

Towards Evaluating the Impact of Ontologies on the Quality of a Digital Library Alerting System

Alfons Huhn, Peter Höfner, Werner Kießling

Angaben zur Veröffentlichung / Publication details:

Huhn, Alfons, Peter Höfner, and Werner Kießling. 2005. "Towards Evaluating the Impact of Ontologies on the Quality of a Digital Library Alerting System." Augsburg: Universität Augsburg.

Nutzungsbedingungen / Terms of use:

licgercopyright

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under these conditions:


Deutsches Urheberrecht

Weitere Informationen finden Sie unter: / For more information see:

<https://www.uni-augsburg.de/de/organisation/bibliothek/publizieren-zitieren-archivieren/publiz/>



UNIVERSITÄT AUGSBURG




**Towards Evaluating the Impact of
Ontologies on the Quality of a Digital
Library Alerting System**

A. Huhn P. Höfner W. Kießling

Report 2005-7

März 2005



INSTITUT FÜR INFORMATIK

D-86135 AUGSBURG

Towards Evaluating the Impact of Ontologies on the Quality of a Digital Library Alerting System

Alfons Huhn, Peter Höfner, Werner Kießling

Institut für Informatik, Universität Augsburg,
86159 Augsburg, Germany
{huhn, hoefner, kiessling}@informatik.uni-augsburg.de

Abstract. Advanced personalization techniques are required to cope with novel challenges posed by attribute-rich digital libraries. At the heart of our deeply personalized alerting system is one extensible preference model that serves all purposes. In this paper we focus on ontology and quality assessment in conjunction with our search technology Preference XPath and XML-based semantic annotations of digital library multimedia objects. We evaluate the impacts of automatic query expansion by ontologies by embedding our alerting system P-News as a black box or a glass box in a test lab. It changes configuration parameters on its own, feeds test cases to P-News, compares the results of different configurations, and stores the result set for further evaluations. The most important indications of this work in progress are: The use of ontologies improves the quality of the result set, generates further results of higher quality, and implies the use of knowledge to reduce a loss of focus.

1 Introduction

Alerting subscribed users about arrivals of new information, articles, or books is a challenging task for every digital library (DL), even with a growing commercial impact as exemplified by *Springer Alerts*. Our experimental P-News system [16] is tailored to alert users, when user-relevant documents newly arrive at the DL. Its main feature is a deep personalization which is achieved by a highly flexible preference methodology with powerful query capabilities. Let's delineate the scope of P-News in a scenario taken from [16]:

Example 1: Cathy is a professor at university and besides research projects also manages a spin-off business. Of course she wants to know about news related to her research and is interested in specific business news. Assume a suitable document arrives at P-News, e.g. a new research article. The alerting process then has to recognize from the representation of Cathy's topical preferences that the quality of a new article justifies a notification.

This example manifests a relation between the users' wishes formulated as preferences and the delivery of the result set. This relation has to be evaluated with respect to the components of ontology [15] and quality assessment.

There are different evaluation approaches like economical, sociological, user-centric, and system-centric [4, 14]. User-centric approaches can be used to get feedback by the customers about the relevancy of the results. System-centric approaches rely on the underlying software system, and evaluate the effectiveness of systems such as content-oriented XML retrieval systems reported in [7]. The main problem of a system-centric approach is the tedious generation of queries as well as their respective relevance assessments. Thus instead of applying a quantitative system-centric approach with relevance measures as precision and recall, we want to discuss the impacts of ontology-based query expansion at a qualitative level, i.e., we are using a test lab which embeds our alerting system in the usual manner of a black box as well as in the unusual manner of a glass box. Since work is still in progress, we postpone the quantitative approach until we will get the input of everyday customers in our next project to have a sound statistical foundation.

The paper is organized as follows: In section 2 we revisit the preference concepts of P-News. XML-based annotations are the key to access just the interesting parts of DL objects. We present the benefits of our search technology ‘Preference XPath’ and especially of the expansion by ontology. Section 3 follows the evaluation scheme of [14] and describes how the P-News system is framed by a test lab to get a proper stimulus-reaction-scheme for a reproducible test suite. Via the test lab we pinpoint those sample cases demonstrating the impacts of ontology. The related challenges are discussed in section 4 in detail. Section 5 contains a summary and outlook.

2 Preference Concepts in P-News

In the sequel we want to give a short survey about our alerting system P-News and its components such as quality assessment, use of ontology, as well as Preference XPath, and XML-based metadata.

2.1 Shifting towards an Attribute-rich Search in DLs

Nowadays DLs contain compound documents of text, images, audio or even video files. Obviously without any adequate annotation of these documents, no intelligent search-engine can be realized. There exist multiple approaches to describe and standardize annotations for multimedia documents. We use the multimedia content description interface MPEG-7 to annotate our test database of digital videos [16]. This XML-based standard provides a hierarchical data organization with lots of attributes to describe the meta-information of digital documents. Each video segment is described by many different aspects, like *CreationInformation* (e.g. title, creator), structural aspects (e.g. segments and subsegments) and conceptional aspects (e.g. text annotation, semantics).

Additionally pragmatic information, like *MediaInformation* (e.g. storage format, visual coding) and *UsageInformation* (e.g. access, rights) are also expressed by attributes. With this standard the focus in searching DL or evaluating a document’s relevance shifts from pure text-matching to a more semantically rich search on typed attributes.

2.2 Personalized Search Technology: Preference XPath

MPEG-7 descriptions are XML data containing predefined attributes for annotating digital documents. So we need some technology for searching and filtering XML data. Query languages for XML like XPath or XQuery are used to formulate precise queries. Thus these languages only provide the syntax and the semantic for *hard condition* queries. But standard query languages cannot handle *soft conditions*, as ‘File size should be around 300 MBytes’ or semantics as ‘I like A more than B’.

In [1, 9] an approach for preference modeling is proposed utilizing strict partial orders. This approach was recently extended in [10]. Using this formalism, XPath is augmented by preferences [9, 12]. Preference XPath can handle queries supporting soft conditions as well as hard conditions. We deliver all *best matching* objects (BMO), but not necessarily exact matches. As an essential feature a set of predefined preference constructors is used for modeling users’ wishes and disfavours. For example, *AROUND* and *BETWEEN* are base preference constructors on numerical attributes. On the other hand *CONTAINS* and *POS* are base preference constructors for non-numerical values. The first constructor describes a subset-relation and the second provides an opportunity to model sets of desirable values. For inductively combining preferences several complex constructors are available in particular the *PARETO* preference denoted as \otimes . In addition the *nested* preference was developed in [16] to handle tree-structured XML data. The importance of Preference XPath for personalization issues is also shown in a practical evaluation [2]. A more detailed list of all available preferences together with their definitions can be found in [9, 10].

2.3 Quality of Preference Query Results

If we cannot deliver an exact match for a query, the BMO query model of Preference XPath automatically retrieves best matching objects as alternatives. Thus a quality assessment for such alternatives is essentially needed. We are using an intuitive linguistic model [3, 11] for the quality of query results. To each result tuple of a preference query under the BMO-model one out of five linguistic quality terms can be assigned. We use *perfect*, *very good*, *good*, *acceptable*, and *sufficient*. These quality terms model different degrees of relevance. Note that in contrast each result tuple of a database query under the conventional exact-match query model has quality *perfect*, whereas tuples not in the query result are rated as *insufficient*.

Example 2: How to calculate the quality for any preference-based query? What is the quality if e.g. the preferred video size should be around 300 MBytes and the offered size is 350?

The answer depends on several factors of the respective situation, for example on the usual range of the video size in an application. As shown in [3, 11] a quality function $QUAL_{P,s}$ can always be generated, for every base preference P depending on a situation s. These functions have to preserve the following fundamental fairness principle: Given a preference $P = (A, <_P)$ and an attribute A with a strict partial order $<_P$ on its domain $dom(A)$, then:

$$x <_P y \text{ implies } QUAL_{P,s}(x) \leq QUAL_{P,s}(y) \text{ for every situation } s.$$

In the above example we might have:

$$\text{QUAL}_{\text{AROUND}(\text{F},300),s}(\text{t}) := \begin{cases} \text{perfect} & , \quad \text{t[F]} = 300 \\ \text{very good} & , \quad 280 \leq \text{t[F]} < 300 \vee 300 < \text{t[F]} \leq 320 \\ \text{good} & , \quad 260 \leq \text{t[F]} < 280 \vee 320 < \text{t[F]} \leq 340 \\ \text{acceptable} & , \quad 240 \leq \text{t[F]} < 260 \vee 340 < \text{t[F]} \leq 360 \\ \text{sufficient} & , \quad \text{t[F]} < 240 \vee 360 < \text{t[F]} \end{cases}$$

Here F is an abbreviation of the attribute ‘FileSize’ and t is a result tuple of the BMO-set. The situations involve situated preferences. A meta model of situation and the interaction with the associated preference repository is described in [5]. Later on we refer to the following quality function for an attribute A of the POS preference:

$$\text{QUAL}_{\text{POS}(\text{A},\text{POS-set}),s}(\text{t}) := \begin{cases} \text{perfect} & , \quad \text{t[A]} \in \text{POS-set} \\ \text{sufficient} & , \quad \text{t[A]} \notin \text{POS-set} \end{cases}$$

After calculating the quality of the involved base preferences, we define the qualities of the complex preferences by combining the qualities of the base preferences. Depending on the strategy, the quality of a Pareto preference $P_1 \otimes \dots \otimes P_n$ can be calculated as in the following examples:

- Optimistic valuation: $\text{QUAL}_{P,s}(\text{t}) := \max\{\{\text{QUAL}_{P_j,s}(\text{t}) \mid j = 1, \dots, n\}\}$
- Pessimistic valuation: $\text{QUAL}_{P,s}(\text{t}) := \min\{\{\text{QUAL}_{P_j,s}(\text{t}) \mid j = 1, \dots, n\}\}$

Furthermore, equidistant linguistic average valuation and median quality valuation exist. For details, examples, and more functions calculating the quality of complex preferences, the reader is referred to [3].

We use the quality assessment on the one hand to give the user a feedback of the quality of the delivered results. On the other hand the user may express a *notification restriction* to be alerted only by those results which have a quality better than his personalized notification threshold.

2.4 Use of Context by Ontologies

Ontologies are mostly modeled by XML-based formats as MPEG-7, DAML+OIL and Web Ontology Language, which is included in the Semantic Web. In ontologies any agent can find easily *synonyms*, *hyponyms* (subordinates), *hypernym* (superordinates) as well as complex relationships. Describing the ontology by a directed graph, each concept together with their associated synonyms are identified by nodes. Hypernyms describe the ancestors (especially parents are hypernyms) and hyponyms the descendants. A keyword-driven search is used in the selected ontology (see Fig. 1) in comparison to the attribute-rich search for DL objects.

Example 3: If a user is interested in ‘invasive software’, a conventional exact-match search engine may deliver an empty result set, because many digital documents are annotated in more detail, like perhaps ‘viruses’ or ‘Trojan horses’, i.e., all tuples of the database are considered as *insufficient*. On the other hand, we could formulate the user’s interests as a preference using the POS-constructor:

POS (title, { ‘invasive software’ }).

Since no perfect match exists, all tuples would be now rated as *sufficient*. However, this might be overly pessimistic in some situations.

Certainly, the user is also interested in more specialized topics. We experimentally exploit the expansion by ontology by the following preference order:

siblings $<_p$ parents $<_p$ descendants $<_p$ synonyms $<_p$ self,

where ‘a $<_p$ b’ describes the preference relation.

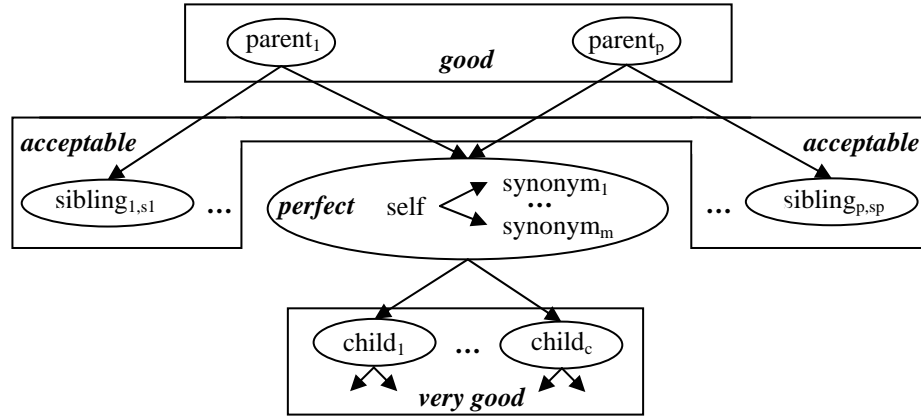


Fig. 1. Quality valuation of expansion by ontology

Corresponding to the preference order, the quality assessment assigns *perfect* to self and synonyms, and so on, till *acceptable* is assigned to the siblings. This mapping also preserves the fundamental fairness principle.

Example 3 (cont.): The original term ‘invasive software’ corresponds to self as well as ‘viruses’ or ‘Trojan horses’ map to child₁ and child₂ in the above ontology. The search expression
 POS(title, { ‘invasive software’ })
 is expanded by ontology to
 ONTO_POS (title, { ‘invasive software’ }, OntologySources).

Since no perfect match exists, the overly pessimistic quality assignment is replaced by a more optimistic one with the help of the ontology.

Ontology enhances the BMO related preference order of ‘invasive software’ and ‘all others’ by raising the quality level of some concepts classified as *sufficient*, because they have a closer semantic relationship with the original concept ‘invasive software’ than others. Now having better alternatives than ‘all others’, the result set better fits the query with a higher level of quality (see Fig. 2).

Depending on the involved attributes of the query, P-News decides which ontologies have to be used for the expansion. This attribute-aware procedure sharpens the focus by just applying those ontologies which semantically model the domain of the attribute. The replacement or extension of ontologies is easily handled by changing

the mapping from the attributes to external ontology files or online resources like WordNet. The quality of the results of the preference POS and its extension ONTO_POS are compared in the following figure:

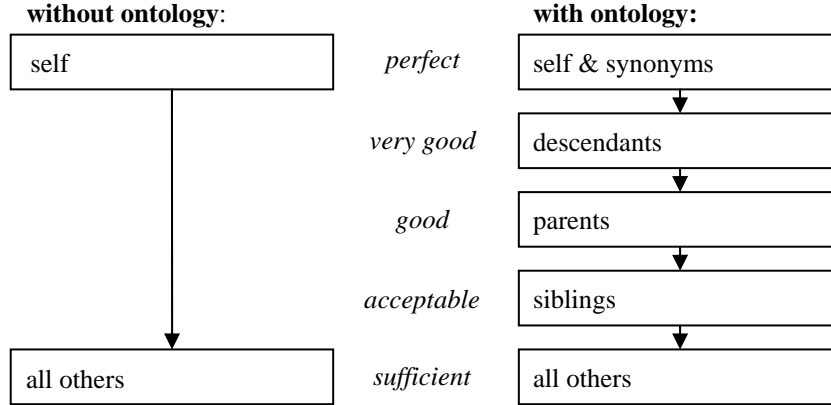


Fig. 2. Quality refinement of POS to ONTO_POS

Analogous to the above ONTO_POS preference, the preferences ONTO_NEG and ONTO_POS/NEG are defined.

3 Test Lab for P-News

Following the approach in [14], a methodology for constructing the test lab for P-News is as follows:

Construct for Evaluation is our alerting system P-News. The component ‘Query Expansion by Ontology’ is involved in the evaluation. *Context of Evaluation* is a deeply personalized alerting service. We use a system-centered approach, by which the tested system is put into a black box, or we use a glass box to change components to measure the impacts. *Criteria reflecting performance* take care of the user-specific wishes as input and the generated links to interesting video segments as output. We use set-theoretic operations like difference, intersection and subset to analyze the results of different configurations qualitatively (see section 1). *Methodology* for doing evaluation is implemented by a test lab, which guarantees ease of use and reproducibility of the experimentation.

3.1 System Architecture

The system architecture of the test lab is shown in Fig. 3. The tested system P-News can be tuned by a configuration file to switch on or off any components, especially the ontology.

In the start-up phase, the system binds the core system with the selected component(s) to produce an agent system starting und interacting with P-News. The test lab

frames the ad-hoc generated P-News system and just takes care of the correlated input and output sets of the underlying tested system.

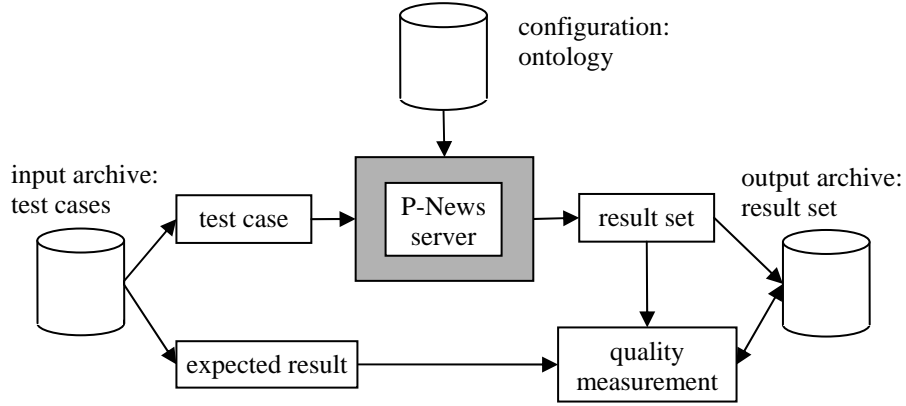


Fig. 3 System-centric test lab for P-News

After starting P-News in a specified configuration, the test lab scans the output to check, if a corresponding output result set has been created for each input test case. If the test suite passed the P-News server, the server process is shut down and a new instance of P-News is created with the next configuration. The test stops, when all selected configurations of the system have been tested by the test suite.

Example 4: We consider the following complex Pareto preference:

ONTO_POS (FileFormat, { 'MPG' }, OntologySources) \otimes
 BETWEEN (FileSize [2.9E8, 3.1E8])

A customer wants to get MPEG-encoded videos with a file size between 290 MB and 310 MB. The encoding as well as the size has the same importance. Looking at this example, the attribute 'FileFormat' may contain values like 'MPG' or 'MPEG'. They are synonyms, and users normally don't know which abbreviation is used inside a system. P-News generates the following result set:

```

<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE ResultSet SYSTEM "result.dtd">
<ResultSet size="2" config="ontology+quality assessment" source="query31">
<element Quality="good">
  ServerURI?file=FileURI_1&begin=0&end=408.01
</element>
<element Quality="good">
  ServerURI?file=FileURI_2&begin=407.97&end=561.79
</element>
</ResultSet>
  
```

This example shows that two video segments addressed by the attributes 'begin' and 'end' are found in two different videos (FileURI_1, FileURI_2). Both have a quality level of *good*.

3.2 Influence of Ontologies on Quality Assessment

Now we take a closer look at the interrelation between quality assessment and use of ontology. Analyzing the results and queries, which have been formulated by an independent testing person, we noticed an increase of the quality level. Not surprisingly the increase of quality arises only in those test cases, where we did not have perfect matches and where the user's terms are also be found in the involved ontology. If we get a hit in the ontology, then we receive a semantic context of other concepts semantically related to the original one. Descendants correspond to more specialized concepts, ancestors state more generic concepts, and the siblings have a closer relationship to the selected concept than all the other concepts except for children and parents.

The query expansion by ontology generates the aforementioned terms and uses them in a strict order preferring more specific concepts (see section 2.4). Therefore we may receive further matches, which stem from the expansion and raise the quality level. If the original query already yields a perfect hit, no improvements are achievable, because 'self' is always better than any possible semantic context.

Example 4 (cont.): We reuse the complex Pareto preference of example 4 to point out, how the ontology influences the quality assessment. Assuming that no perfect hits are found for the selected scale of 'FileSize', the BMO model may deliver alternatives rated with the quality level *good* among others. We assume also that the database has no objects with 'FileFormat' of value 'MPG'. Thus for this attribute the BMO model delivers a set of alternatives which is completely rated as *sufficient*. If the customer has chosen a notification constraint of *very good* quality, he gets an empty result set because all alternatives have worse quality then specified.

After using ontology the original query is expanded by the term 'MPEG' as well as by others. This term is classified as a synonym to the original 'MPG' and rated as *perfect*. Thus the result set of the Pareto query gets a compound quality level of *perfect* using the optimistic valuation, and therefore passes the filter of the user-specific quality level.

We conclude that the results are rated higher, if the expansion by ontology generates additional semantics related terms, which yield higher quality hits in the corpus.

4 Practical Challenges and Related Work

For the system-centric evaluation of P-News, we constructed an IT-related ontology from different sources. Our ontology describes 1616 relations between 1503 different concepts, of which 72 concepts have in total 154 synonyms. This ontology is still rather small in comparison with commercially used ones (e.g. the 'Unified Medical Language System' (UMLS) includes 975,354 concepts and 2.4 million concept names), but it fits our about 400 multimedia objects dealing with computer science topics. Currently the test suite consists of about 200 cases.

As stated in [8] recall for the ontology-based query expansion outperforms recall for keyword-based techniques. For query cases expanded by the concepts of children,

usually the precision increases, but may decrease, if ancestors' concepts are included. Context queries (see section 4.1) always ameliorate recall and precision. For disambiguation (see section 4.2) the correlation between selected concepts based on semantic closeness is determined. In [13] also the positive impacts of synonyms and hyponyms are stated, but the terms of the sense definition yield the greatest increase.

Analyzing the reasons, why the expansion failed or even was not applied, some linguistic and context related requirements have been found by the help of the test suite for further improvements to be still implemented.

4.1 Linguistic Challenges

Lots of entries in an ontology are annotated in the plural to describe a concept, but obviously singular is also used in the attributes or by users. This phenomenon can be reduced by stemming [6]. Other grammatical categories like verbs can also be derived from the linguistically correct root. In ontology nouns (e.g. optimization) are used as entries, but many full-text attributes are annotated by sentences or even short stories, in which often verbs (e.g. optimize) describe the issue.

There is also a mismatch of concepts between the unstructured string domain of a full-text attribute and the hierarchically structured concepts of an ontology, if compound nouns or nominal phrases ('A's B' or 'B of A') are used by the user as search terms in the ontology. Neglecting the inherent structure of the user input and only interpreting it as flat string, conflicts arise from the hierarchical structure of the ontology. Let's look at an excerpt of the Computer Science related ontology, to clarify that an expansion by ontology may rely on the context of ancestors of the selected concept:




Fig. 4. Concept depending on ancestors

Example 5: A query gets no impact from the above ontology, if the query is asking for 'Applications of Artificial Intelligence', 'Applications in Artificial Intelligence', or 'Artificial Intelligence Applications' for any attribute, because at first view these compound terms constitute no concepts in the ontology.

But if the flat terms of the query are transformed to a nested (context-sensitive) expression like already done in P-News [16], the ontology revamps its contribution.

Let 'A'/'B' denote that 'B' must be a descendant of 'A'. The original term is transformed to ONTO_POS (Attribute, {'Artificial Intelligence'/'Applications'}, OntologySources) and the query expansion may generate results of higher quality.

4.2 Contextual Challenges

If the expansion yields several hits in the ontology, then at present the *fusion* of the context of all instances of the input concept is implemented. In some cases it seems not to be appropriate to use the union of all instances, but to analyze the context of each instance, in order to select the best matching instance driven by content. Otherwise *loss of focus* (as mentioned in [17]) may happen.

The same effect is produced by disjoint semantic meanings of a single term. Linguistically spoken, the term is a homograph (see Fig. 5).

Example 6: If ‘Logic Programming’ is used as the search key in the following ontology, then it’s apparent that two superseded concepts ‘Programming Techniques’ and ‘Theorem Proving’ are applicable. Supposing that we receive no hit for descendants and parents, the expansion takes care of the siblings and adds ‘Resolution’ to the query. Unfortunately the concept ‘Resolution’ also belongs as a child to the concept ‘Screen’ in the meaning of ‘screen resolution’.




Fig. 5. Ontology prone to a loss of focus

We get a loss of focus, if the context is neglected. In the case of multiple parents of a sibling, we have to focus on those concepts which are more related to the original one. One possible way to implement this feature consists in deteriorating the quality level of the sibling from *acceptable* to *sufficient*. In the worst case ambiguous concepts are rated as *sufficient* (see Fig. 6).

Example 6 (cont.): We split the BMO-set corresponding to ‘Resolution’ into two parts. One set includes the terms which correspond to the context related concepts of ‘Logic Programming’ and ‘Resolution’, i.e., ‘Computer Science’//‘Resolution’ is still rated as *acceptable* whereas ‘Screen’//‘Resolution’ drops to *sufficient* (see Fig. 6). The second set contains the remaining elements rated as *sufficient*.

Assuming that a sibling has multiple parents, loss of focus may arise. To handle this challenge we exploit the context. We identify the common parents of *self* and this sibling. All the parents of the common parents (restricted set of *self*’s grandfathers) are used as the POS-set of a POS preference applied to the BMO-set of the sibling. The common parents must not be considered, because they already have a better rating according to the ONTO_POS preference.

Example 6 (cont.):

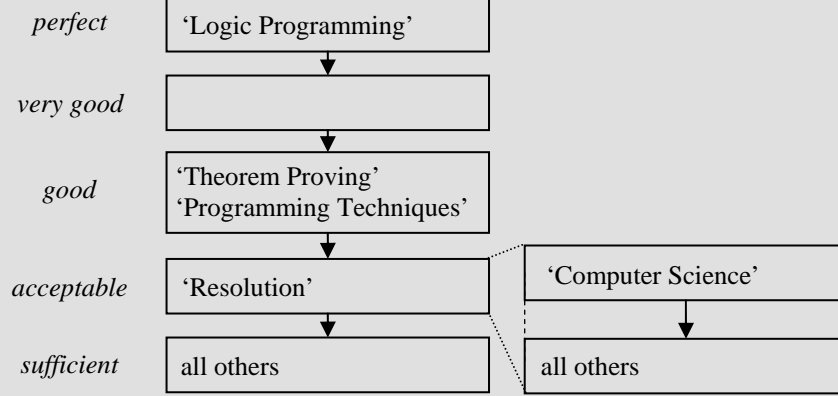


Fig. 6. Quality refinement reducing a loss of focus

Thus the preference model also offers a way to reduce a loss of focus by exploiting the context generated by ontology-based query expansion.

5 Summary and Outlook

In this paper we have investigated the impact of preference query expansion by ontology on the quality of our experimental P-News alerting system. We have implemented a system-centric test lab for P-News to facilitate the tedious evaluation task. From a first series of test cases we could find strong evidence for the following conjectures: The use of ontologies can improve the quality of the BMO query result, and it can be controlled to reduce a loss of focus during query expansion. The unique feature of our work is demonstrated by the easy and consistent integration of new requirements into the preference model framework to meet the deeply personalized needs of each client.

Our work in progress will explore these issues systematically in depth to get more comprehensive results on the benefits of using ontologies for preference query expansion. We are also planning to implement an on-the-fly-check of ontology generated concepts with the 'world knowledge' available online to identify ambiguous concepts or even to sort out distracting concepts. For this purpose WordNet of the university of Princeton is best suited. Depending on the domain of the attributes we want to take account of extended application-specific ontologies (e.g. for medicine: UMLS and MeSH, or for goods and services: UNSPSC and WAND).

Acknowledgements

P-News is funded within the German Research Foundation's strategic research initiative 'Distributed Processing and Delivery of Digital Documents ($V^III D^II$)'. We also like to thank Georg Berky for implementing the test lab.

References

1. Chomicki, J.: Preference formulas in relational queries. In *ACM Transactions on Database Systems (TODS 2003)*, Vol. 28/4, 427-466
2. Döring, S., Fischer, S., Kießling, W., Preisinger, T.: Optimizing the Catalog Search Process for E-Procurement Platforms. To be published in *Proc. of the Int. Workshop on Data Engineering Issues in E-Commerce (DEEC 2005)*, Tokyo, Japan
3. Fischer, S.: Personalized Query Result Presentation and Offer Composition for E-Procurement Applications. Dissertation (2004), Univ. of Augsburg (submit. for publicat.)
4. Fuhr, N., Hansen, P., Micsik, A., Solvberg, I.: Digital Libraries: A Generic Classification and Evaluation Scheme. In Constantopoulos, P., Solvberg, I. (eds.): *Research and Advanced Technology for Digital Libraries (ECDL 2001)*, Darmstadt, Germany, 187-199
5. Holland, S., Kießling, W.: Situated Preferences and Preference Repositories for Personalized Database Applications. In *Proc. 23rd Int. Conf. on Conceptual Modeling (ER 2004)*, Shanghai, China, 511-523
6. Hull, D. A.: Stemming algorithms: a case study for detailed evaluation. *Journal of the American Society for Information Science*, 47(1), (1996), 70-84
7. Kazai, G., Lalmas, M., Fuhr, N., Gövert, N.: A report on the first year of the Initiative for the Evaluation of XML retrieval (INEX'02). In *Journal of the American Society for Information Science and Technology*, Vol. 55/6, 551-556
8. Khan, L., McLeod, D., Hovy, E.: Retrieval effectiveness of an ontology-based model for information selection. In the *VDLB Journal* (2004) 13, 71-85
9. Kießling, W.: Foundations of Preferences in Database Systems. In *Proc. Int. Conf. on Very Large Databases (VLDB 2002)*, Hong Kong, China, 311-322
10. Kießling, W.: Preference Queries with SV-Semantics. In J. R. Haritsa, T. M. Vijayaraman (eds.): *Advances in Data Management 2005 (COMAD 2005)*, Goa, India, 15-26
11. Kießling, W., Fischer, S., Döring, S.: COSIMA B2B – Sales Automation for E-Procurement. In *Proc. 6th IEEE Conference on E-Commerce Technology (CEC'04)*, San Diego, CA, USA, 59-68
12. Kießling, W., Hafenrichter, B., Fischer, S., Holland, S.: Preference XPATH, A Query Language for E-Commerce. In: Buhl, H. U., Huther, A., Reitwiesner, B. (eds.): *Information Age Economy*. Physika-Verlag, Heidelberg (2001), 427-440
13. Navigli, R., Velardi, P.: An Analysis of Ontology-based Query Expansion Strategies. Workshop on Adaptive Text Extraction and Mining (ATEM 2003) in the 14th European Conference on Machine Learning (ECML 2003), Cavtat-Dubrovnik, Croatia
14. Saracevic, T.: Evaluation of digital libraries: An overview. Presentation at the DELOS WP7 Workshop on the Evaluation of Digital Libraries, 4-5 October 2004, Padua, Italy
15. Staab S., Studer, R.: *Handbook on Ontologies*. Springer 2004
16. Wang, Q., Balke, W.-T., Kießling, W., Huhn, A.: P-News: Deeply Personalized News Dissemination for MPEG-7 based Digital Libraries. In R. Heevy, L. Lyon (eds.): *Research and Advanced Technology for Digital Libraries (ECDL2004)*, Bath, UK, 256-268
17. Weikum, G., Graupmann, J., Schenkel, R., Theobald, A.: Towards a Statistically Semantic Web. In P. Atzeni, W. Chu, H. Lu, et al. (eds.): *Conceptual Modeling - ER 2004: 23rd International Conference on Conceptual Modeling (ER 2004)*, 6