An Ontology for WWW Summarization in Bone Marrow Transplantation (BMT): an Interim Report

Brigitte Endres-Niggemeyer[1], Bernd Hertenstein[2], Claudia Villiger[1], and Carsten Ziegert[2]

[1] Fachhochschule Hannover/University of Applied Sciences,
Dept. of Information and Communication, Ricklinger Stadtweg 120
D-30459 Hanover, Germany
{Brigitte.Endres-Niggemeyer, Claudia.Villiger}@ik.fh-hannover.de
http://summit-bmt.fh-hannover.de/
[2] Medizinische Hochschule Hannover/Hanover Medical School
Dept. of Hematology and Oncology,
D-30623 Hanover, Germany
Hertenstein.Bernd@mh-hannover.de
Carsten.Ziegert@ik.fh-hannover.de

Abstract. In an interim report, we describe an ontology for WWW summarization in Bone Marrow Transplantation. It is text-based and qualifies as a grounded ontology. In addition, it is user-centered. For stating medical knowledge, we use first-order logic (FOL) extended with contexts. The ontology is prepared for serving several purposes in a summarization system. It will be stored and managed by an XML database server. Currently, it is still under construction, but developed so far that its features can be demonstrated.

1 Introduction

In this paper, we describe an ontology under construction. Its aim is to support efficient search for documents and summarization from the WWW for Bone Marrow Transplantation (BMT), a specialized and life-critical area of hematology.

We are presenting our BMT ontology in an interim report when about 20% of it is realized. The basic argument for doing so is that we should present the ideas at an early stage: nothing really new will come up by waiting for an implementation rate of some 80%, but only more implementation of the same. We cannot expect to one day report about completed work because in a real-world environment, ontology engineering is a never-ending task, sliding softly from a construction mode to the later update mode.

1.1 Background

To our best knowledge, there is no BMT ontology that might support summarizing. While we can use technical features and in part material from existing thesauri and ontologies, we are far from taking over the bulk of the ontology from external sources. A main content reason is that existing ontologies or thesauri, MeSH[1] (Medical Subject Headings) and UMLS (Unified Medical Language System)[2] included, do not reach the detail needed for summarizing in BMT. This is easily explained: They contain BMT concepts without specifically focusing on them. Their authors certainly never intended to provide concepts for

[1] The database of Medical Subject Headings (MeSH) is available at http://www.nlm.nih.gov/mesh/MBrowser.html.

[2] The metathesaurus of UMLS contains concepts and classifications from more than 60 vocabularies. The metathesaurus is available at http://umlsks3.nlm.nih.gov/.

text knowledge processing at a sentence level. For their aims, the vocabulary can be more general. A second main reason is related to text processing. Our ontology is a concept dictionary that must encompass the knowledge for dealing with concept occurrences in text, where the concepts may occur under a different form, in the most simple case as a synonym. This sort of task-specific knowledge is not found in existing ontologies or thesauri. A third problem resides in the reliability of ontologies and thesauri which have only indirect relations to the data they describe. One may doubt whether they have any empirical grounding in current texts of the domain that satisfies standard demands of corpus linguistics. Since text summarization must pave the way from concrete texts to the formal representation and the processing of their meaning, an empirically constructed ontology that is rooted in texts of its domain is the safer approach.

The situation is illustrated by the MIHMA project [15] knowledge representation for bmt line, a web site that serves as the main reference point to BMT material on the Internet for a diversified audience. The knowledge representation has features of an ontology and it focuses on the BMT domain. It draws upon UMLS and is used for retrieving web ressources. It is left to a domain expert to judge their relevance and integrate them into bmt line. Applause for a group that shared our domain cannot conceal the fact that their purposes and their ontology content were different from what we need.

Together with many colleagues, Hahn and Reimer [12] favor the use of linguistic and domain knowledge in text understanding. In practice, however, ontologies are not yet often used for the purpose of summarization. WordNet (concepts of common English) is by far the most popular ontology in use, applied for instance in SUMMARIST [14].

Technically speaking, today's choices include WWW-resident representation techniques, for instance XML-based ones inspired by SHOE [13]. We apply them in a situation where different and earlier approaches compete with newer web–driven ones. We adhere to logic representation formalisms. We state our facts and rules in first-order predicate logic extended by contexts [17] and enable safe [8] and simple conclusions.

1.2 Use of the BMT ontology for summarization

Our ontology is a dense representation of domain knowledge in Bone Marrow Transplantation (BMT) tailored to the needs of domain experts. Most of the factual and linguistic knowledge is incorporated in the ontology. For every domain concept, the ontology also encompasses statements of relevant knowledge about it. From an information retrieval perspective, these statements exclude wrong concept combinations. During summarization, a match with propositions helps to establish the relevance of candidate statements in retrieved texts with respect to the current question. For concept identification in running text, the BMT ontology stores lexical equivalents of concepts and the paraphrases by which they are expressed. Like a classic thesaurus, the ontology also conveys information for human users, in particular for query scenario formulation. In the following it is described in some detail how the ontology supports query formulation for information retrieval, text passage retrieval, and summarization proper.

- **Query scenario formulation and query expansion**
  For question-oriented summarization, we need reasonably precise queries. The better the question, the better chances the system has to come up with a helpful response. In order to have users state their queries in a well-structured fashion that the system can interpret, we provide scenarios, query forms specific for types of questions and situations. Into these forms, users fill concepts of the ontology. The scenarios accept statements and questions. Users can browse the concept, synonyms, hypernyms, hyponyms and description fields of ontology records. We provide definitions and descriptions of concepts assembled from many WWW sources and a few others.

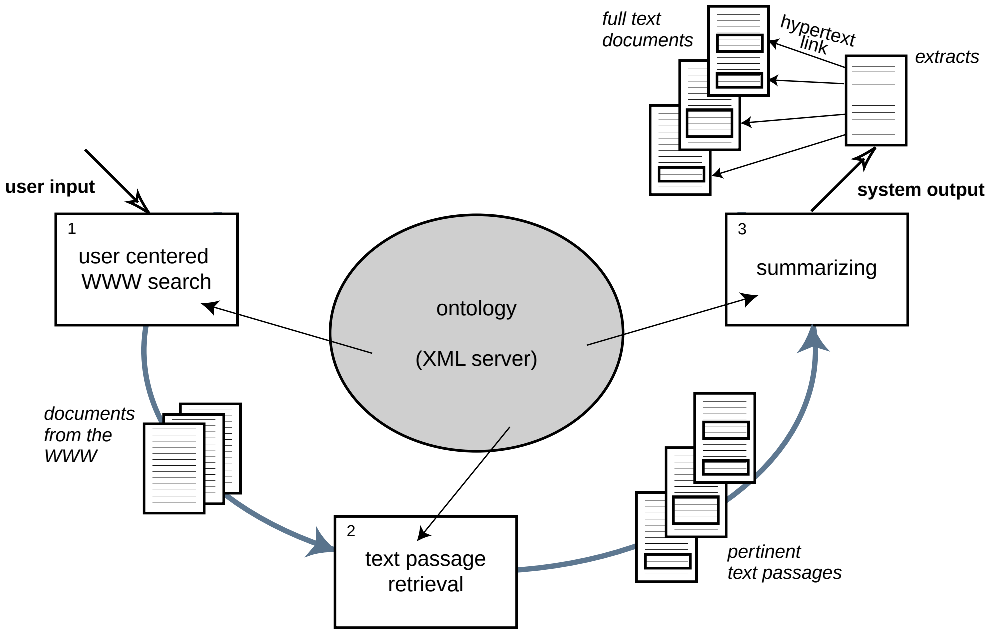

Fig. 1. The ontology as the knowledge server for the system components

After being reformulated in predicate logic form, the query scenario is enriched with related concepts from the ontology, and adapted to the search form of involved search engines (see processing step 1 in figure 1).

– **Text passage retrieval**
  As soon as retrieved documents arrive, they are roughly checked for relevance by scanning them for passages that include the query terms. The preliminary relevance check is performed using the concepts of the ontology (see processing step 2 in figure 1).

– **Summarizing**
  Summarizing means reducing information to its most important (relevant) points [4]. Most often, the source information is given as a text, and the resulting summary is a short text. Summarization mostly operates at sentence level, and summaries are composed of relevant statements (often sentences), as opposed to passages including unrelated relevant terms as in text passage retrieval. When the summary is to answer a question, only statements that correspond to the question at a sentence level can be relevant, all others are not. So passages that contain ontology concepts used in the query are checked for question-related relevance at sentence level. For this purpose, the ontology must store propositional knowledge with every concept, practically speaking logical propositions that use the current concept as an argument. One of the main differences between thesauri and ontologies is that ontologies state propositional knowledge about their concepts, whereas thesauri are restricted to core paradigmatic relations such as generic ones.
  Our system under development uses summarization agents that apply the strategies of human summarizers [4]. They are modelled in terms of cognitive agents (see step 3 in figure 1).

In this paper, we focus on content engineering problems of our ontology. First we describe the empirical method applied for concept acquisition and the organization of the ontology (see section 2.1). We proceed with a detailed discussion of the formalization process (see section 2.2). After that, we give some technical implementation details.

2 The BMT Ontology for Summarization

2.1 Empirical Ontology Development

In line with others [2, 9] we promote an approach to knowledge modeling based on texts. We argue that papers published in core journals of the domain are safer than other knowledge sources, especially oral information sources such as interviews or expert meeting results, because their contents and wording have been checked by the reviewers of the journal. Following an empirical procedure adapted from thesaurus building [5] we refer to classic thesaurus building methodology [1] where empirical and knowledge organization issues are discussed supported by decades of practical experience. We apply the principles of qualitative field research, especially the framework for inductive theory development [9]. According to the grounded theory framework, our ontology is a grounded model of its domain because all concepts are justified by and connected with their evidence in texts.

Because of the ontology's high visibility for users and its central role for the future system's practical value and acceptance, it is only consistent to expand the realm of user-centered system design [19, 22] from the user interface to the ontology, the system's main knowledge representation. This follows from the core idea of user-centered system design that we have to think the system from its user interface, because for users, the system is what the interface presents. User roles with active participation are common in user–centered and participatory design. According to this practice, we try to integrate our users as responsible subject matter co-authors of the ontology.

The ontology is built up according to the ideal 13 steps procedure [5] presented in table 1, with some adaptations to practical conditions.

1	2–3 current relevant papers or book chapters are exploited to obtain an initial stock of concepts
2	concepts are supplemented with WordNet knowledge
3	if available, MeSH descriptors are added
4	the meaning of the concepts is made explicit and they are formalized and represented for the use of different players
5	users set up search scenarios
6	from user search scenarios queries are derived, the search engines are started
7	the found documents are summarized
8	the summarization results are integrated into the question/answer scenarios
9	summaries are checked for failures by physicians and technical team members
10	the knowledge representation is improved
11	agents are adapted or created
12	back to step 5 as often as needed
13	a new partial ontology is integrated into the existing one

Table 1. The step-by-step procedure of empirical ontology construction and evaluation [5]

The first adjustment was stimulated by a domain expert. Instead of exploiting small numbers of papers from each subdomain, he proposed choosing two sorts of papers for the whole domain: BMT papers from Blood, a core journal of the domain where current issues are discussed, and BMT educational papers from the American Society of Hematology

(ASH) which present fundamental topics in a tutorial style. We are now arriving at stage 5 of the procedure in table 1, with the restriction that our first round of scenario testing has not yet given rise to a WWW search.

Experience of thesaurus builders shows that concepts in literature may systematically differ from concepts asked for in user queries. Several reasons for this phenomenon are known, for instance different interests in information demand and information offer (with demand sometimes preceding offer) or in knowledge background, as with knowledge producers rooted in scientific theories and application–oriented users referring to commercial categories. With this in mind, we feed our ontology also from user query scenarios. With an additional knowledge base where domain experts can state their own knowledge, we arrive at three knowledge sources (see figure 2):

– Papers of a core journal and educational papers in BMT are the prevailing knowledge source.
– Text–based knowledge is supplemented by concepts from user scenarios.
– Concept records provided by domain experts are added.

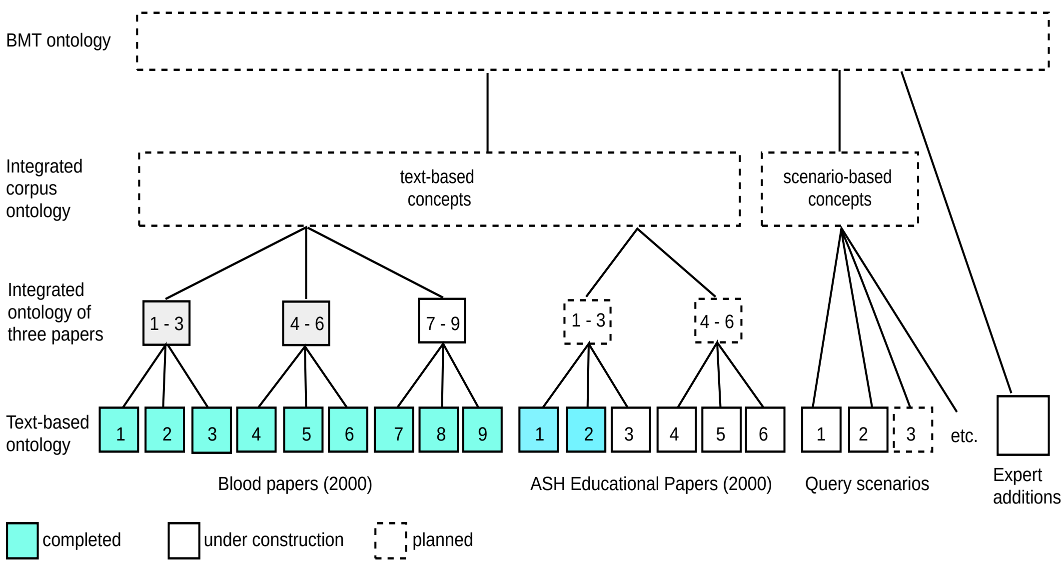

Fig. 2. The empirical structure of the BMT ontology

Figure 2 gives an overview of the empirical structure of the BMT ontology under construction. Partial ontologies with a light gray background are completed, planned ones have broken lines and clear backgrounds. All others with normal lines and white backgrounds are currently under construction.

Basic ontology records include the following main fields:

– concept
– synonyms
– hypernyms
– hyponyms
– description
– occurrences
– assertions

We add descriptions of concepts to the concept record, because when formulating questions, users must have the opportunity to check whether they have selected the appropriate concept. During knowledge acquisition from a paper all concept occurrences with their contexts, normally a sentence, are copied to the field 'occurrences'. If these contexts convey valid knowledge, more concise statements are entered into the assertions field of the record. They are transformed into predicate logic expressions later on (see section 2.2).

The ontology is grounded in so far as all concepts are justified by and connected with their evidence in text. As soon as we integrate text–based ontologies, the bulky occurrence data are left behind, but they can be recovered from higher levels of ontology integration. For the sake of ease and safety, integration is done by bundles of three text-based ontologies at a time.

During integration, conceptualizations are checked and adapted if necessary, double records are removed, and statements found in the source databases are entered into the respective record in the integrated part of the ontology. If necessary and possible, the concept descriptions are improved. A return to the source paper is sometimes necessary for checking and changing conceptualizations.

When we glean concepts from user query scenarios, the empirical evidence is somewhat weaker than elsewhere. Often, the physicians' scenario descriptions are very short, context information is poor or almost missing. To make user–defined scenarios fit for WWW retrieval, we have to modularize them and to fill them up with concepts from external sources. All concepts needed for stating usable scenarios are entered into the respective scenario ontology.

After illustrating the way the ontology is set up, we give a quantitative account of the concepts acquired so far. Table 2 shows the number of concepts derived from Blood (2000) papers. The scenario ontology of the first survey turn contains 211 concepts resulting from 28 scenarios. We have started lexical equivalences recently, the same is true for formalized contexts (see section 2.2).

Text-based Ontologies	Records	Integrated Ontologies	Records
Blood1	235		
Blood2	173		
Blood3	166	Blood1-3	479
Blood4	160		
Blood5	415		
Blood6	241	Blood4-6	781
Blood7	134		
Blood8	227		
Blood9	133		

Table 2. Overview of ontology records

SummIt top, the upper model of the ontology (its first two levels are shown in figure 3), has been conceived deductively by drawing from the ontologies comprised in the metathesaurus of UMLS (Unified Medical Language System). Just adopting one of these classifications proved to be awkward. The MeSH classification for example does not fit our needs:

- The MeSH taxonomy is not strictly generic which is a prerequisite for our ontology with its strict isa-relations for inheritance management.
- The MeSH taxonomy is broad and its granularity does not suffice for our purposes.

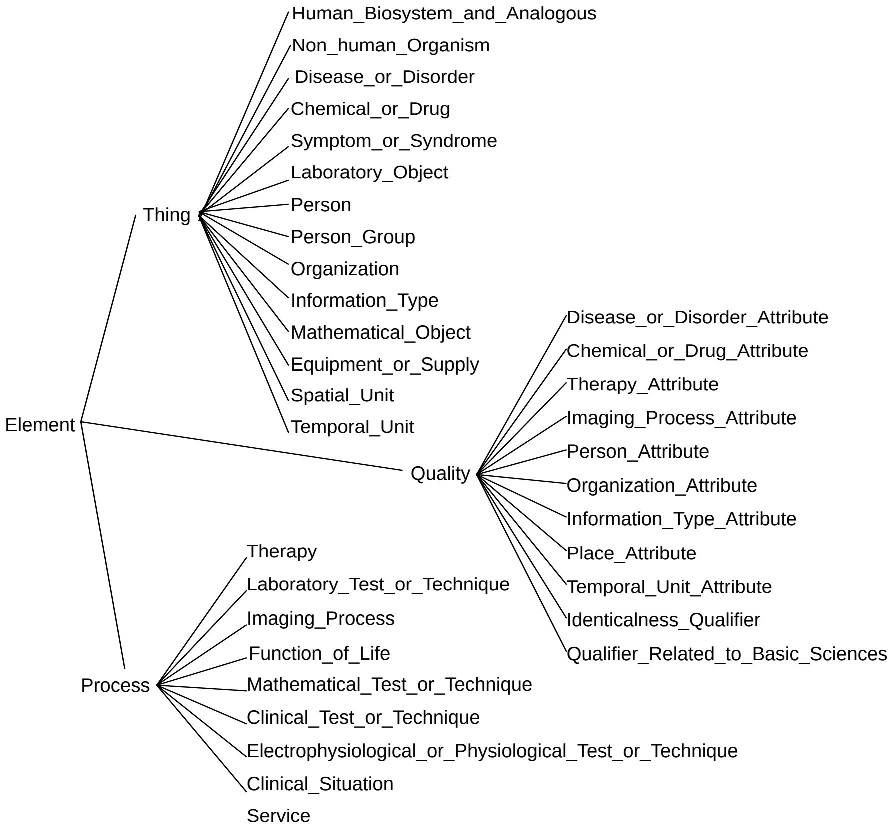

Fig. 3. The ontology top (SummIt top) in its current state

For the BMT taxonomy an expert selected those MeSH classes at the two upmost levels of this classification which fit the requests of our task. If requisite, they were reformulated so that the expression of generic relations was facilitated. We enhanced the classes by assigning 455 concepts to them and by inserting classes (e.g. clinical situation) if concepts could not be assigned to the existing taxonomy. If necessary, class designations and the internal structuring of classes were changed. The resulting classification was compared with other ontologies comprised in the metathesaurus of UMLS.

We also considered the Generalized Upper Model (GUM) [3] and the BigTop proposed by Sowa [20]. As we are using the ontology for text processing, we had an eye upon the classes that show up in natural language text and are recognized there by our parser[3].

The preliminary top structure is useful because it sets an anchor to the generic hierarchies, but it is also often enough disputed by our empirical data. Experts of the domain verify concepts gathered by the text–based approach as well as the internal structure of the resulting classes. The taxonomy is adapted inductively as needed.

2.2 From Textual Context to First-Order Context Expressions

In medicine many statements are valid only if the limits of their scope are respected. Perhaps more than elsewhere we have to represent presuppositions. Furthermore, many of the semantic problems in NLP involve using a context to determine the meaning of an utterance.

[3] The parser is provided by Conexor (http://www.conexor.fi).

To meet both the medical and the linguistic requirements, we apply the context formalization proposed by [11, 18]. Context expressions are composed of several propositions or predicates. While some of them set up the frame of thinking about the core propositions, the core predicates state the important points. A context expression as a whole must be true in domain expert eyes. Normally, it is the meaning equivalent of a sentence. When checking whether a context expression from input satisfies the conditions set by the query, summarizing agents first look for inconsistencies in context. If they find none, they examine the core of the context expression.

The context expressions and our predicate logic expressions are not physically stored in individual concept records, because almost all of them are shared by several concepts. Hence they are stored in and allocated from a central database. The information for recognizing concepts in running text (i.e. sets of textual equivalents) is stored in the central pool together with the predicates.

Context expressions are the main propositional knowledge representation format of the BMT ontology. They are obviously useful structures for stating medical knowledge because they limit the scope of an assertion. Context expressions assert that the proposition p is true ('ist') in the context c:

$$ist(c, p)$$

Both context and core propositions are predicate logic expressions. Following the example of Prolog, we accept only conjunctions and implications within context and core propositions. Disjunctions are expressed by alternative facts or rules. All predicates use ontology concepts as arguments.

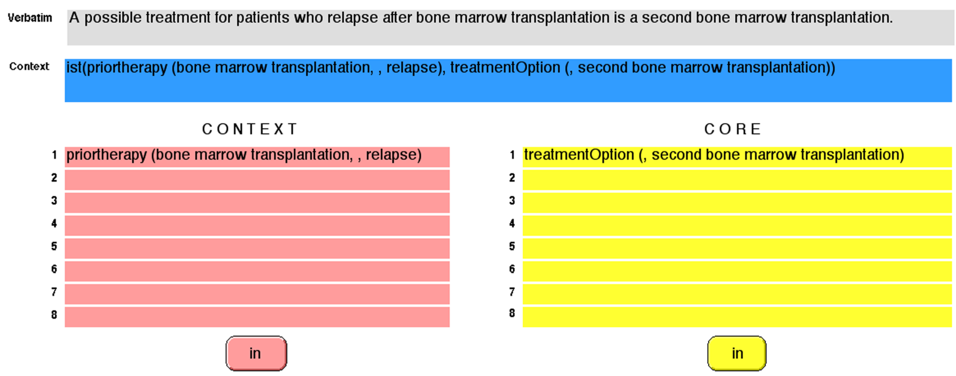

Fig. 4. A simple context expression is generated from context propositions on the left and core propositions on the right side

Figure 4 shows a context expression that is ready to enter the context database. It is accompanied by its formulation in natural language (top line). The context base accepts predicate expressions from the predicate base that are transferred either into the context area or into the core area. From the arriving context and core propositions, the context expression (second line) is generated.

Our context expressions are grounded upon knowledge found in the papers of our corpus. We extract contexts, often sentences, but also longer sequences if necessary for stating a context in stand-alone manner. While there are often quite explicit statements of medical knowledge, this is by no means always the case. We present a counterexample. It demonstrates that we can mine more and more relevant knowledge from a text by skilled interpretation. In Blood3 [19] we find:

> DLI therapy would seem an acceptable alternative to second allogeneic BMT, however, as the morbidity and mortality of second BMT is high and prohibitive for many patients.

> However, the incidence of acute or chronic GVHD and marrow aplasia may be acceptable compared with other potential treatment options such as second marrow transplantation.

> However, given the excessive toxicity anticipated from second BMT within 1 year of the initial transplant, a trial of UDLI or other investigational therapy seems warranted.

Among many other things, we learn here that DLI therapy is an alternative to second allogeneic BMT. Doing a second bone marrow transplantation is presupposed three times. Readers can conclude this without bothering about the medical terms our example sentences are teeming with. The formulation in our ontology (cf. top line in figure 4) states the point as follows:

> A possible treatment of patients who relapse after bone marrow transplantation is a second bone marrow transplantation.

Although this statement does not show up literally, it is a knowledge item whose validity is not confined to the paper at hand. It had to be derived by a competent human understander. She or he restates good assertions, elaborating implicit text knowledge if needed. Investing this intellectural work should enhance the quality of our ontology. Obviously, systems that can replace humans at this task are out of sight, only some technical support is feasible.

Textual formulations are stripped of rhetorical accessoires (example replacement rule: incidence of $X \Rightarrow X$) and prepared for later formalization. On their way to the context knowledge base they pass some normalization steps:

- Their concepts are standardized: only ontology concepts, no synonyms are accepted.
- They are reconstructed according to some syntax from the syntax base holding the syntaxes for predicate expressions. Thus they are normalized according to the number and kind of arguments.
- The resulting predicates are complemented with a natural language translation and stored in the predicate base for later reuse.
- The predicates are entered into the context expression, either into the core or the context position.

Stocks of context expressions are still modest at the time of writing: Some 243 of them are ready for use and approximately 340 predicates are equipped with equivalences.

2.3 Technical Implementation of the Ontology

From a technical point of view, ontology construction is supported by general–purpose tools. Source documents are PDF files and the selection of concepts from texts is facilitated by Conc[4], a freeware concordancer. The ontology in the current state of development is stored in a set of FileMaker databases[5]. The Filemaker application environment supports ontology construction, also the development of the taxonomy and the formulation of context expressions with the underlying predicates and their argument structures (syntaxes).

[4] The Conc concordancer is provided by the Summer Institute of Linguistics and is available at http://www.sil.org/computing/conc/.

[5] FileMaker is a relational database server with a GUI for the Apple Macintosh environment. FileMaker is a registered trademark of FileMaker, Inc. ©1984-2000 (http://www.filemaker.com).

For consistency checking of the taxonomy, the corresponding FileMaker entries are imported into an XML database server (see below) and then exported to Graphviz, a free tool for displaying graph structures[6]. It helps to verify the correctness of generic relations.

Although there are existing tools for ontology acquisition like Ontolingua [6] and upcoming standards for ontology representation like OIL [7] we decided to use plain XML and developed our own DTD for ontology representation. Thus we can use the XML database server dbXML[7] which provides a Java API for programming tasks and on that way retain the flexibility to import the ontology data into standardized tools later. Semantic checks may then be feasible or we will improve our tools in such a way that these checks are provided. When defining the DTD for the XML structure being used we restricted ourselves to the limitations propagated by the authors of SHOE [13] who demand semantic interoperability and therefore generality in the names of the XML tags. Thus, our tags are called <concept>, or <proposition> holding subtags like <synonym> or <predicate_name> instead of using tags like <cancer> and implementing concepts in a manner like <cancer> leukemia </cancer> as dismissed by the SHOE team.

3 Conclusion

We have presented an ontology for WWW summarization in Bone Marrow Transplantation. Due to the requirements of the task, the users, and the domain, the ontology comes with some extraordinary features. It is text–based — it qualifies as a grounded ontology — and it is user-centered. Currently, the ontology is still growing but its innovative features are set. We proceed now towards test and application. We have invested much intellectual work, work that is difficult for humans. As long as the core ontology is being constructed, we do not think that the human intellectual effort can be replaced by machine approaches without a quality loss that is not acceptable in a medical domain. Further extension of the ontology and its maintainance may be facilitated by semi-automatic ontology acquisition [16] and by concept input suggested by users. Compared with the development of thesauri, which may be much larger than the BMT ontology and as demanding in intellectual effort, our own intellectual investment is, after all, far from exiting.

Acknowledgements

Implementation of the SummIt-BMT system is supported by the German Science Foundation (DFG) under grant EN 186/6-1 and HE 2927/2-1, by the German Federal Ministry of Education and Research (bmbf) under grant 1701200, and by the Ministry of Science and Culture of Lower Saxony under grant 1999.384.

The authors would like to thank Andrea Köpcke and Juliane Topp for their contributions to this paper. Thanks also to the reviewers, especially to the one who confronted us with a long list of well-deliberated questions.

References

1. Aitchison, J., Gilchrist, A.: Thesaurus Construction and Use: A Practical Manual. 3rd edn. Aslib, London (1997)
2. Aussenac-Gilles, N., Biébow, B., Szulman, S.: Corpus analysis for conceptual modelling. EKAW 2000. `http://www.irit.fr/ACTIVITES/EQ_SMI/GRACQ/WSEKAW2000/accepted.html` (2000)

[6] Graphviz is provided by AT&T (http://www.research.att.com/sw/tools/graphviz).
[7] dbXML has been made available by The dbXML Group (http://www.dbxmlgroup.com). The server is still under active development.

3. Bateman, J. A., Henschel, R., Rinaldi, F.: Generalized Upper Model. http://www.darmstadt.gmd.de/publish/komet/gen-um/newUM.html (1995)
4. Endres-Niggemeyer, B.: Summarizing Information. Springer, Berlin (1998)
5. Endres-Niggemeyer, B.: Empirical Methods for Ontology Engineering in Bone Marrow Transplantation. International Workshop on Ontological Engineering on the Global Information Infrastructure, Dagstuhl Castle, May 25, 1999. http://www.ik.fh-hannover.de/ik/person/ben/OntoBone.pdf (1999)
6. Farquhar, A., Fikes, R., Pratt, W., Rice, J.; Collaborative Ontology Construction for Information Integration. Knowledge Systems Laboratory Department of Computer Science, KSL-95-63 (1995)
7. Fensel, D., Horrocks, I., Van Harmelen, F., Decker, S., Erdmann, M., Klein, M.: OIL in a Nutshell. In: Knowledge Acquisition, Modeling, and Management, In: Dieng, R. et al. (eds.): Proceedings of the European Knowledge Acquisition Conference (EKAW-2000). Lecture Notes in Artificial Intelligence, LNAI. Springer, Berlin (2000)
8. Fox, J., Das, S. K.: Safe and Sound: Artificial Intelligence in Hazardous Applications. MIT Press, Cambridge MA (2000)
9. Glaser, B. G., Strauss, A. L.: The Discovery of Grounded Theory: Strategies for Qualitative Research. 11th edn. Aldine Atherton, New York (1980)
10. Golebiowska, J., Dieng-Kuntz, R., Corby, O., Mousseau, D.: Building and Exploiting Ontologies for an Automobile Project Memory, IJCAI 01 Workshop on Ontologies for Information Sharing. http://www.tzi.de/buster/IJCAIwp/Finals/golebiowska.pdf (2001)
11. Guha, R. V.: Contexts: A Formalization and Some Applications. PhD thesis, Computer Science Department, Stanford University, Stanford, CA (1991). Also technical report ACT-CYC-423-91. Microelectronics and Computer Technology Corporation, Austin, TX (1991)
12. Hahn, U., Reimer, U.: Knowledge-Based Text Summarization: Salience and generalization Operators for Knowledge Base Abstraction. In: Mani, I., Maybury M.T. (eds.): Advances in Automatic Text Summarization. MIT Press, Cambridge, Mass. (1999)
13. Heflin, J., Hendler, J., Luke, S.: SHOE: A Knowledge Representation Language for Internet Applications. Technical Report CS-TR-4078 (UMIACS TR-99-71), Dept. of Computer Science, University of Maryland at College Park (1999) Also http://www.cs.umd.edu/projects/plus/SHOE
14. Hovy, E., Lin, C.-Y.: Automated Text Summarization in SUMMARIST. In: Mani, I., Maybury M.T. (eds.): Advances in Automatic Text Summarization. MIT Press, Cambridge, Mass. (1999)
15. Kindermann, C. Hoppe, T., Marwinski, D., Oertel, H., Paulus, O. K., Heimann, S., Schmiedel, A., Tolksdorf, R., Volle, P.: The MIHMA demonstrator application: bmt line. KIT–Report 132. Technische Universität Berlin, Berlin (1996) Also http://flp.cs.tu-berlin.de/~kit/reportliste/kitlistehtml.html
16. Maedche, A., Staab, S.: Semi-automatic engineering of ontologies from text, Proceedings of the 12th International Conference on Software and Knowledge Engineering (SEKE 2000), Chicago, USA, 6-8 July 2000. http://www.aifb.uni-karlsruhe.de/WBS/publications/2000/seke_amaetal_2000.pdf (2000)
17. McCarthy, J., Buvač, S.: Formalizing Context (Expanded Notes). In: Aliseda, A., Glabbeek, R. van, Westerstohl, Dag (eds.): Computing Natural Language. Stanford Universiy, Stanford (1997) 13–50. Also http://www-formal.stanford.edu/buvac/formalizing-context.ps
18. Norman, D. A., Draper, S. W. eds.: User Centered System Design. Erlbaum, London (1986)
19. Porter, D.L., Collins Jr, R. H., Hardy, C., Kernan, N. A., Drobyski, W. R., Giralt, S., Flowers, M. E. D., Casper, J., Leahey, A. , Parker, P. , Mick, R., Bate-Boyle, B., King, R., Antin, J. H.: Treatment of Relapsed Leukemia after Unrelated Donor Marrow Transplantation with Unrelated Donor Leukocyte Infusions. Blood Vol. 95(4) (2000) 1214–1221
20. Sowa, J. F.: Knowledge Representation: Logical, Philosophical, and Computational Foundations. Brooks Cole Publishing Co., Pacific Grove, CA (2000)
21. Winograd, T. (ed.): Bringing Design to Software. Addison-Wesley, Reading MA (1996)