

SEAL — A Framework for Developing SEmantic Web PortALs

^{1,3}Alexander Maedche, ^{1,2}Steffen Staab, ¹Nenad Stojanovic, ^{1,2,3}Rudi Studer,
and ^{1,2}York Sure

¹Institute AIFB, University of Karlsruhe, D-76128 Karlsruhe, Germany
<http://www.aifb.uni-karlsruhe.de/WBS>
{ama,sst,nst,rst,ysu}@aifb.uni-karlsruhe.de

²Ontoprise GmbH, Haid-und-Neu Straße 7, 76131 Karlsruhe, Germany
<http://www.ontoprise.de>

³FZI Research Center for Information Technologies,
Haid-und-Neu Straße 10-14, 76131 Karlsruhe, Germany
<http://www.fzi.de/wim>

Abstract. The core idea of the Semantic Web is to make information accessible to human and software agents on a semantic basis. Hence, web sites may feed directly from the Semantic Web exploiting the underlying structures for human and machine access. We have developed a generic approach for developing semantic portals, *viz.* SEAL (SEmantic portAL), that exploits semantics for providing and accessing information at a portal as well as constructing and maintaining the portal.

In this paper, we discuss the role that semantic structures make for establishing communication between different agents in general. We elaborate on a number of intelligent means that make semantic web sites accessible from the outside, *viz.* semantics-based browsing, semantic querying and querying with semantic similarity, and machine access to semantic information at a semantic portal. As a case study we refer to the AIFB web site — a place that is increasingly driven by Semantic Web technologies.

1 Introduction

The widely-agreed core idea of the Semantic Web is the delivery of data on a semantic basis. Intuitively the delivery of semantically apprehended data should help with establishing a higher quality of communication between the information provider and the consumer. How this intuition may be put into practice is the topic of this paper.

We discuss means to further communication on a semantic basis. For this one needs a theory of communication that links results from semiotics, linguistics, and philosophy into actual information technology. We here consider *ontologies* as a sound semantic basis that is used to define the meaning of terms and hence to support intelligent access, *e.g.* by semantic querying [5] or dynamic hypertext views [19].

Thus, ontologies constitute the foundation of our SEAL (SEmantic portAL) approach. The origins of SEAL lie in Ontobroker [5], which was conceived for semantic search of knowledge on the Web and also used for sharing knowledge on the Web [2].

It then developed into an overarching framework for search and presentation offering access at a portal site [19]. This concept was then transferred to further applications [1, 21, 24] and is currently extended into a commercial solution¹.

We here describe the SEAL core modules and its overall architecture (Section 3). Thereafter, we go into several technical details that are important for human and machine access to a semantic portal.

In particular, we describe a general approach for semantic ranking (Section 4). The motivation for semantic ranking is that even with accurate semantic access, one will often find too much information. Underlying semantic structures, *e.g.* topic hierarchies, give an indication of what should be ranked higher on a list of results.

Finally, we present mechanisms to deliver and collect machine-understandable data (Section 5). They extend previous means for better digestion of web site data by software agents. Before we conclude, we give a short survey of related work.

2 Ontology and knowledge base

For our AIFB intranet, we explicitly model relevant aspects of the domain in order to allow for a more concise communication between agents, *viz.* within the group of software agents, between software and human agents, and — last not least — between different human agents. In particular, we describe a way of modeling an ontology that we consider appropriate for supporting communication between human and software agents.

2.1 Ontologies for communication

Research in ontology has its roots in philosophy dealing with the nature and organisation of being. In computer science, the term ontology refers to an engineering artifact, constituted by a specific vocabulary used to describe a particular model of the world, plus a set of explicit assumptions regarding the intended meaning of the words in the vocabulary. Both, vocabulary and assumptions, serve human and software agents to reach common conclusions when communicating.

Reference and meaning. The general context of communication (with or without ontology) is described by the meaning triangle [15]. The meaning triangle defines the interaction between symbols or words, concepts and things of the world (*cf.* Figure 1).

The meaning triangle illustrates the fact that although words cannot completely capture the essence of a reference (= concept) or of a referent (= thing), there is a correspondence between them. The relationship between a word and a thing is indirect. The correct linkage can only be accomplished when an interpreter processes the word invoking a corresponding concept and establishing the proper linkage between his concept and the appropriate thing in the world.

¹ *cf.* <http://www.time2research.de>




Fig. 1. The Meaning Triangle

Logics. An ontology is a general logical theory constituted by a vocabulary and a set of statements about a domain of interest in some logic language. The logical theory specifies relations between signs and it apprehends relations with a semantics that restricts the set of possible interpretations of the signs. Thus, the ontology reduces the number of mappings from signs to things in the world that an interpreter who is committed to the ontology can perform — in the ideal case each sign from the vocabulary eventually stands for exactly one thing in the world.

Figure 2 depicts the overall setting for communication between human and software agents. We mainly distinguish three layers: First of all, we deal with things that exist in the real world, including in this example human and software agents, cars, and animals. Secondly, we deal with symbols and syntactic structures that are exchanged. Thirdly, we analyze models with their specific semantic structures.




Fig. 2. Communication between human and/or software agents

Let us first consider the left side of Figure 2 without assuming a commitment to a given ontology. Two human agents HA_1 and HA_2 exchange a specific sign, *e.g.* a word like “jaguar”. Given their own internal model each of them will associate the sign to his

own concept referring to possibly two completely different existing things in the world, *e.g.* the animal *vs.* the car. The same holds for software agents: They may exchange statements based on a common syntax, however, they may have different formal models with differing interpretations.

We consider the scenario that both human agents commit to a specific ontology that deals with a specific domain, *e.g.* animals. The chance that they both refer to the same thing in the world increases considerably. The same holds for the software agents SA_1 and SA_2 : They have actual knowledge and they use the ontology to have a common semantic basis. When agent SA_1 uses the term “jaguar”, the other agent SA_2 may use the ontology just mentioned as background knowledge and rule out incorrect references, *e.g.* ones that let “jaguar” stand for the car. Human and software agents use their concepts and their inference processes, respectively, in order to narrow down the choice of referents (*e.g.*, because animals do not have wheels, but cars have).

A new model for ontologies. Subsequently, we define our notion of ontology. However, in contrast to most other research about ontology languages it is not our purpose to invent a new logic language or to redescribe an old one. Rather what we specify is a way of *modeling* an ontology that inherently considers the special role of signs (mostly strings in current ontology-based systems) and references.

Our motivation is based on the conflict that ontologies are for human and software agents, but logical theories are mostly for mathematicians and inference engines. Formal semantics for ontologies is a *sine qua non*. In fact, we build our applications on a well-understood logical framework, *viz.* F-Logic [10]. However, in addition to the benefits of logical rigor, user and developer of an ontology-based system profit from ontology structures that allow to elucidate possible misunderstandings.

For instance, one might specify that the sign “jaguar” refers to the union of the set of all animals that are jaguars and the set of all cars that are jaguars. Alternatively, one may describe that “jaguar” is a sign that may either refer to a concept “animal-jaguar” or to a concept “car-jaguar”. We prefer the second way. In conjunction with appropriate GUI modules (*cf.* Sections 3ff) one may avoid presentations of ‘funny symbols’ to the user like “animal-jaguar”, while avoiding ‘funny inference’ such as may arise from artificial concepts like the union of the sets denoted by ‘animal-jaguar’ and ‘car-jaguar’.

2.2 Ontology *vs.* knowledge base

Concerning the general setting just sketched, the term ontology is defined — more or less — as some piece of formal knowledge. However, there are several properties that warrant the distinction of knowledge contained in the ontology *vs.* knowledge contained in the so-called *knowledge base*, which are summarized in Table 1.

The ontology constitutes a general logical theory, while the knowledge base describes particular circumstances. In the ontology one tries to capture the general conceptual structures of a domain of interest, while in the knowledge base one aims at the specification of the given state of affairs. Thus, the ontology is (mostly) constituted by *intensional* logical definitions, while the knowledge base comprises (mostly) the *extensional* parts. The theory in the ontology is one which is mostly developed during the set up (and maintenance) of an ontology-based system, while the facts in the knowledge

Table 1. Distinguishing ontology and knowledge base

	Ontology	Knowledge base
Set of logic statements	yes	yes
Theory	general theory	theory of particular circumstances
Statements are mostly	intensional	extensional
Construction	set up once	continuous change
Description logics	T-Box	A-Box

base may be constantly changing. In description logics, the ontology part is mostly described in the T-Box and the knowledge base in the A-Box. However, our current experience is that it is not always possible to distinguish the ontology from the knowledge base by the logical statements that are made. In the conclusion we will briefly mention some of the problems referring to some examples of following sections.

The distinctions (“general” vs. “specific”, “intensional” vs. “extensional”, “set up once” vs. “continuous change”) indicate that for purposes of development, maintenance, and good design of the software system it is reasonable to distinguish between ontology and knowledge base. Also, they describe a rough shape of where to put which parts of a logical theory constraining the intended semantic models that facilitate the referencing task for human and software agents. However, the reader should note that none of these distinctions draw a clear cut borderline between ontology and knowledge base in general. Rather, it is typical that in a few percent of cases it depends on the domain, the view of the modeler, and the experience of the modeler, whether she decides to put particular entities and relations into the ontology or into the knowledge base.

Both following definitions of ontology and knowledge base specify constraints on the way an ontology (or a knowledge base) should be modeled *in a particular logical language* like F-Logic or OIL:

Definition 1 (Ontology). *An ontology is a sign system $\mathcal{O} := (\mathcal{L}, \mathcal{F}, \mathcal{G}, \mathcal{C}, \mathcal{H}, \mathcal{R}, \mathcal{A})$, which consists of*

- A **lexicon**: The lexicon contains a set of signs (lexical entries) for concepts, \mathcal{L}^c , and a set of signs for relations, \mathcal{L}^r . Their union is the lexicon $\mathcal{L} := \mathcal{L}^c \cup \mathcal{L}^r$.
- Two **reference functions** \mathcal{F}, \mathcal{G} , with $\mathcal{F} : 2^{\mathcal{L}^c} \mapsto 2^{\mathcal{C}}$ and $\mathcal{G} : 2^{\mathcal{L}^s} \mapsto 2^{\mathcal{S}}$. \mathcal{F} and \mathcal{G} link sets of lexical entries $\{L_i\} \subset \mathcal{L}$ to the set of concepts and relations they refer to, respectively, in the given ontology. In general, one lexical entry may refer to several concepts or relations and one concept or relation may be referred to by several lexical entries. Their inverses are \mathcal{F}^{-1} and \mathcal{G}^{-1} .

In order to map easily back and forth and because there is a n to m mapping between lexicon and concepts/relations, \mathcal{F} and \mathcal{G} are defined on sets rather than on single objects.

- A set \mathcal{C} of **concepts**: About each $C \in \mathcal{C}$ exists at least one statement in the ontology, viz. its embedding in the taxonomy.
- A **taxonomy** \mathcal{H} : Concepts are taxonomically related by the irreflexive, acyclic, transitive relation \mathcal{H} , ($\mathcal{H} \subset \mathcal{C} \times \mathcal{C}$). $\mathcal{H}(C_1, C_2)$ means that C_1 is a subconcept of C_2 .

- A set of binary **relations** \mathcal{R} : \mathcal{R} denotes a set of binary relations.² They specify pairs of domain and ranges (D, R) with $D, R \in \mathcal{C}$.
The functions d and r applied to a binary relation Q yield the corresponding domain and range concepts D and R , respectively.
- A set of ontology axioms, A .

The reader may note that the structure we propose is very similar to the WordNet model described by Miller [14]. WordNet has been conceived as a mixed linguistic / psychological model about how people associate words with their meaning. Like WordNet, we allow that one word may have several meanings and one concept (synset) may be represented by several words. However, we allow for a seamless integration into logical languages like OIL or F-Logic by providing very simple means for definition of relations and for knowledge bases.

We define a knowledge base as a collection of object descriptions that refer to a given ontology.

Definition 2 (Knowledge Base). We define a knowledge base as a 7-tupel $\mathcal{KB} := (\mathcal{L}, \mathcal{J}, \mathcal{I}, \mathcal{W}, \mathcal{S}, A, \mathcal{O})$, that consists of

- a **lexicon** containing a set of signs for instances, \mathcal{L} .
- A **reference function** \mathcal{J} with $\mathcal{J} : 2^{\mathcal{L}} \mapsto 2^{\mathcal{I}}$. \mathcal{J} links sets of lexical entries $\{L_i\} \subset \mathcal{L}$ to the set of instances they correspond to.
Thereby, names may be multiply used, e.g. “Athens” may be used for “Athens, Georgia” or for “Athens, Greece”.
- a set of **instances** \mathcal{I} . About each $I_k \in \mathcal{I}, k = 1, \dots, l$ exists at least one statement in the knowledge base, viz. a membership to a concept C from the ontology \mathcal{O} .
- A **membership function** \mathcal{W} with $\mathcal{W} : 2^{\mathcal{I}} \mapsto 2^{\mathcal{C}}$. \mathcal{W} assigns sets of instances to the sets of concepts they are members of.
- **Instantiated relations**, \mathcal{S} , are described, viz. $\mathcal{S} \subseteq \{(x, y, z) | x \in \mathcal{I}, y \in \mathcal{R}, z \in \mathcal{I}\}$.
- A set of knowledge base axioms, A .
- A reference to an ontology \mathcal{O} .

Overall the decision to model some relevant part of the domain in the ontology vs. in the knowledge base is often based on gradual distinctions and driven by the needs of the application. Concerning the technical issue it is sometimes even useful to let the lexicon of knowledge base and ontology overlap, e.g. to use a concept name to refer to a particular instance in a particular context. In fact researchers in natural language have tackled the question how the reference function J can be dynamically extended given an ontology, a context, a knowledge base and a particular sentence.

3 SEAL infrastructure and core modules

The aim of our intranet application is the presentation of information to human and software agents taking advantage of semantic structures. In this section, we first elaborate on the general architecture for SEAL (SEmantic PortAL), before we explain functionalities of its core modules.

² Here at the conceptual level, we do not distinguish between relations and attributes.

3.1 Architecture

The overall architecture and environment of SEAL is depicted in Figure 3:




Fig. 3. AIFB Intranet - System architecture

The *backbone* of the system consists of the *knowledge warehouse*, *i.e.* the data repository, and the *Ontobroker* system, *i.e.* the principal inferencing mechanism. The latter functions as a kind of middleware run-time system, possibly mediating between different information sources when the environment becomes more complex than it is now.

At the front end one may distinguish between three types of *agents*: *software agents*, *community users* and *general users*. All three of them communicate with the system through the *web server*. The three different types of agents correspond to three primary modes of interaction with the system.

First, remote applications (*e.g.* software agents) may process information stored at the portal over the internet. For this purpose, the *RDF generator* presents RDF facts through the web server. Software agents with *RDF crawlers* may collect the facts and, thus, have direct access to semantic knowledge stored at the web site.

Second, community users and general users can access information contained at the web site. Two forms of accessing are supported: navigating through the portal by exploiting hyperlink structure of documents and searching for information by posting queries. The hyperlink structure is partially given by the portal builder, but it may be

extended with the help of the *navigation* module. The navigation module exploits inferencing capabilities of the inference engine in order to construct conceptual hyperlink structures. Searching and querying is performed via the *query* module. In addition, the user can personalise the search interface using the *semantic personalization* preprocessing module and/or rank retrieved results according to semantic similarity (done by the postprocessing module for *semantic ranking*). Queries also take advantage of the Ontobroker inferencing.

Third, only community users can provide data. Typical information they contribute includes personal data, information about research areas, publications, activities and other research information. For each type of information they contribute there is (at least) one concept in the ontology. Retrieving parts of the ontology, the *template* module may semi-automatically produce suitable HTML forms for data input. The community users fill in these forms and the template modules stores the data in the knowledge warehouse.

3.2 Core modules

The core modules have been extensively described in [19]. In order to give the reader a compact overview we here shortly survey their function. In the remainder of the paper we delve deeper into those aspects that have been added or considerably extended recently, *viz.* semantic ranking (Section 4), and semantic access by software agents (Section 5).

Ontobroker. The Ontobroker system [6] is a deductive, object-oriented database system operating either in main memory or on a relational database (via JDBC). It provides compilers for different languages to describe ontologies, rules and facts. Beside other usage, in this architecture it is also used as an inference engine (server). It reads input files containing the knowledge base and the ontology, evaluates incoming queries, and returns the results derived from the combination of ontology, knowledge base and query.

The possibility to derive additional factual knowledge from given facts and background knowledge considerably facilitates the life of the knowledge providers and the knowledge seekers. For instance, one may specify that if a person belongs to a research group of institute AIFB, he also belongs to AIFB. Thus, it is unnecessary to specify the membership to his research group *and* to AIFB. Conversely, the information seeker does not have to take care of inconsistent assignments, *e.g.* ones that specify membership to an AIFB research group, but that have erroneously left out the membership to AIFB.

Knowledge warehouse. The knowledge warehouse [19] serves as repository for data represented in the form of F-Logic statements. It hosts the ontology, as well as the data proper. From the point of view of inferencing (Ontobroker) the difference is negligible, but from the point of view of maintaining the system the difference between ontology definition and its instantiation is useful. The knowledge warehouse is organised around a relational database, where facts and concepts are stored in a reified format. It states relations and concepts as first-order objects and it is therefore very flexible with regard to changes and amendments of the ontology.

Navigation module. Beside the hierarchical, tree-based hyperlink structure which corresponds to hierarchical decomposition of domain, the navigation module enables complex graph-based semantic hyperlinking, based on ontological relations between concepts (nodes) in the domain. The conceptual approach to hyperlinking is based on the assumption that semantic relevant hyperlinks from a web page correspond to conceptual relations, such as `memberOf` or `hasPart`, or to attributes, like `hasName`. Thus, instances in the knowledge base may be presented by automatically generating links to all related instances. For example, on personal web pages (*cf.* Figure 5) there are hyperlinks to web pages that describe the corresponding research groups, research areas and project web pages.

Query module. The query module puts an easy-to-use interface on the query capabilities of the F-Logic query interface of Ontobroker. The portal builder models web pages that serve particular query needs, such as querying for projects or querying for people. For this purpose, selection lists that restrict query possibilities are offered to the user. The selection lists are compiled using knowledge from the ontology and/or the knowledge base. For instance, the query interface for persons allows to search for people according to research groups they are members of. The list of research groups is dynamically filled by an F-Logic query and presented to the user for easy choice by a drop-down list (*cf.* snapshot in Figure 4).




Fig. 4. Query form based on definition of concept Person

Even simpler, one may apprehend a hyperlink with an F-Logic query that is dynamically evaluated when the link is hit. More complex, one may construct an `isA`, a `hasPart`, or a `hasSubtopic` tree, from which query events are triggered when particular nodes in the tree are navigated.

Personalization module. The personalization component allows to provide check-box personalization and preference-based personalization (including profiling from semantics-based log files). For instance, one may detect that user group A is particularly interested in all pages that deal with nature-analog algorithms, e.g. ones about genetic algorithms or ant algorithms.

Template module. In order to facilitate the contribution of information by community users, the template module generates an HTML form for each concept that a user may instantiate. For instance, in the AIFB intranet there is an input template (cf. Figure 5, upper left) generated from the concept definition of `person` (cf. Figure 5, lower left). The data is later on used by the navigation module to produce the corresponding person web page (cf. Figure 5, right hand side).



Fig. 5. Templates generated from concept definitions

In order to reduce the data required for input, the portal builder specifies which attributes and relations are derived from other templates. For example, in our case the

portal builder has specified that project membership is defined in the project template. The co-ordinator of a project enters information about which persons are participants of the project and this information is used when generating the person web page taking advantage of a corresponding F-Logic rule for inverse relationships. Hence, it is unnecessary to input this information in the person template.

Ontology lexicon. The different modules described here make extensive use of the lexicon component of the ontology. The most prevalent use is the distinction between English and German (realized for presentation, though not for the template module, yet). In the future we envision that one may produce more adaptive web sites making use of the explicit lexicon. For instance, we will be able to produce short descriptions when the context is sufficiently narrow, *e.g.* working with ambiguous acronyms like ASP³ or SEAL⁴

4 Semantic Ranking

This section describes the architecture component “Semantic Ranking” which has been developed in the context of our application. First, we will introduce and motivate the requirement for a ranking approach with a small example we are facing. Second, we will show how the problem of semantic ranking may be reduced to the comparison of two knowledge bases. Query results are reinterpreted as “query knowledge bases” and their similarity to the original knowledge base without axioms yields the basis for semantic ranking. Thereby, we reduce our notion of similarity between two knowledge bases to the similarity of concept pairs [23, 11].

Let us assume the following ontology:

- 1 : $\text{Person} :: \text{Object}[\text{WORKSIN} \Rightarrow \text{Project}]$.
 - 2 : $\text{Project} :: \text{Object}[\text{HASTOPIC} \Rightarrow \text{Topic}]$.
 - 3 : $\text{Topic} :: \text{Object}[\text{SUBTOPICOF} \Rightarrow \text{Topic}]$.
 - 4 : $\text{FORALL } X, Y, Z \ Z[\text{HASTOPIC} \rightarrow Y] \leftarrow X[\text{SUBTOPICOF} \rightarrow Y]$
and $Z[\text{HASTOPIC} \rightarrow X]$.
- (1)

To give an intuition of the semantic of the F-Logic statements, in line 1 one finds a concept definition for a **Person** being an **Object** with a relation **WORKSIN**. The range of the relation for this **Person** is restricted to **Project**.

Let us further assume the following knowledge base:

- 5 : $\text{KnowledgeManagement} : \text{Topic}$.
 - 6 : $\text{KnowledgeDiscovery} : \text{Topic}[\text{SUBTOPICOF} \rightarrow \text{KnowledgeManagement}]$.
 - 7 : $\text{Gerd} : \text{Person}[\text{WORKSIN} \rightarrow \text{OntoWise}]$.
 - 8 : $\text{OntoWise} : \text{Project}[\text{HASTOPIC} \rightarrow \text{KnowledgeManagement}]$.
 - 9 : $\text{Andreas} : \text{Person}[\text{WORKSIN} \rightarrow \text{TelekomProject}]$.
 - 10 : $\text{TelekomProject} : \text{Project}[\text{HASTOPIC} \rightarrow \text{KnowledgeDiscovery}]$.
- (2)

³ Active server pages *vs.* active service providers.

⁴ “SouthEast Asian Linguistics Conference” *vs.* “Conference on Simulated Evolution and Learning” *vs.* “Society for Evolutionary Analysis in Law” *vs.* “Society for Effective Affective Learning” *vs.* some other dozens — several of which are indeed relevant in our institute.

Definitions of instances in the knowledge base are syntactically very similar to the concept definition in F-Logic. In line 6 the instance `KnowledgeDiscovery` of the concept `Topic` is defined. Furthermore, the relation `SUBTOPICOF` is instantiated between `KnowledgeDiscovery` and `KnowledgeManagement`. Similarly in line 7, it is stated that `Gerd` is a `{concPerson working in OntoWise}`. Ontology axioms like given in line 4 (1) use this syntax to describe regularities. Line 4 states that if some Z has topic X and X is a subtopic of Y then Z also has topic Y .

Now, an F-Logic query may ask for all people who work in a knowledge management project by:

$$\begin{aligned} \text{FORALL } Y, Z \leftarrow Y[\text{WORKSIN} \rightarrow Z] \text{ and} \\ Z : \text{Project}[\text{HASTOPIC} \rightarrow \text{KnowledgeManagement}] \end{aligned} \quad (3)$$

which may result in the tuples $M_1^T := (\text{Gerd}, \text{OntoWise})$ and $M_2^T := (\text{Andreas}, \text{TelekomProjekt})$. Obviously, both answers are correct with regard to the given knowledge base and ontology, but the question is, what would be a plausible ranking for the correct answers. This ranking should be produced from a given query without assuming any modification of the query.

4.1 Reinterpreting queries

Our principal consideration builds on the definition of semantic similarity that we have first described in [23, 11]. There, we have developed a measure for the similarity of two knowledge bases. Here, our basic idea is to reinterpret possible query results as a “query knowledge base” and compute its similarity to the original knowledge base while abstracting from semantic inferences. The result of an F-Logic query may be re-interpreted as a *query knowledge base* (QKB) by the following approach.

An F-Logic query is of the form or can be rewritten into the form⁵:

$$\text{FORALL } \overline{X} \leftarrow \overline{P}(\overline{X}, \overline{k}), \quad (4)$$

with \overline{X} being a vector of variables (X_1, \dots, X_n) , \overline{k} being a vector of constants, and \overline{P} being a vector of conjoined predicates. The result of a query is a two-dimensional matrix M of size $m \times n$, with n being the number of result tuples and m being the length of \overline{X} and, hence, the length of the result tuples. Hence, in our example above $\overline{X} := (Y, Z)$, $\overline{k} := (\text{'knowledge management'})$, $\overline{P} := (P_1, P_2)$, $P_1(a, b, c) := a[\text{WORKSIN} \rightarrow b]$, $P_2(a, b, c) := b[\text{HASTOPIC} \rightarrow c]$ and

$$M := (M_1, M_2) = \begin{pmatrix} \text{Gerd} & \text{Andreas} \\ \text{OntoWise} & \text{TelekomProjekt} \end{pmatrix}. \quad (5)$$

Now, we may define the query knowledge base i (QKB_i) by

$$QKB_i := \overline{P}(M_i, \overline{k}). \quad (6)$$

⁵ Negation requires special treatment.

The similarity measure between the query knowledge base and the given knowledge base may then be computed in analogy to [23]. An adaptation and simplification of the measures described there is given in the following together with an example.

4.2 Similarity of knowledge bases

The similarity between two objects (concepts and or instances) may be computed by considering their relative place in a common hierarchy H . H may, but need not be a taxonomy \mathcal{H} . For instance, in our example from above we have a categorization of research topics, which is not a taxonomy!

Our principal measures are based on the cotopies of the corresponding objects as defined by a given hierarchy H , e.g. an ISA hierarchy \mathcal{H} , an part-whole hierarchy, or a categorization of topics. Here, we use the *upwards cotopy* (UC) defined as follows:

$$\text{UC}(O_i, H) := \{O_j | H(O_i, O_j) \vee O_j = O_i\} \quad (7)$$

UC is overloaded in order to allow for a set of objects M as input instead of only single objects, *viz.*

$$\text{UC}(M, H) := \bigcup_{O_i \in M} \{O_j | H(O_i, O_j) \vee O_j = O_i\} \quad (8)$$

Based on the definition of the upwards cotopy (UC) the object match (OM) is defined by:

$$\text{OM}(O_1, O_2, H) := \frac{|\text{UC}(O_1, H) \cap \text{UC}(O_2, H)|}{|\text{UC}(O_1, H) \cup \text{UC}(O_2, H)|}. \quad (9)$$

Basically, OM reaches 1 when two concepts coincide (number of intersections of the respective upwards cotopies and number of unions of the respective cotopies is equal); it degrades to the extent to which the discrepancy between intersections and unions increases (a OM between concepts that do not share common superconcepts yields value 0).

Example. We here give a small example for computing UC and OM based on a given categorization of objects H . Figure 6 depicts the example scenario. The upwards cotopy $\text{UC}(\text{knowledge discovery}, H)$ is given by $\{\text{knowledge discovery}, \text{knowledge management}\}$. The upwards cotopy $\text{UC}(\text{optimization}, H)$ computes to $\{\text{optimization}\}$.

Computing the object match OM between `KnowledgeManagement` and `Optimization` results in 0, the object match between `KnowledgeDiscovery` and `CSCW` computes to $\frac{1}{3}$.

The match introduced above may easily be generalized to relations using a relation hierarchy H_R . Thus, the predicate match (PM) for two n-ary predicate P_1, P_2 is defined by a mean value. Thereby, we use the geometric mean in order to reflect the intuition that if the similarity of one of the components approaches 0 the overall similarity between two predicates should approach 0 — which need not be the case for the arithmetic mean:




Fig. 6. Example for computing UC and OM

$$PM(P_1(I_1, \dots, I_n), P_2(J_1, \dots, J_n)) := \sqrt[n+1]{OM(P_1, P_2, \mathcal{H}_R) \cdot OM(I_1, J_1, H) \cdot \dots \cdot OM(I_n, J_n, H)}. \quad (10)$$

This result may be averaged over an array of predicates. We here simply give the formula for our actual needs, where a query knowledge base is compared against a given knowledge base KB:

$$Simil(QKB_i, KB) = Simil(\overline{P}(M_i, \overline{k}), KB) := \frac{1}{|\overline{P}|} \sum_{P_j \in \overline{P}} \max_{Q(M_i, \overline{k}) \in KB.S} PM(P_j(M_i, \overline{k}), Q(M_i, \overline{k})). \quad (11)$$

For instance, comparing the two result tuples from our example above with the given knowledge base: First, $M_1^T := (Gerd, OntoWise)$. Then, we have the query knowledge base (QKB_1):

$$\begin{aligned} &Gerd[WORKSIN \rightarrow OntoWise]. \\ &OntoWise[HASTOPIC \rightarrow KnowledgeManagement]. \end{aligned} \quad (12)$$

and its relevant counterpart predicates in the given knowledge base (KB) are:

$$\begin{aligned} &Gerd[WORKSIN \rightarrow OntoWise]. \\ &OntoWise[HASTOPIC \rightarrow KnowledgeManagement]. \end{aligned} \quad (13)$$

This is a perfect fit. Therefore $Simil(QKB_1, KB)$ computes to 1.

Second, $M_2^T := (Andreas, TelekomProject)$. Then, we have the query knowledge base (QKB_2):

$$\begin{aligned} &Andreas[WORKSIN \rightarrow TelekomProject]. \\ &TelekomProject[HASTOPIC \rightarrow KnowledgeManagement]. \end{aligned} \quad (14)$$

and its relevant counterpart predicates in the given knowledge base (KB) are:

$$\begin{aligned} &Andreas[WORKSIN \rightarrow TelekomProject]. \\ &TelekomProject[HASTOPIC \rightarrow KnowledgeDiscovery]. \end{aligned} \quad (15)$$

Hence, the similarity of the first predicates indicates a perfect fit and evaluates to 1, but the congruency of $TelekomProject[HASTOPIC \rightarrow KnowledgeManagement]$ with

`TelekomProject[HAS TOPIC → KnowledgeDiscovery]` measures less than 1. The instance match of `KnowledgeDiscovery` and `KnowledgeManagement` returns $\frac{1}{2}$ in the given topic hierarchy. Therefore, the predicate match returns $\sqrt[3]{1 \cdot 1 \cdot \frac{1}{2}} \approx 0.79$. Thus, overall ranking of the second result is based on $\frac{1}{2}(1 + 0.79) = 0.895$.

Remarks on semantic ranking. The reader may note some basic properties of the ranking: (i) similarity of knowledge bases is an asymmetric measure, (ii) the ontology defines a conceptual structure useful for defining similarity, (iii) the core concept for evaluating semantic similarity is cotopy defined by a dedicated hierarchy. The actual computation of similarity depends on which conceptual structures (e.g. hierarchies like taxonomy, part-whole hierarchies, or topic hierarchies) are selected for evaluating conceptual nearness. Thus, similarity of knowledge bases depends on the view selected for the similarity measure.

Ranking of semantic queries using underlying ontological structures is an important means in order to allow users a more specific view onto the underlying knowledge base. The method that we propose is based on a few basic principles:

- Reinterpret the combination of query and results as query knowledge bases that may be compared with the explicitly given information.
- Give a measure for comparing two knowledge bases, thus allowing rankings of query results.

Thus, we may improve the interface to the underlying structures without changing the basic architecture. Of course, the reader should be aware that our measure may produce some rankings for results that are hardly comparable. For instance, results may differ slightly because of imbalances in a given hierarchy or due to rather random differences of depth of branches. In this case, ranking may perhaps produce results that are not better than unranked ones — but the results will not be any worse either.

5 RDF outside — From a Semantic Web Site to the Semantic Web

In the preceding sections we have described the development and the underlying techniques of the AIFB semantic web site. Having developed the core application we decided that RDF-capable software agents should be able to understand the content of application. Therefore, we have built an automatic RDF GENERATOR that dynamically generates RDF statements on each of the static and dynamic pages of the semantic knowledge portal. Our current AIFB intranet application is “Semantic Web-ized” using RDF facts instantiated and defined according to the underlying AIFB ontology. On top of this generated and formally represented metadata, there is the RDF CRAWLER, a tool that gathers interconnected fragments of RDF from the internet.

5.1 RDF GENERATOR — an example

The RDFMAKER established in the ONTOBROKER framework (cf. [5]) was a starting point for building the RDF GENERATOR. The idea of RDFMAKER was, that from ONTOBROKER’S internal data base, RDF statements are generated.

RDF GENERATOR follows a similar approach and extends the principal ideas. In a first step it generates an RDF(S)-based ontology that is stored on a specific XML namespace, *e.g.* in our concrete application

<http://ontobroker.semanticweb.org/ontologies/aifb-onto-2001-01-01.rdfs>. Additionally, it queries the knowledge warehouse. Data, *e.g.* for a person, is checked for consistency, and, if possible, completed by applying the given F-Logic rules. We here give a short example of what type of data may be generated and stored on a specific homepage of a researcher:

```
<rdf:RDF
  xmlns:rdf = "http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:aifb = "http://ontobroker.semanticweb.org/aifb-2001-01-01.rdfs#">

  <aifb:PhDStudent rdf:ID="per:ama">
    <aifb:name>Alexander Maedche</aifb:name>
    <aifb:email>ama@aifb.uni-karlsruhe.de</aifb:email>
    <aifb:phone>+49- (0) 721-608 6558</aifb:phone>
    <aifb:fax>+49- (0) 721-608 6580</aifb:fax>
    <aifb:homepage>http://www.aifb.uni-karlsruhe.de/WBS/ama</aifb:homepage>
    <aifb:supervisor
      rdf:resource = "http://www.aifb.uni-karlsruhe.de/studer.html#per:rst"/>
    </aifb:PhDStudent>
</rdf:RDF>
```

RDF GENERATOR is a configurable tool, in some cases one may want to use inferences to generate materialized, complete RDF descriptions on a home page, in other cases one may want to generate only ground facts of RDF. Therefore, RDF GENERATOR allows to switch axioms on and off in order to adopt the generation of results to varying needs.

5.2 RDF CRAWLER

The RDF CRAWLER⁶ is a tool which downloads interconnected fragments of RDF from the internet and builds a knowledge base from this data. Building an external knowledge base for the whole AIFB (its researcher, its projects, its publications, ...) becomes easy using the RDF CRAWLER and machine-processable RDF data currently defined on AIFB's web. We here shortly describe the underlying techniques of our RDF CRAWLER and the process of building a knowledge base. In general, RDF data may appear in Web documents in several ways. We distinguish between pure RDF (files that have an extension like "*.rdf"), RDF embedded in HTML and RDF embedded in XML. Our RDF CRAWLER uses RDF-API⁷ that can deal with different embeddings of RDF described above.

One problem of crawling is the applied filtering mechanism: Baseline crawlers are typically restricted by a given depth value. Recently several new research work on so-called *focused crawling* has been published (*e.g. cf.* [3]). In their approach, they use a set of predefined documents associated with topics in a Yahoo like taxonomy to build a focused crawler. Two hypertext mining algorithms constitute the core of their approach.

⁶ RDF CRAWLER is freely available for download at <http://ontobroker.semanticweb.org/rdfcrawler>.

⁷ RDF-API is freely available at <http://www-db.stanford.edu/~melnik/rdf/api.html>.

A classifier evaluates the relevance of a hypertext document with respect to the focus topics and a distiller identifies hypertext nodes that are good access points to many relevant pages within a few links. In contrast, our approach uses ontological background knowledge to judge the relevance of each page. If a page is highly relevant, the crawler may follow the links on the particular web site. If RDF data is available on a page, we judge relevance with respect to the quantity and quality of available data and by the existing URI's.

Example: Erdoes numbers. As mentioned above we here give a small example of a nice application that may be easily built using RDF metadata taken from AIFB using the RDF CRAWLER. The so-called *Erdoes numbers* have been a part of the folklore of mathematicians throughout the world for many years⁸.

Scientific papers are frequently published with co-authors. Based on information about collaboration one may compute the Erdoes number (denoted $PE(R)$) for a researcher R . In the AIFB web site the RDF-based metadata allows for computing estimates of Paul Erdoes numbers of AIFB members. The numbers are defined recursively:

1. $PE(R) = 0$, iff R is Paul Erdoes
2. $PE(R) = \min\{PE(R_1) + 1\}$ else, where R_1 varies over the set of all researchers who have collaborated with R , *i.e.* have written a scientific paper together.

To put this into work, we need lists of publications annotated with RDF facts. The lists may be automatically generated by the RDF GENERATOR. Based on the RDF facts one may crawl relevant information into a central knowledge base and compute these numbers from the data.

6 Related work

This section positions our work in the context of existing web portals and also relates our work to other basic methods and tools that are or could be deployed for the construction of community web portals, especially to related work in the area of semantic ranking of query results.

Related Work on Knowledge Portals. One of the well-established web portals on the web is Yahoo⁹. In contrast to our approach Yahoo only utilizes a very light-weight ontology that solely consists of categories arranged in a hierarchical manner. Yahoo offers keyword search (local to a selected topic or global) in addition to hierarchical navigation, but is only able to retrieve complete documents, *i.e.* it is not able to answer queries concerning the contents of documents, not to mention to combine facts being found in different documents or to include facts that could be derived through ontological axioms. Personalization is limited to check-box personalization. We get rid of these shortcomings since our portal is built upon a rich ontology enabling the portal to

⁸ The interested reader may have a look at <http://www.oakland.edu/~grossman/erdoshp.html> for an overall project overview.

⁹ <http://www.yahoo.com>

give integrated answers to queries. Furthermore, our semantic personalization features provide more flexible means for adapting the portal to the specific needs of its users.

A portal that is specialized for a scientific community has been built by the Math-Net project [4]. At <http://www.math-net.de/> the portal for the (German) mathematics community is installed that makes distributed information from several mathematical departments available. This information is accompanied by meta-data according to the Dublin Core¹⁰ Standard [25]. The Dublin Core element “Subject” is used to classify resources as conferences, as research groups, as preprints etc. A finer classification (*e.g.* via attributes) is not possible except for instances of the publication category. Here the common MSC-Classification¹¹ is used that resembles a light-weight ontology of the field of mathematics. With respect to our approach Math-Net lacks a rich ontology that could enhance the quality of search results (*esp.* via inferencing), and the smooth connection to the Semantic Web world that is provided by our RDF generator.

The Ontobroker project [5] lays the technological foundations for the AIFB portal. On top of Ontobroker the portal has been built and organizational structures for developing and maintaining it have been established. Therefore, we compare our system against approaches that are similar to Ontobroker.

The approach closest to Ontobroker is SHOE [7]. In SHOE, HTML pages are annotated via ontologies to support information retrieval based on semantic information. Besides the use of ontologies and the annotation of web pages the underlying philosophy of both systems differs significantly: SHOE uses description logic as its basic representation formalism, but it offers only very limited inferencing capabilities. Ontobroker relies on Frame-Logic and supports complex inferencing for query answering. Furthermore, the SHOE search tool neither provides means for a semantic ranking of query results nor for a semantic personalization feature. A more detailed comparison to other portal approaches and underlying methods may be found in [19].

Related Work on Semantic Similarity. Since our semantic ranking is based on the comparison of the query knowledge base with the given ontology and knowledge base, we relate our work to the comparison of ontological structures and knowledge bases (covering the same domain) and to measuring the similarity between concepts in a hierarchy. Although there has been a long discussion in the literature about evaluating knowledge-bases [13], we have not found any discussion about comparing two knowledge bases covering the same domain that corresponds to our semantic ranking approach. Similarity measures for ontological structures have been investigated in areas like cognitive science, databases or knowledge engineering (*cf. e.g.*, [17, 16, 18, 9]). However, all these approaches are restricted to similarity measures between lexical entries, concepts, and template slots within one ontology.

Closest to our measure of similarity is work in the NLP community, named semantic similarity [17] which refers to similarity between two concepts in a *isA*-taxonomy such as the WordNet or CYC upper ontology. Our approach differs in two main aspect from this notion of similarity: Firstly, our similarity measure is applicable to a hierarchy which may, but not need be a taxonomy and secondly it is taking into account not

¹⁰ <http://www.purl.org/dc>

¹¹ *cf.* Mathematical Subject Classification; <http://www.ams.org/msc/>

only commonalities but also differences between the items being compared, expressing both in semantic-cotopy terms. This second property enables the measuring of self-similarity and subclass-relationship similarity, which are crucial for comparing results derived from the inferencing processes, that are executed in the background.

Conceptually, instead of measuring similarity between isolated terms (words), that does not take into account the relationship among word senses that matters, we measure similarity between “words in context”, by measuring similarity between Object-Attribute-Value pairs, where each term corresponds to a concept in the ontology. This enables us to exploit the ontological background knowledge (axioms and relations between concepts) in measuring the similarity, which expands our approach to a methodology for comparing knowledge bases.

From our point of view, our community portal system is rather unique with respect to the collection of methods used and the functionality provided. We have extended our community portal approach that provides flexible means for providing, integrating and accessing information [19] by semantic personalization features, semantic ranking of generated answers and a smooth integration with the evolving Semantic Web. All these methods are integrated into one uniform system environment, the SEAL framework.

7 Conclusion

In this paper we have shown our comprehensive approach SEAL for building semantic portals. In particular, we have focused on three issues.

First, we have considered the ontological foundation of SEAL. There, we have made the experience that there are many big open issues that have hardly been dealt with so far. In particular, the step of formalizing the ontology raises very principal problems. The issue of where to put relevant concepts, *viz.* into the ontology *vs.* into the knowledge base, is an important one that deeply affects organization and application. However, there exist no corresponding methodological guidelines to base the decision upon so far. For instance, we have given the example ontology and knowledge base in (1) and (2). Using description logics terminology, we have equated the ontology with the “T-Box” and we have put the topic hierarchy into the knowledge base (“A-Box”). An alternative could have been to formalize the topic hierarchy as an *isa*-hierarchy, which however it isn’t and put it into the T-Box. We believe that both alternatives exhibit an internal fault, *viz.* the ontology should not be equated with the T-Box, but rather should its scope be independent from an actual formalization with particular logical statements. Its scope should to a large extent depend on soft issues, like “Who updates a concept?” and “How often does a concept change?” such as already indicated in Table 1. Second, we have described the general architecture of the SEAL approach, which is also used for our real-world case study, the AIFB web site. The architecture integrates a number of components that we have also used in other applications, like Ontobroker, navigation or query module. Third, we have extended our semantic modules to include a larger diversity of intelligent means for accessing the web site, *viz.* semantic ranking and machine access by crawling.

For the future, we see a number of new important topics appearing on the horizon. For instance, we consider approaches for ontology learning [12] in order to semi-

automatically adapt to changes in the world and to facilitate the engineering of ontologies. Currently, we work on providing intelligent means for providing semantic information, *i.e.* we elaborate on a semantic annotation framework that balances between manual provisioning from legacy texts (*e.g.* web pages) and information extraction [22]. Given a particular conceptualization, we envision that one wants to be able to use a multitude of different inference engines taking advantage of different inferencing capabilities (temporal, non-monotonic, high scalability, etc.). Then, however, one needs means to change from one representation paradigm to the next [20].

Finally, we envision that once semantic web sites are widely available, their automatic exploitation may be brought to new levels. Semantic web mining considers the level of mining web site structures, web site content, and web site usage on a semantic rather than at a syntactic level yielding new possibilities, *e.g.* for intelligent navigation, personalization, or summarization, to name but a few objectives for semantic web sites [8].

Acknowledgements. The research presented in this paper would not have been possible without our colleagues and students at the Institute AIFB, University of Karlsruhe, and Ontoprise GmbH. We thank Jürgen Angele, Kalvis Apsitis (now: RITI Riga Information Technology Institute), Nils Braeunlich, Stefan Decker (now: Stanford University), Michael Erdmann, Dieter Fensel (now: VU Amsterdam), Siegfried Handschuh, Andreas Hotho, Mika Maier-Collin, Daniel Oberle, and Hans-Peter Schnurr. Research for this paper was partially financed by Ontoprise GmbH, Karlsruhe, Germany, by US Air Force in the DARPA DAML project “OntoAgents”, by EU in the IST-1999-10132 project “On-To-Knowledge” and by BMBF in the project “GETESS” (01IN901C0).

References

1. J. Angele, H.-P. Schnurr, S. Staab, and R. Studer. The times they are a-changin’ — the corporate history analyzer. In D. Mahling and U. Reimer, editors, *Proceedings of the Third International Conference on Practical Aspects of Knowledge Management. Basel, Switzerland, October 30-31, 2000*, 2000. <http://www.research.swisslife.ch/pakm2000/>.
2. V. Richard Benjamins and Dieter Fensel. Community is knowledge! (KA)². In *Proceedings of the 11th Workshop on Knowledge Acquisition, Modeling, and Management (KAW ’98), Banff, Canada, April 1998*, 1998.
3. S. Chakrabarti, M. van den Berg, and B. Dom. Focused crawling: a new approach to topic-specific web resource discovery. In *Proceedings of WWW-8*, 1999.
4. W. Dalitz, M. Grötschel, and J. Lügger. Information Services for Mathematics in the Internet (Math-Net). In A. Sydow, editor, *Proceedings of the 15th IMACS World Congress on Scientific Computation: Modelling and Applied Mathematics*, volume 4 of *Artificial Intelligence and Computer Science*, pages 773–778. Wissenschaft und Technik Verlag, 1997.
5. S. Decker, M. Erdmann, D. Fensel, and R. Studer. Ontobroker: Ontology Based Access to Distributed and Semi-Structured Information. In R. Meersman et al., editors, *Database Semantics: Semantic Issues in Multimedia Systems*, pages 351–369. Kluwer Academic Publisher, 1999.
6. D. Fensel, S. Decker, M. Erdmann, and R. Studer. Ontobroker: The Very High Idea. In *Proceedings of the 11th International Flairs Conference (FLAIRS-98), Sanibel Island, Florida, May, 1998*.

7. J. Heflin and J. Hendler. Searching the web with shoe. In *Artificial Intelligence for Web Search. Papers from the AAAI Workshop. WS-00-01*, pages 35–40. AAAI Press, 2000.
8. A. Hotho and G. Stumme, editors. *Semantic Web Mining — Workshop at ECML-2001 / PKDD-2001, Freiburg, Germany*, 2001.
9. E. Hovy. Combining and standardizing large-scale, practical ontologies for machine translation and other uses. In *Proc. of the First Int. Conf. on Language Resources and Evaluation (LREC)*, 1998.
10. M. Kifer, G. Lausen, and J. Wu. Logical Foundations of Object-Oriented and Frame-Based Languages. *Journal of the ACM*, 42:741–843, 1995.
11. A. Maedche and S. Staab. Discovering conceptual relations from text. In *Proceedings of ECAI-2000*. IOS Press, Amsterdam, 2000.
12. A. Maedche and S. Staab. Ontology learning for the semantic web. *IEEE Intelligent Systems*, 16(2), 2001.
13. T.J. Menzis. Knowledge maintenance: The state of the art. *The Knowledge Engineering Review*, 10(2), 1998.
14. G. Miller. Wordnet: A lexical database for English. *CACM*, 38(11):39–41, 1995.
15. C.K. Odgen and I.A. Richards. *The Meaning of Meaning: A Study of the Influence of Language upon Thought and of the Science of Symbolism*. Routledge & Kegan Paul Ltd., London, 10 edition, 1923.
16. R. Rada, H. Mili, E. Bicknell, and M. Blettner. Development and application of a metric on semantic nets. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(1), 1989.
17. P. Resnik. Knowledge maintenance: The state of the art. In *Proceedings of IJCAI-95*, pages 448–453, Montreal, Canada, 1995.
18. R. Richardson, A. F. Smeaton, and J. Murphy. Using wordnet as knowledge base for measuring semantic similarity between words. Technical Report CA-1294, Dublin City University, School of Computer Applications, 1994.
19. S. Staab, J. Angele, S. Decker, M. Erdmann, A. Hotho, A. Maedche, H.-P. Schnurr, R. Studer, and Y. Sure. Semantic community web portals. *Proc. of WWW9 / Computer Networks*, 33(1-6):473–491, 2000.
20. S. Staab, M. Erdmann, and A. Maedche. Engineering ontologies using semantic patterns. In A. Preece, editor, *Proc. of the IJCAI-01 Workshop on E-Business & the Intelligent Web*, 2001.
21. S. Staab and A. Maedche. Knowledge portals — ontologies at work. *AI Magazine*, 21(2), Summer 2001.
22. S. Staab, A. Maedche, and S. Handschuh. An annotation framework for the semantic web. In *Proceedings of the First Workshop on Multimedia Annotation, Tokyo, Japan, January 30-31, 2001*, 2001.
23. S. Staab, A. Maedche, and S. Handschuh. Creating metadata for the semantic web: An annotation framework and the human factor. Technical Report 412, Institute AIFB, University of Karlsruhe, 2001.
24. Y. Sure, A. Maedche, and S. Staab. Leveraging corporate skill knowledge - From ProPer to OntoProper. In D. Mahling and U. Reimer, editors, *Proceedings of the Third International Conference on Practical Aspects of Knowledge Management. Basel, Switzerland, October 30-31, 2000*, 2000. <http://www.research.swisslife.ch/pakm2000/>.
25. S. Weibel, J. Kunze, C. Lagoze, and M. Wolf. *Dublin Core Metadata for Resource Discovery*. Number 2413 in IETF. The Internet Society, September 1998.