Transitive Meronymy

Automatic Concept-Based Query Expansion Using Weighted Transitive Part-Whole Relations

Katrin Weller and Wolfgang G. Stock, Düsseldorf

In a theoretical view, we specify part-whole relations which are transitive. We will discuss the application of meronymy for automatic concept-based query expansion tasks in information retrieval. For practical reasons, we propose to specify meronymic relations and apply different weightings for query expansion purposes. For the construction of knowledge organization systems we point out that on a hyponymy concept ladder a term transmits all its parts (and all transitive parts of those) to its narrower terms.

Transitive Meronymie. Automatische begriffsbasierte Suchanfrageerweiterung unter Nutzung gewichteter transitiver Teil-Ganzes-Relationen.
Unsere theoretisch orientierte Arbeit isoliert transitive Teil-Ganzes-Beziehungen. Wir diskutieren den Einsatz der Meronymie bei der automatischen begriffsbasierten Suchanfrageerweiterung im Information Retrieval. Aus praktischen Gründen schlagen wir vor, die Bestandsrelationen zu spezifizieren und die einzelnen Arten mit unterschiedlichen Gewichtungswerten zu versehen, die im Retrieval genutzt werden. Für das Design von Wissensordnungen ist bedeutsam, dass innerhalb der Begriffsleiter einer Abstraktionsrelation ein Begriff alle seine Teile (sowie alle transitiven Teile der Teile) an seine Unterbegriffe vererbt.

1 Introduction

Ontologies as well as classification systems and thesauri consist of concepts and the paradigmatic semantic relations between them. Paradigmatic relations are document-independent relations used to model a domain of interest; in contrast to syntagmatic relations, which are merely based on the co-occurrence of concepts within single documents (Stock & Stock, 2008, pp. 68-70). Paradigmatic relations constitute the fundamental backbone of any controlled vocabulary. The "classical" semantic relations used in knowledge representation are synonymy, hierarchy and the unspecific association relation (Khoo & Na, 2006; Storey, 1993). Ontologies may define more specific relations. Still, the core semantic relation of every knowledge organization system is hierarchy. There are two kinds of hierarchic relations, which should be distinguished: hyponymy (is-a relation) (Cruse, 2002) and meronymy (part-of relation) (Winston, Chaffin, & Herrmann, 1987; Artale, Franconi, Guarino, & Pazzi, 1996; Iris, Litowitz, & Evens, 1989). Relations can furthermore obtain different properties, which allow drawing additional conclusions from the presence of certain relations. An important property of semantic relations is transitivity (others are mainly reflexivity or symmetry). Given x, y, z are concepts and ρ is a semantic relation, the relation ρ is transitive if

$$[(x\ \rho\ y) \wedge (y\ \rho\ z)] \rightarrow (x\ \rho\ z).$$

Understanding the nature of different semantic relations helps to create more accurate knowledge representations and to use them appropriately for information retrieval purposes. Knowledge about transitivity between concepts is needed to allow query expansion with concepts over more than one level in a semantic net (Stock, 2007, pp. 480-481). If transitivity is not given, this option is not per se available.
Transitivity of hyponymic relations is commonly not put into question. But how does meronymy behave? Consider two examples:

(1) Düsseldorf *is part of* North Rhine-Westphalia.
North Rhine-Westphalia *is part of* Germany.
If transitivity is given, we can infer:
Düsseldorf *is part of* Germany.

(2) Sonja's finger *is part of* Sonja.
Sonja *is part of* the Information Science Department.
If transitivity is given as well, we can infer: Sonja's finger *is part of* the Information Science Department
(Winston, Chaffin, & Herrmann, 1987).

In example (1) transitivity is given, in case (2) it is obviously not. In the way meronymy is currently applied, it cannot be regarded as generally transitive or generally intransitive. Our examples seem to show that there are some special part-whole cases, which are transitive, and some other, which are intransitive. But is meronymy really only *one* relation? Or do we misleadingly sum up different kinds of meronymy into one general part-whole relation? What would happen, if part-of-relations were separated into specific types of meronymic relations (Weller & Peters, 2007)? Is each of these meronymic sub-relations transitive? The transitivity of meronymy is a well known *problem* (Cruse, 1979), but there are still no satisfactory *solutions* for the practical use of transitive meronymy in information retrieval systems. Besides approaches in linguistics (Winston, Chaffin, & Herrmann, 1987; Cruse, 1979) and philosophy (Johansson, 2004; Varzi, 2006) there are some studies on partitive reasoning in the medical domain (Bernauer, 1996; Hahn, Schulz, & Romacker, 1999; Schulz, 2001; Schulz & Hahn, 2005; Schulz, Romacker, & Hahn, 1998), but to our knowledge the problem is not broadly discussed in information retrieval research. In retrieval research, our problem is localized in the area of concept- (or ontology-)based query expansion (Bhogal, Macfarlane, & Smith, 2007; Järvelin, Kekäläinen, & Niemi, 2001).
This paper is a theoretical approach in information science to clarify some problems concerning transitivity in meronymy in information retrieval and knowledge representation. It is arranged as follows: After reviewing basic problems of concept-based query expansion (section 2) we will discuss meronymic relations in section 3 and create a set of specified part-whole relations which are all transitive. In section 4 we show the benefits of transitive meronymy for the application of automatic semantic query expansion

in information retrieval tasks. Based upon this, we discuss whether semantic query expansion is in need of weighted semantic relationships (section 5) and whether meronyms are heritable downwards the concept ladder of hyponymy (section 6). All this is relevant for today's information retrieval scene (section 7).

2 Concept-based query expansion

Let us have a look at two hierarchical concept ladders! The first list is an example from the Medical Subject Headings (MeSH):

Diagnostic Techniques and Procedures
Diagnostic Techniques, Surgical
Endoscopy
Endoscopy, Digestive System
Endoscopy, Gastrointestional
Colonoscopy
Sigmoidoscopy.

This concept ladder is a pure hyponymy relation. A sigmoidoscopy *is a* colonoscopy, a colonoscopy *is a* gastrointestional endoscopy and so on. Assume a user searching for colonoscopy! If there are only few hits on his request, it should be possible to expand the query term with its narrower terms and – if there are still few hits – with its broader terms and with the other narrower terms of these concepts (the siblings of the start term). If the relations are transitive, there is no theoretical limit to expand the query. (The only limit is the size of the hit sets.) Based on the knowledge about hyponyms and hyperonyms the retrieval system will automatically create new queries, for example:

colonoscopy OR sigmoidoscopy
colonoscopy OR sigmoidoscopy OR "endoscopy, gastrointestional".

The second example is from the UNESCO thesaurus:

Social and human sciences
Human settlements and land use
Rural areas
Rural population
Rural women.

There is a lack of transitivity in this concept ladder. Rural women *are parts of* rural population and rural population *is part of* rural areas, but are rural women in deed *part of* rural areas (besides cattle, farm houses and soil)? Here we can expand a query on "rural women" with "rural population" but by no means with "rural areas".
There is no problem for concept-based query expansion to add terms from the neighboring levels (the direct narrower as well as the direct broader terms). But to add terms from further hierarchical levels automatically (i.e. without further enquiry on the user) is only allowed if transitivity is given. So the system can enhance a query for "endoscopy" with the term "sigmoidoscopy", but it would be a mistake to add "rural areas" to the search concept "rural women".
Some online information suppliers e.g., the German medical information provider DIMDI (Stock & Stock, 2003), offer a search option "CT DOWN" (controlled terms down), which takes all narrower terms (and the narrower terms of these terms and so on down to the bottom terms) into account. This is a very helpful option, but only if all hierarchical relations are transitive.

3 Meronymic relations

While hyponymy structures hierarchical concepts according to logical aspects, meronymy reflects a physical point of view (Khoo & Na, 2006, pp. 176). Concepts are subdivided according to their components; a hierarchical structure exists between a concept representing a wholeness (*holonym*) as upper class and concepts representing parts of it as lower classes (*meronyms*).
Meronymy is an inverse relation. So we have to define a relation on the way top-down in a concept ladder (whole – part: "has part" or "consists of") and the inverse relation on the way bottom – up (part – whole: "is part of").
To subdivide a concept into its components, only constitutive characteristics should be the determining factor. Yet, these constitutive characteristics may have different qualities, which produce different kinds of partitive relations respectively.

According to Gerstl and Pribbenow we may define two general classes of partitions: parts of structured and parts of non-structured objects (arbitrary parts) (Gerstl & Pribbenow, 1996; Pribbenow, 2002). Part-whole relations of structured objects refer to precisely given partitions inherent to objects. In contrast, non-structured objects do not provide fixed segmentations and are rather arbitrarily split up.

Figure 1 identifies a hierarchy of meronymic relations based on these two upper classes. Yet, defining the nature of a partitive relation is not always unproblematic. Thus, this classification should be subject to further discussions and future refinements.

Classical examples for structured parts (that do not pose problems) are geographic entities in form of geographic (administrative) units and their subunits, e.g. a city and its districts. The member-collection relation is intended for non-social

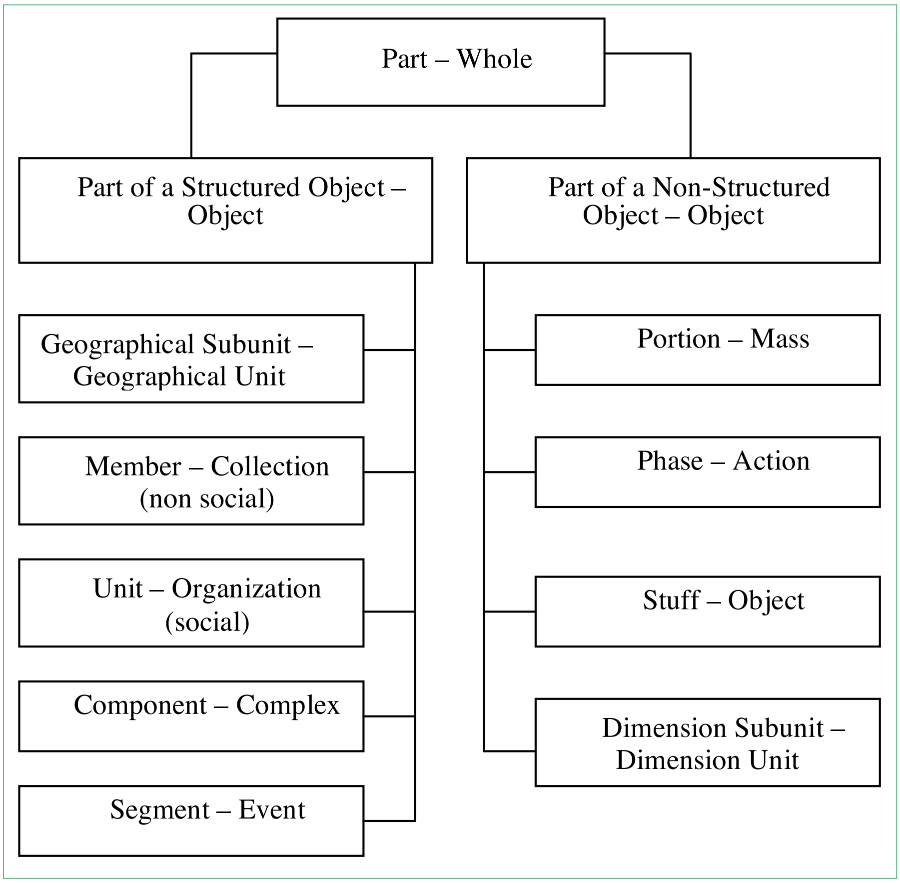

Figure 1. Hierarchy of meronymic relations.

collections only (e.g. tree – forest, ship – fleet) (Winston, Chaffin, & Herrmann, 1987) and is distinguished from the unit-organization relation, which comprises social groups, down to single persons belonging to an organized group (e.g. researcher – department – faculty – university). One very common and frequent relation is that between a complex and its components, like wheel – car, roof – house.

Sometimes actions and processes can also be subdivided, the borders between structured and non-structured segmentations in these cases are often blurred. We use the segment-event relation for describing parts of pre-structured processes (e.g. trapeze act – circus performance) and the phase-action relation for unstructured processes (e.g. paying – buying).

Partitions of non-structured objects often make use of external criteria, such as dimension units (e.g. millilitre – litre) or portions (e.g. slice – pie). Alternatively, internal features of an object may be considered for its segmentation (e.g. steel – bike) (Winston, Chaffin, & Herrmann, 1987; Gerstl & Pribbenow, 1996; Pribbenow, 2002).

Sometimes different dimensions are applicable for dividing a concept into its parts: for example a book may be considered as consisting of single pages and a cover or, if the content is regarded, of different chapters.

The different kinds of meronymic relations have to be specified during the ontology engineering process (Schulz, Romacker, & Hahn, 1998), which demands for careful consideration on the nature of existing relations. For practical reasons, ontology designers should not create too many different meronymic relations, for all those relations and connected concepts must be handled in a knowledge representation as well as in an information retrieval system. But they should create as many as necessary to secure that every single part-of relation is transitive, and that intransitivity which is merely based on inaccuracy is avoided.

Current ontologies do not sufficiently make use of the possibilities to define specific relations (sometimes even meronymy and hyponymy are not distinguished consequently). With the growing interest in ontology engineering this should change in near future. We propose to begin analyses on the types of relations in use as well as their frequencies. We expect that certain relations will be more frequent than others. Some relations are more likely to be used for a hierarchical chain with more than one level (as seen in the example for the unit-organization relation, from researcher up to university) than others (e.g. the portion-mass relation).

4 Query expansion using specified, transitive meronymy

Now that we have identified different types of part-whole relations, we can use them for more precise knowledge models. Accurate models may for example support enhanced query expansion systems. Let us go back to the problem of transitivity as described in the introduction. A query, not retrieving enough results, should be expanded with concepts meaningfully related to the original query terms. This is unproblematic for direct hyponyms, hyperonyms, holonyms and meronyms; if only one hierarchical level is considered, we can assume that a meaningful relation exists. If more than one hierarchical level should be included in the expanded query, meaningful and applicable extensions can only be guaranteed for transitive relations. As we have shown, unspecified meronymy cannot be stated as generally transitive or intransitive. But we now have the means for describing partitive relations more precisely, and thus to better identify cases in which transitivity is given.

In the case of the relation between Germany, North Rhine-Westphalia and Düsseldorf, we can see that there is the same kind of meronymy used for both levels: the geographical subunit relation (fig. 2). But in the example of Sonja, who is part of the Information Science Department and who has a finger as part of herself, we may now distinguish two different kinds of relations (fig. 3).

One first step to improve existing automatic query expansion techniques would thus be, to make use of relational chains with more than one hierarchical level, as long as the type of relation stays the same and transitivity is given.

Specified knowledge relations may also be of use for semi-automatic expansion systems. A user could be supplied with suggestions for expanding the original query terms that also include information about how the new terms are related to the given query (Stojanovic, 2005). This should help the user in rendering the query more precisely.

Additional improvements for automatic approaches can be achieved by applying weightings to the relations, i.e. to com-

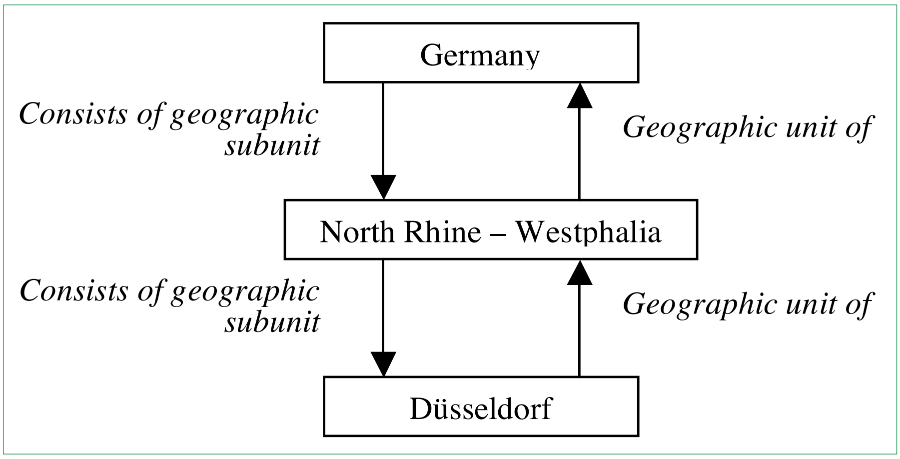

Figure 2. Example of transitive meronymy.

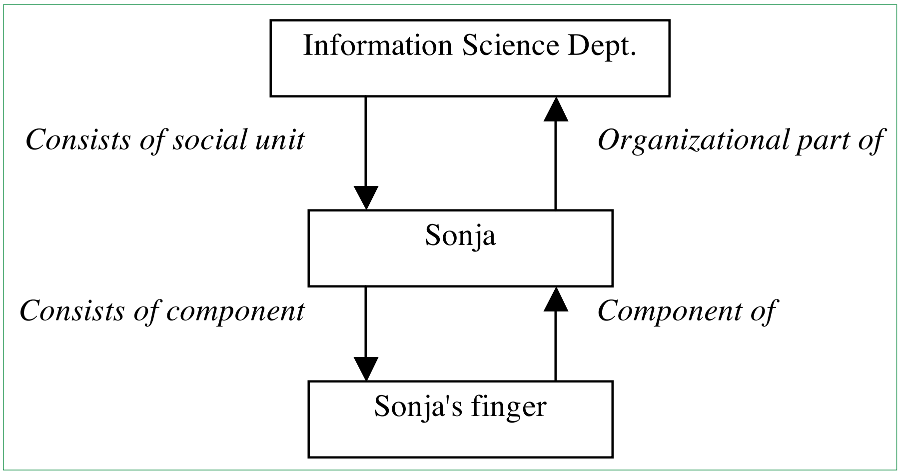

Figure 3. Example of intransitive meronymy, consisting of two different kinds of relations.

bine controlled vocabularies with quantitative calculations. In this case we do not only account whether a meaningful relation exists between two concepts or not, but we may also consider different qualities of relatedness.

5 Weighted meronymic relations

Working with (unweighted) transitive relations in the case of hyponymy is well known (Agrawal, Borgida, & Jagadish, 1989). To calculate the similarity of a query term with other concepts it is possible to count the edges on the shortest way from term A to term B in the given semantic network (Rada, Mili, Bicknel, & Blettner, 1989; Yang & Powers, 2005). For automatic query expansion mechanisms, additional weighting of relations may provide a more adequate way for meaningful extensions (Stock, 2007, p. 286). We may for example define a threshold and allow only concepts with values above it to be used for expanding a query.

Weightings may be assigned regarding (a) the distance between related terms, (b) the direction of the path for expanding (upwards vs. downwards), (c) the nature of the relation, (d) the combination of different relations.

Of course these different aspects may also be used in combination. More, as consequence of our previous explanations, distance-based weights should always be used in combination with more detailed information. The basis then would be to use weightings according to the (shortest possible) distances between related terms, e.g. a value of 1.0 for directly related concepts and gradually lower values for relations of higher degrees. Theoretically, these weightings might be applied automatically without regarding the nature of concepts or relations. But practically, at least transitivity has to be considered, so that more than one level of expansion is only allowed for transitive relations. To avoid intransitive cases, one might for example include only hyponymic relations and synonyms.

More precise values are achieved, by giving different weights to different types of relations. For example the retrieval software Convera works with the following default values (Bayer et al., 2005): Synonyms: 1.0 – Narrower terms: 0.8 – Broader terms: 0.5 – Related terms: 0.4. These values may be adjusted, new types of relations may be considered. We see, that also the directions of relations are already considered, as narrower terms are assigned with higher values than broader terms. For query expansion, a distinction between hierarchical relations to narrower terms and those to broader terms is reasonable.

What is still missing in this example is a differentiated view on hyponymy and meronymy and on the subtypes of partitive relations. We propose to assign slightly higher values for hyponymic relations than for meronymic ones.

For weighted meronymic relationships, we again encourage to use smaller values on the way upwards a concept ladder and higher values on the way down. For each transitive meronymy we then have to find different, appropriate values. For example, the stuff-object relation is rated to have minimal impact for query expansion mechanisms, the value should be set very low respectively (to include a query term "steel" when expanding a search for "bikes" does not promise a higher recall of relevant results, though the stuff object relation may be useful for query refinements and specification if to many results are retrieved; expanding a search for "steel" with "bike" the other way round is even less promising).

In contrast, the component-complex relation is of high importance and should be weighted with accordingly high values. In figure 4 this would mean, that relations going downwards along continuous lines (representing components of complexes) are rated with high values, going downwards these lines (complexes a component belongs to) is weighted less; the relations represented as dotted lines (representing stuff-object relations) would obtain only minimal (downwards) or null values (upwards).

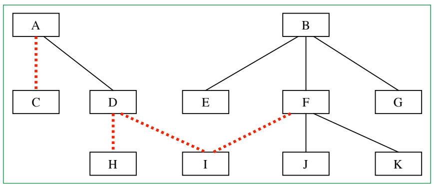

Figure 4. Concept system with component-complex-relation (continuous lines) and stuff-object-relation (dotted lines).

Finally, we may also take into account how different types of relations are used in combination. This may be of use for specific constructions that are not transitive (due to different relational types in use) but may still provide meaningful relations for expanding a query.

6 Heritable parts

Meronymy and hyponymy interact in the area of feature inheritance. Hyponyms inherit properties from their hyperonyms, this also holds for properties expressed as partitive relations (Horrocks & Sattler, 1999). This is an important aspect for the accurate construction of knowledge representation models. When creating hyponymies, it has to be considered, that concepts pass on their related meronyms to their narrower terms (Miller, 1998, p. 38; Stock, 2007, p. 279).

In figure 5 we have a hyponymic relation between A, B and C (B is narrower term of A; C is narrower term of B), also related to A is a meronymic relation chain, i.e. P1 is part of A (P2 is part of P1 and P3 is part of P2, all transitive). In this case the properties has_part P1 to has_part P3 are passed on from A to B and C. If we consider for example the hyponymic structure of vertebrates (A) and their constitutional part, the vertebral column (P1): each hyperonym (e.g. from gnathostomata (B) down to cats (skipping C and going further down)) inherits the vertebral column and also its components, e.g. the coccyx (P2) and single vertebrae (P3) as parts.

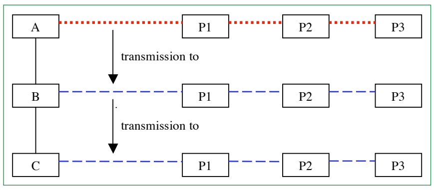

Figure 5. Heritable parts in the hyponymy relation. Hyponymy: continuous lines, transitive meronymy: dotted lines.

7 Application Fields

Knowledge relations play different roles for information retrieval applications. As an essential part of elaborated knowledge representation systems or controlled vocabularies, they can be part of the query formulation as well as the document indexing process.

Professional databases (e.g. Medline, INSPEC, CAS) are usually domain specific and thus make use of specialized controlled vocabularies. These vocabularies are used to create document surrogates for retrieval purposes and are also utilized by the system's users to express their search queries. These applications rely heavily on the quality of the vocabularies in use. The performance of search results also depends on the accuracy in indexing and on the user's ability to express his information needs. Mechanisms to broaden a query to retrieve more than the direct matches for a concept can be useful.

Web search engines normally do use knowledge representation methods neither for indexing nor for query formulation, in order to reduce production costs and to provide maximum convenience for web-users. Yet they may apply automatic query expansion mechanisms based on controlled vocabularies for improved results. This may also be applied to document collections indexed with social tagging (folksonomies) (Kolbitsch, 2007).

Besides these two most common application fields, also less well-established approaches are of interest: For example, concept-based indexing systems (e.g.

Convera) assign index terms to a document based on word occurrences in texts and their semantic environment. A usual case for automatic indexing would be that a term A appears in a document and is also declared as descriptor in the underlying controlled vocabulary, with the result that A is used to index the document. More complex variations also consider knowledge relations; if term A1 and A2 appear in a text and are both listed as narrower terms of A, we may also add A to the documents surrogate. Again, additional weightings would help to decide, whether bindings are strong enough to include certain related terms.

8 Conclusions and Future Work

Specified semantic relations are one means to improve knowledge representation systems, which can be applied for query (re-)formulation and document indexing in information retrieval.
We discussed the specification of partitive relations as a way to provide transitivity for meronymy. Kinds of meronymy have to be made explicit in practical applications. The subtypes suggested here will have to be realized and evaluated, and probably modified respectively. In a next step, different weightings may be determined and applied for weighted query expansion mechanisms.
In this paper we focused particularly on meronymy as one major semantic relation. For practical applications in information retrieval other types of relationships will have to be considered accordingly. Particularly, the associative relations will have to be analyzed and specified, respectively. Furthermore, it is necessary to implement retrieval systems with semantic relations and with query expansion tools based on these relations. Thus, measures of retrieval system quality can be conducted: Does the application of transitive meronyms really end in higher recall or in higher precision?

References

Agrawal, R., *Borgida*, A., & *Jagadish*, H.V. (1989). Efficient management of transitive relationships in large data and knowledge bases. In: Proceedings of the 1989 ACM SIGMOD International Conference on Management of Data (pp. 253-262). New York: ACM.

Artale, A., *Franconi*, E., *Guarino*, N., & *Pazzi*, L. (1996). Part-whole relations in object-centered systems: An overview. Data & Knowledge Engineering 20(1996)3, 347-383.

Bayer, O., *Höhfeld*, S., *Josbächer*, F., *Kimm*, N., *Kradepohl*, I., *Kwiatkowski*, M., *Puschmann*, C., *Sabbagh*, M., *Werner*, N., & *Vollmer*, U. (2005). Evaluation of an ontology-based knowledge-management-system. A case study of Convera RetrievalWare 8.0. Information Services & Use 25(2005), 181-195.

Bernauer, J. (1996). Analysis of part-whole relation and subsumption in the medical domain. Data & Knowledge Engineering 20(1996)3, 405-415.

Bhogal, J., *Macfarlane*, A., & *Smith*, P. (2007). A review of ontology-based query expansion. Information Processing & Management 43(2007)4, 866-886.

Cruse, D.A. (1979). On the transitivity of the part-whole-relation. Journal of Linguistics 15(1979)1, 29-38.

Cruse, D.A. (2002). Hyponymy and its varieties. In: Green, R., Bean, C.A., & Myaeng, S.H. (eds.): The Semantics of Relationships (pp. 3-21). Dordrecht: Kluwer.

Gerstl, P., & Pribbenow, S. (1996). A conceptual theory of part-whole relations and its applications. Data & Knowledge Engineering 20(1996)3, 305-322.

Hahn, U., *Schulz*, S., & *Romacke*r, M. (1999). Part-whole reasoning. A case study in medical ontology engineering. IEEE Intelligent Systems & Their Applications 14(1999)5, 59-67.

Horrocks, I., & *Sattle*r, U. (1999). A description logic with transitive and inverse roles and role hierarchies. Journal of Logic and Computation 9(1999)3, 385-410.

Iris, M.A., *Litowitz, B.E., & Evens*, M. (1989). Problems of the part-whole-relation. In: Evens, M. (ed.): Relational Models of the Lexicon (pp. 261-288). New York: Cambridge Univ. Pr.

Järvelin, K., *Kekäläinen*, J., & *Niemi*, T. (2001). ExpansionTool: Concept-based query expansion and construction. Information Retrieval 3(2001), 231-255.

Johansson, I. (2004). On the transitivity of the parthood relations. In: Hochberg, H., & Mulligan, K. (eds.): Relations and Predicates (pp. 161-181). Frankfurt: Ontos.

Khoo, C.S.G., & *Na, J.C.* (2006). Semantic relations in information science. In: Annual Review of Information Science and Technology 40(2006), 157-228.

Kolbitsch, J. (2007). WordFlickr. A solution to the vocabulary problem in social tagging systems. In: Proceedings of I-MEDIA '07 and I-SEMANTICS '07, Graz, Austria, September 5-7, 2007 (pp. 77-84).

Miller, G.A. (1998). Nouns in Wordnet. In Fellbaum, C. (ed.): WordNet. An Electronic Lexical Database (pp. 23-46). Cambridge, Mass., London: MIT Press.

Pribbenow, S. (2002). Meronymic relationships: From classical mereology to complex part-whole relations. In: Green, R., Bean, C.A., & Myaeng, S.H. (eds.): The Semantics of Relationships (pp. 35-50). Dordrecht: Kluwer.

Rada, R., *Mili*, H., *Bicknel, E.* & *Blettner*, M. (1989). Development and application of a metric on semantic nets. IEEE Transactions on Systems, Man and Cybernetics 19(1989)1, 17-30.

Schulz, S. (2001). Bidirectional mereological reasoning in anatomical knowledge bases. In Proceedings AMIA Symp. (pp. 607-611).

Schulz, S., & *Hahn*, U. (2005). Part-whole representation and reasoning in formal biomedical ontologies. Artificial Intelligence in Medicine 34(2005)3, 179-200.

Schulz, S., *Romacker*, M., & *Hahn*, U. (1998). Part-whole reasoning in medical knowledge bases using description logics. Lecture Notes in Computer Science 1504(1998), 237-248.

Stock, M., & *Stock*, W.G. (2003). DIMDI: Eine deutsche Informationsbehörde für Gesundheit und Medizin. Password (2003)9, 18-27.

Stock, W.G. (2007). Information Retrieval. Informationen suchen und finden. München, Wien: Oldenbourg.

Stock, W.G., & *Stock*, M. (2008). Wissensrepräsentation. Informationen auswerten und bereitstellen. München, Wien: Oldenbourg.

Stojanovic, N. (2005). On the query refinement in the ontology-based searching for information. Information Systems 30(2005), 543-563.

Storey, V.C. (1993). Understanding semantic relationships. VLDB Journal 2(1993), 455-488.

Varzi, A.C. (2006). A note on the transitivity of parthood. Applied Ontology 1(2006)2, 141-146.

Weller, K., & *Peters*, I. (2007). Reconsidering relationships for knowledge representation. In: Proceedings of I-KNOW '07. Graz, Austria, September 5-7, 2007 (pp. 493-496).

Winston, M.E., *Chaffin*, R., & *Herrmann*, D. (1987). A taxonomy of part-whole relations. Cognitive Science (1987)11, 417-444.

Yang, D., & *Powers*, D.M.W. (2005). Measuring semantic similarity in the taxonomy of WordNet. In: Proceedings of the 28th Australasian Conference on Computer Science, Vol. 38 (pp. 315-322). Darlinghurst: Australian Computer Society.

Meronymy, semantic relation, part-whole-relation, transitivity, query expansion, weighted relation.
Gewichtung, Relation, Recherchestrategie, Information Retrieval, Meronymie

The Authors

Katrin Weller, M.A.

is researcher at the Heinrich-Heine-University, Düsseldorf, and works in the Department of Information Science. Currently she is involved in the research project Ontoverse, where she is concerned with Ontologies as new forms of knowledge organization systems, knowledge engineering processes, and collaborative ontology management.
weller@uni-duesseldorf.de

Univ.-Prof. Dr. Wolfgang G. Stock

is the head of the Information Science Department of the Heinrich-Heine-University Düsseldorf (Germany).
stock@phil-fak.uni-duesseldorf.de