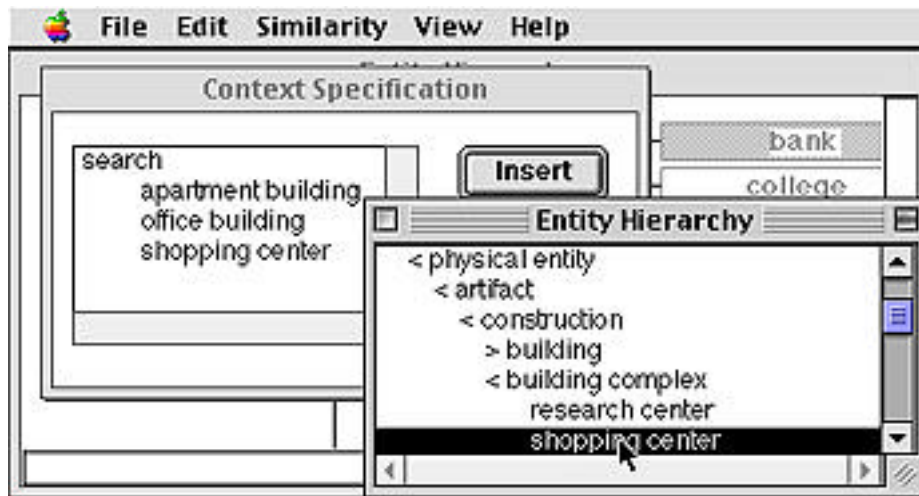


Figure 3: Context-3: Context specification with an operation and specific entity classes (low level in the hierarchy) as arguments.

As a result of the different context scenarios, different weights are generated for the similarity of parts, functions, and attributes. Besides the first scenario (Context-0), the other three scenarios define weights as a function of the variability of features within the application domain. The second scenario (Context-1) defines an application domain with entity classes that contain *play* as a function (*stadium*, *sports arena*, *athletic field*, and *park*). The third scenario (Context-2) results in a domain application with the entity class *downtown* and all entity classes semantically related to *downtown* by is-a or part-of relations (e.g., *artifact*, *construction*, *facility*, and *travelway*). The last scenario defines a small application domain that contains only the specific entity classes listed in the context specification (i.e., *apartment building*, *office building*, *shopping center*). Figure 4 shows an example of a partial view of the hierarchy with highlighted (dark lines) entities classes that belong to the application domain, the context specification window (for Context-1), and the window with the weights for each type of distinguishing features.

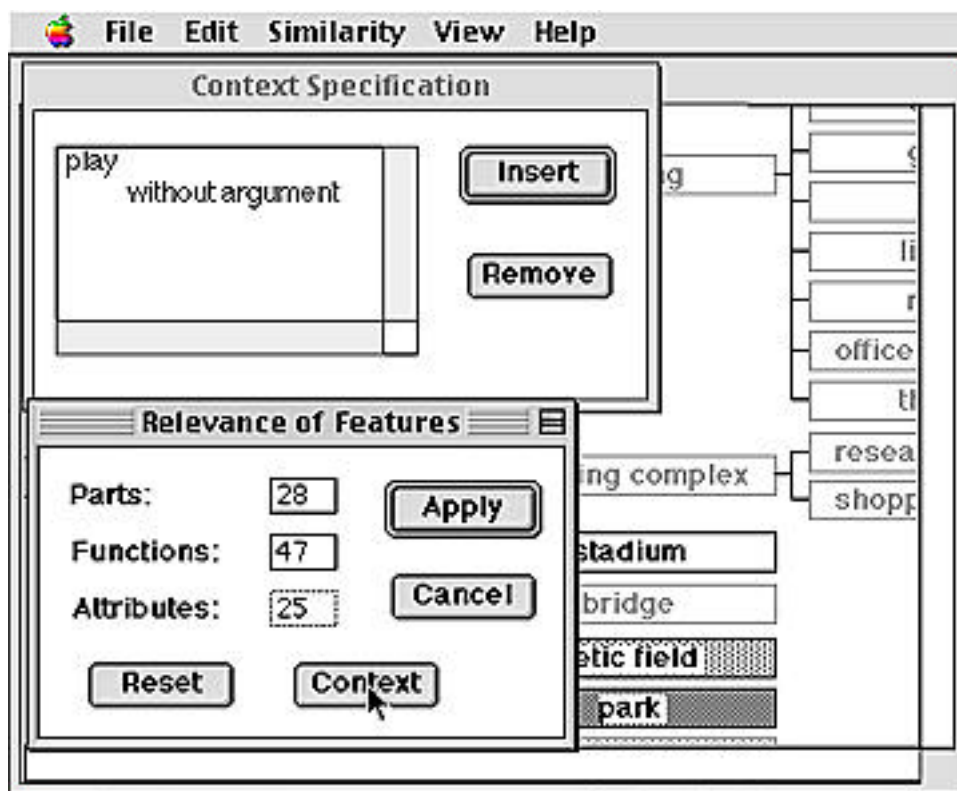


Figure 4: Partial view of the hierarchy with highlighted entity classes in the application domain.

Table 1 displays the sets of weights for parts, functions, and attributes that result from the definition of the five scenarios. While Context-1 highlights functions, Context-2 emphasizes attributes, and Context-3 gives about equal importance to parts and attributes.

Weights	Context-0	Context-1	Context-2	Context-3
Parts	33	23	32	38
Functions	33	51	23	22
Attributes	33	26	45	40

Table 1: Weights (%) for different specification of context.

Table 2 presents results of the similarity evaluation between a *stadium* and a subset of the entire knowledge base. Some of these entity classes may not be part of the application domain that results from the specification of each scenario.

Similarity (Stadium, X)				
X	Context-0	Context-1	Context-2	Context-3
Sports arena	0.63 (1)	0.69 (2)	0.62 (1)	0.59 (1)
Athletic field	0.60 (2)	0.72 (1)	0.60 (2)	0.54 (2)
Theater	0.33 (3)	0.24 (4)	0.41 (3)	0.38 (3)
Park	0.29 (4)	0.40 (3)	0.25 (5)	0.23 (5)
Museum	0.28 (5)	0.21 (5)	0.35 (4)	0.32 (4)

Table 2: Example of semantic similarity values for each of the four scenarios. Numbers in parenthesis refer to the similarity rank (1: best match, 5: worst match).

5. Discussion

Examples with different context specifications show that context has an influence on the determination of the relative importance of distinguishing features and, therefore, the similarity assessment. Context specification affects not only the similarity values, but also the relative location of entity classes in a ranking of similarity. The significance of the context specification may vary depending on the knowledge base to which the similarity model applies. If a knowledge base contains only definitions of entity classes that are part of the domain of the application, no context specification is necessary. Furthermore, the entity class definitions can also embed a context, since only the relevant distinguishing features for a particular application may have been listed. Therefore, context specification becomes more useful when a large and general-purpose knowledge base exists and a very detailed or specific application domain is desired.

The domain of an application reflects the underlying hierarchical structure of the knowledge base. For example, a context specification with an operation's arguments that are entity classes located at the upper level of the hierarchy (e.g., Context-2) leads to an application domain that contains a large number of entity classes (67% of our global knowledge base). Since the top levels of the hierarchy are usually less detailed (i.e., less distinguishing features are described), what matters for the variability of distinguishing features is the number of entity classes in the application domain that belong to the more detailed levels.

We performed an informal test with 21 human subjects, asking them to rank the six entity types with respect to a *stadium* under the four different contexts. People agreed strongly with the ranks of the similarity among entities for Context-0, confirming the best two matches and the least similar entity (Table 3).

Similarity (Stadium, X)				
X	Most similar	Second most similar	Third most similar	Least similar
Sports arena	16	3		1
Athletic field	4	13	2	
Theater		4	12	1
Museum			2	9
Park			5	4

Table 3: Human subjects' answers for similarity assessment with Context-1 (in number of people).

A similarly good match was obtained for Context-2. For Context-3, subjects agreed with the top and bottom ranks, but differed about the ranks of the second and third most similar cases. The responses for this context showed much more variation in the judgements and, therefore, small differences—often a single vote—influenced the rankings. One possible explanation for these less focused assessments may be that Context-3 led to a similarity assessment was done over entity classes that are outside of the application domain.

The strongest disagreement between the model and the human subjects was found for Context-1 (play). There was a clear tendency to make the functional characteristics of the entity class a relevant factor for the similarity evaluation. Although our model also emphasizes functional characteristics and the three most similar entity classes are the same as the three most similar entities given in people answers, the ranking among these three entity classes differs (Table 4).

Similarity (Stadium, X)

X	Most similar	Second most similar	Third most similar	Least similar
Sports arena	4	4	12	
Athletic field	8	12	1	
Theater				4
Museum				13
Park	9	5	7	

Table 4: Human subjects’ answers for similarity assessment with Context-1 (in number of people).

The main difference between the model and people’s judgment is associated with the ranking of *park*. The matching-distance model finds a stadium to be less similar to a park than a stadium to a theater within the context of comparing downtowns. One possible explanation is that these differences are due to the ambiguity of the term *park*. While the model considers the definition of a park as a piece of land in an urban area with recreational purpose, subjects may have chosen the different concept of a ballpark.

6. Conclusions and Future Work

Context has an important effect on the determination of semantically similar entity classes. Two entity classes may be more closely related to each other in one context than in another, since the importance of their distinguishing features varies with the context. In this paper we developed an approach to integrate context into a model for the assessment of semantic similarity among entity classes. This approach emphasizes the semantics derived from the use of entity classes and defines contextual information as the set of a user’s intended operations. A preliminary human-subject test over a set of four context scenarios showed that that model determined in three out of four cases the best match from a list of five terms, as well as the least similar term.

As future work, we will test and calibrate the matching-distance model by performing a formal human-subject test. Like entity classes, operations can be also semantically interrelated [38]. Thus, a further study should explore semantic relations among verbs to allow for a flexible and more accurate definition of entity classes that belong to the domain of the application.

A different approach to the determination of the relevance of distinguishing features is to analyze the degree of informativeness of a feature. Following the standard argument of information theory [35], the information content of a feature is defined as being inversely proportional to the probability of occurrence of this feature. While the degree of informativeness highlights variability, this work emphasizes commonalities. A combination of both approaches (i.e., variability and commonalities) could be explored to account for cases when the similarity evaluation is done

exclusively among the entity classes within the domain of discourse. In such a case, what characterizes the application domain becomes less relevant for distinguishing entity classes within the domain.

Although we have focused on the use of context for similarity assessment, the specification of the user's intended operations may play an important role for a wider range of processes in information systems. Instead of the user's intended operations, operations associated with the data modeled in a database have been suggested for solving problems of interoperability [32]. A description of the user's intended operation can create a bridge between user's expectation and information retrieved from individual or multiple databases.

7. Acknowledgments

Discussions with Werner Kuhn, Bob Rugg, and Gio Wiederhold helped with the development of the similarity model and the context specifications.

8. References

1. Tversky, A., 1977, Features of Similarity. *Psychological Review* 84(4): 327-352.
2. Voorhees, E., 1998, Using WordNet for Text Retrieval, in: C. Fellbaum (editor), *WordNet: An Electronic Lexical Database*, pp. 285-303, The MIT Press: Cambridge, MA.
3. Ginsberg, A., 1993. A Unified Approach to Automatic Indexing and Information Retrieval. *IEEE Expert* 8(5): 46-56.
4. Lee, J., M. Kim, and Y. Lee, 1993, Information Retrieval Based on Conceptual Distance in IS-A Hierarchies. *Journal of Documentation* 49(2): 188-207.
5. Bishr, Y., *Semantic Aspects of Interoperable GIS*. 1997, Wageningen Agricultural University and ITC, The Netherlands.
6. Bouguettaya, A., B. Benatallah, and A. Elmagarmid, Interconnecting Heterogeneous Information Systems. 1998, in: A. Elmagarmid (editor), *Advances in Database Systems*, Kluwer: Norwell, MA.
7. Rodríguez, A., M. Egenhofer, and R. Rugg, 1999, Assessing Semantic Similarity Among Geospatial Entity Class Definitions, in: A. Vckovski, K. Brassel, and H.-J. Schek, (editors), *Interoperating Geographic Information Systems INTEROP99, Zurich, Switzerland*, Lecture Notes in Computer Science 1580, pp. 189-202, Springer-Verlag: Berlin.
8. Krumhansl, C., 1978, Concerning the Applicability of Geometric Models to Similarity Data: The Interrelationship Between Similarity and Spatial Density. *Psychological Review* 85(5): 445-463.
9. Goldstone, R., D. Medin, and J. Halberstadt, 1997, Similarity in Context. *Memory and Cognition* 25(2): 237-255.
10. Miller, G., R. Reckwith, C. Fellbaum, D. Gross, and K. Miller, 1990, Introduction to WordNet: An On-Line Lexical Database. *International Journal of Lexicography* 3(4): 235-244.
11. Dahlgren, K., 1988, *Naive Semantics for Natural Language Understanding*, Kluwer: Norwell, MA.
12. Sheth, A., Data Semantics: What, Where, and How?, 1979, in: R. Meersman and L. Mark (editors), *Database Application Semantics*, Chapman and Hall.
13. Collins, A. and M. Quillian, 1969, Retrieval Time From Semantic Memory. *Journal of Verbal Learning and Verbal Behavior* 8: 240-247.
14. Smith, J. and D. Smith, 1977, Database Abstractions: Aggregation and Generalization. *ACM Transactions of Database Systems* 2(2): 105-133.
15. Winston, M., R. Chaffin, and D. Herrmann, 1987, A Taxonomy of Part-Whole Relations. *Cognitive Science* 11: 417-444.

16. Rosch, E. and C. Mervis, 1995, Family Resemblances: Studies in the Internal Structure of Categories. *Cognitive Psychology*, 7: 573-603.
17. Rosch, E., 1973, On the Internal Structure of Perceptual and Semantic Categories, in: T. Moore (editor), *Cognitive Development and the Acquisition of Language*, Academic Press: New York.
18. Leech, G., 1981, *Semantics: The Study of Meaning*. Penguin: Harmondsworth, UK.
19. McCarthy, J., 1987, Generality in Artificial Intelligence. *Communications of the ACM* 30(12): 1030-1035.
20. Aïmeur, E. and C. Frasson. 1995, Eliciting the Learning Context in Co-Operative Tutoring Systems. in: *Workshop on Modelling Context in Knowledge Representation and Reasoning*. Institute Blaise Pascal: Paris, France.
21. Kashyap, V. and A. Sheth, 1996, Schematic and Semantic Similarities between Database Objects: A Context-based Approach. *Very Large Database Journal* 5(4): 276-304.
22. Hearst, M., 1994, *Context and Structure in Automated Full-Text Information Access*. Computer Science Division, University of California at Berkeley.
23. Akman, V. and M. Surav, 1996, Steps Toward Formalizing Context. *AI Magazine* 17(3): 55-72.
24. Leacock, C. and M. Chodorow, 1998, Combining Local Context and WordNet Similarity for Word Sense Identification, in: C. Fellbaum (editor) *WordNet: An Electronic Lexical Database*. pp. 265-283, The MIT Press: Cambridge, MA.
25. Turner, R., 1998, Context-Mediated Behavior for Intelligent Agents. *International Journal of Human-Computer Studies* 48: 307-330.
26. Dojat, M. and F. Pachet. 1995, Three Compatible Mechanisms for Representing Medical Context Implicitly. in: *Workshop on Modelling Context in Knowledge Representation and Reasoning*, Institute Blaise Pascal: Paris, France.
27. Ouksel, A. and C. Naiman, 1994, Coordinating Context Building in Heterogeneous Information Systems. *Journal of Intelligent Information Systems* 3(1): 151-183.
28. Wiederhold, G. and J. Jannink. (in press), Composing Diverse Ontologies. in: *8th Working Conference on Database Semantics (DS-8)*. Rotorua, New Zealand, IFIP/Kluwer/Chapman & Hall.
29. Hayes, P., 1990, Naive Physics I: Ontology for Liquids, in: D. Weld and J. de Kleer (Editors) *Reading in Qualitative Reasoning about Physical Systems*, pp. 484-502, Morgan Kaufmann Publishers: San Mateo, CA.
30. Egenhofer, M. and D. Mark, 1995, *Naive Geography*. in: A. Frank and W. Kuhn (editors) *Spatial Information Theory A Theoretical Basis for Geographic Information Systems, International Conference COSIT'95*. Semmering, Austria, Lecture Notes in Computer Science 988, pp. 1-15, Berlin: Springer-Verlag.
31. Miller, G. and W. Charles, 1991, Contextual Correlates of Semantic Similarity. *Language and Cognitive Processes* 6(1): 1-28.
32. Kuhn, W. 1994, Defining Semantics for Spatial Data Transfers. in: T. Waugh and R. Healey (editors), *Sixth International Symposium on Spatial Data Handling*, pp. 973-987, Edinburgh, Scotland: International Geographical Union.
33. Fellbaum, C., 1990, English Verbs as a Semantic Net. *International Journal of Lexicography* 3(4): 270-301.
34. Gale, W., K. Church, and D. Yarowsky, 1992, A Method for Disambiguating Word Senses in a Large Corpus, *Computers and Humanities* 26(5/6): 415-450.
35. Ross, S., 1976, *A First Course in Probability*, Macmillan, New York.
36. Miller, G., 1990, Nouns in WordNet: A Lexical Inheritance System. *International Journal of Lexicography* 3(4): 245-264.
37. USGS, *View of the Spatial Data Transfer Standard (SDTS) Document*. 1998, available at <http://mcmcweb.er.usgs.gov/sdts/standard.html>. Last Modification: Friday, 12-Jun-98.
38. Fellbaum, C., 1998, A Semantic Network of English Verbs, in: C. Fellbaum (editor) *WordNet: An Electronic Lexical Database*, pp. 69-104, The MIT Press: Cambridge, MA.