Matching with a Hierarchical Ontology

Yaacov Choueka, Nachum Dershowitz, and Liad Tal

No Institute Given

Abstract. We describe how hierarchical ontologies are used for classifying products, as well as for answering queries. For classifying and scoring product descriptions and queries in an electronic-commerce search engine, we make use of declarative hypernym/hyponym and meronym/holonym hierarchies, such as the sense hierarchies provided by WordNet. Retrieval is also aided by use of the ontology. Results of initial experiments are encouraging.

I now see the hierarchical as both beautiful and necessary.
—Camille Paglia (1992)

1 Introduction

Consider a search engine geared to electronic commerce. Each user query retrieves a ranked list of matching products and a dialogue is initiated to refine the search results down to a short list of quality matches. For example, a search for "black bags" might return hundreds of products. Since there are many different types of bags, in addition to displaying a ranked list of black bags, the system allows an online customer to choose from various categories, such as handbags, briefcases, doctor bags, or backpacks. Products in the store's catalogue are classified (either on-the-fly or in a pre-processing stage) to each of the possible categories and subcategories. See Fig. 1.

We describe in the following sections how hierarchical ontologies are used for classifying products, as well as for answering queries.

2 Hierarchical Ontologies

For classifying and scoring product descriptions, as well as queries, we make use of declarative hypernym/hyponym (*is-a*) and meronym/holonym (*part-of*) hierarchies, such as the sense hierarchies provided by WordNet [3]. We assume these ontologies are organized as directed acyclic graphs. Each node—representing a class of objects (commodities and commodity categories, in our case)—may be related by *is-a* and *part-of* edges to other nodes. For example: a rucksack *is-a* kind of a backpack, which *is-a* kind of bag (in the sense of container), jet black *is-a* shade of black, and a clasp is a *part-of* a handbag, as is a handle.

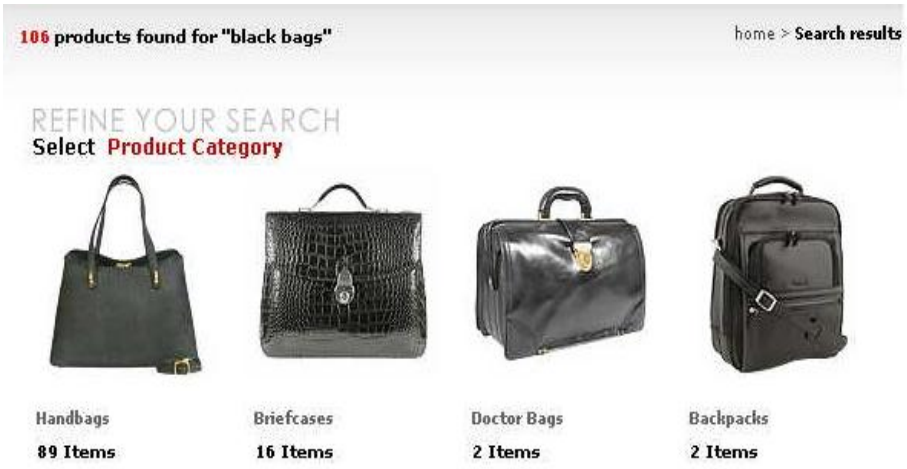

Fig. 1. Sample search dialogue.

3 Classification

Making use of a hyponym taxonomy has long been a tool in hierarchical classification. See, for example, [4, 2, 5].

There are several novel ways in which an ontology can help:

1. Consider rival classifications, B (e.g. bag) and C (container), of some item, be they derived by linguistic or statistical means.
   (a) If C is a hypernym of B, then the more specific classification B is preferred.
   (b) If neither is a hypernym of the other (shirt and cardigan, say), then their closest common ancestor (*least upper bound*) A (garment) is a better classification than either.
   (c) If several sibling meronyms are competing, this suggests their shared holonym, especially when the latter is also a contender.
2. To decide between rival meronym and holonym (lens and camera), additional criteria are called for (e.g. the likelihood that a camera lens is sold alone).
3. Even weak indications of relevance to a hypernym/hyponym/meronym/holonym can help in word sense disambiguation.

For example, consider a product labelled "Zombie Costume," and described as a "Mask from the Buffy the Vampire Slayer collection of masks. From 'Movie Originals' at The Fright Catalog - Halloween 2001." Linguisitic analysis suggests it may be a *costume* or a *mask*. Since *mask* is a hyponym of *costume*, the former, more specific, category is preferred. On the other hand, *mask* may also be a meronym of *costume*, but since only one part of a *costume* is mentioned in the description, the category *mask* wins out.

4 Retrieval

It is standard to organize retrieved items according to a predetermined hierarchy, as illustrated in the example in Fig. 1. See [1].

There are several additional ways in which some search engines take advantage of hierarchical ontologies. These include:

1. Retrieve all hyponyms (e.g. rucksacks and etuis) of query terms (bag), but not their hypernyms.
2. If the query ("orange leather briefcase") is for a subcategory in which there are too few elements, one may wish to suggest siblings (orange leather backpacks; red leather briefcases) or cousins from the hypernym hierarchy.
3. In some contexts, a query for a category might also propose its meronym classes (lenses for "cameras") as also of relevance.

If hypernym/hyponym and meronym/holonym links have meaningful weights attached, they can be used to assign probabilities to matches and improve the quality of the ranking of retrieved items. For example, links from W (e.g. watch) to its hypernyms I (instrument) and J (jewelry) may be weighted by the likelihood that an item classified as W is an I or a J, respectively; the reverse, hyponym link, from J to W, may be qualified by the chance that someone who asked for J will choose an item from W. Similarly, one may label the link from meronym C (clasps) to holonym H (handbags) by the likelihood that a C is part of an H and, in the reverse direction, that an H has part C.

5 Results

Preliminary experiments with many thousands of product descriptions are encouraging. Just using a shallow hypernym/hyponym hierarchy (1a, above), for example, improves classification for about 1% of the products, and almost never harms.

Acknowledgements.

We thank Michael Flor for his large share in this effort and Yael Villa for her comments and encouragement.

References

1. Hao Chen and Susan Dumais. Bringing order to the Web: Automatically categorizing search results. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 145–152, The Hague, The Netherlands, April 2000. ACM Press.
2. Susan T. Dumais and Hao Chen. Hierarchical classification of Web content. In Nicholas J. Belkin, Peter Ingwersen, and Mun-Kew Leong, editors, *Proceedings of SIGIR-00, 23rd ACM International Conference on Research and Development in Information Retrieval*, pages 256–263, Athens, GR, July 2000. ACM Press.

3. Christiane Fellbaum, editor. *WordNet: An Electronic Lexical Database*. MIT Press, Cambridge, MA, 1998.
4. Sam Scott and Stan Matwin. Text classification using WordNet hypernyms. In *Proceedings of the COLING/ACL Workshop on Usage of WordNet in Natural Language Processing Systems*, pages 45–52, Montréal, Canada, August 1998.
5. Aixin Sun, Ee-Peng Lim, and Wee-Keong Ng. Hierarchical text classification methods and their specification. In Alvin T. S. Chan, Stephen C. F. Chan, H. V. Leong, Vincent, and T. Y. Ng, editors, *Cooperative Internet Computing*, pages 236–256. Kluwer Academic, Boston, March 2003.