Yellow Pages on the Semantic Web

Eero Hyvönen[1,2], Kim Viljanen[2,1], and Antti J. Hätinen[1]

[1] University of Helsinki, Department of Computer Science
firstname.lastname@cs.helsinki.fi,
http://www.cs.helsinki.fi/group/seco/
[2] Helsinki Institute for Information Technology (HIIT)

Abstract. Yellow pages catalogs and corresponding directory services on the web are a widely used business concept for helping people to find companies providing services and selling products. When on the web, matching the customer's need with the relevant services offerred by companies is typically based on keyword search, table-based search, a list of service categories listed in some order, a hierarchical category system, or a combination of these. In spite of the versatility of possibilities, it can still be difficult to the end-user to map a need on the services offered. On the other hand, for the catalog advertiser, it may be difficult to index the service in such a way that the prospects would not miss the service. We propose that in order to enhance the recall and precision of yellow page services, the advertisements should be annotated using semantic web ontologies. Based on such conceptual definitions, the user can be provided with new content-based searching and browsing facilities, which makes the service more profitable to the advertisers and the directory service provider. As a first step towards this goal, an experimental semantic yellow page implementation is presented.

1 Introduction

Yellow pages is a widely used and commercially successful service model for matching the need of a customer with the corresponding products and services offered by companies. In this model, the yellow pages service provider maintains a list or a hierarchy of product and service categories, such as "Electronic equipment" or "Car Rental". The provider's customers are companies that buy advertisement space in the catalog within these categories. The business logic of yellow pages is based on the advertising companies' belief on the fact that the catalog is extensively used by their prospects in finding products and services. Yellow page catalogs are an especially important marketing channel for small companies that cannot offer large marketing campaigns but still need get in contact with a large potential customer base. Yellow pages are traditionally published in printed form once a year as a part of local telephone books. On the WWW, lots of similar directory services are provided online.

In this article, we discuss how Semantic Web [9] technologies can be employed in yellow page services. Our idea is to design ontologies [10] classifying the yellow page contents from the usage point of view. By annotating the advertisements

with metadata based on such ontologies, content-based search functions and service recommendation can be provided to the end-user. We focus on the problem of how to match the customer need with what the companies offer in their advertisements. The Helsinki area yellow pages[3] are used as a test case.

In the following, the current state of the art of yellow page services is first outlined. Section 3 presents our vision of how the semantic web technologies can help to solve problems in current online yellow page services. In section 4 we describe our first prototype of a semantic yellow page service. Section 5 summarizes the work and discusses some open problems.

2 Yellow pages: state of the art

Printed yellow pages The Helsinki area yellow pages (HYP) (2002) is a typical example of a yellow page catalog. HYP is based on an alphabetical category list of 1085 products and services. From the user's viewpoint, a problem of using the catalog is that a user's need may be related to several advertisement categories. For example, if one needs help for a headache, the possibly relevant categories in HYP include "pharmacy shops" (in Finnish: apteekkeja), "medical consultancy" (lääkeneuvontaa), "doctors" (lääkäreitä), "medical centers" (lääkärikeskuksia), "medical labs" (lääketieteellisiä laboratorioita), "hospitals" (sairaaloita), perhaps even "nursing services" (sairaanhoitoalan palveluja). On the other hand, a category may be related to several needs. For example, "building equipment" (rakennustarvikkeita) is potentially relevant in many building needs and problems. Furthermore, several categories are partly overlapping, such as "barber's shops" (parturit) and "hair dressers" (kampaamot). In the bilingual HYP, there is also the additional problem that many, but not all, categories have both Finnish and Swedish categories. For example, printing services are advertised under the Finnish label "painopalveluja" and the Swedish term "tryckerier".

From the advertiser's viewpoint, the advertisement should be found in as many reasonable ways as possible without bothering all possible customer needs, categorizations or linguistic terminology issues. In the current category system, this not possible but the company may have to buy advertisement space in several categories for better coverage. This is not always feasible to the company and leads to the situation where many useful categories in the list are practically empty. This is not what the user expects from a good service. For example, the category "garden furniture" (puutarhakalusteita) in HYP contains only one advertisement although many, if not most furniture shops sell also garden furniture. From the user's and advertiser's viewpoint, the challenge is to create a search mechanism with high recall (relevant advertisements found)[4] and high precision (irrelevant advertisements not found)[5].

[3] http://www.keltaisetsivut.fi/

[4] Recall r is defined as the ratio $r = f/a$, where f is the number of retrieved relevant documents and a the number of all relevant documents. [4]

[5] Precision p is defined as the ratio $p = f/n$, where f is the number of retrieved relevant documents and n the number of all retrieved documents. [4]

Online yellow pages Yellow page services online can benefit the user in providing 1) new information content (e.g., hyperlinks to up-to-date business web sites), 2) better searchability of contents, and 3) interactive features (e.g., online enquiry forms and feedback mechanisms) [2]. A typical[6] online yellow page service provides the user with the following functions:

Keyword and table-based search Keyword search is used for finding the category names, the advertisements, and product names (e.g., in Brabys.com). The search may be multi-lingual (e.g., in web.wlwonline.de). Some services, such as SuperPages.com and Keltaisetsivut.fi, provides the user with table-based search along various data fields, such as the business name and location. The search result is usually presented as a hit list of matching advertisements that may be grouped according to their categories (e.g., in SuperPages.com).

Hierarchical category navigation Some services provide the user with a hierarchical navigation facility by which the user can focus from top categories to subcategories and finally to the advertisements (e.g., Brabys.com).

Location services Most of the online directory services provide location-based services that typically include maps and driving directions for finding to a given place.

Editorial content An example of editorial content is the Yell.com's shopping guide for buying a car. It provides links to related categories which may be in the interest of the user, such as car dealers, driving schools, and garage services. In this paper, we will call such semantic associations *recommended links*. In SuperPages.com, recommended links associate a selected category to an other category and in SMARTpages.com pop-up banner windows show advertisements related to the selected category.

Search engines vs. web catalogs An increasing number of companies have their own web sites. Search engines, such as Google[7] and AltaVista[8], are getting ever better in finding relevant information. This means a threat to the well-established business model of yellow pages. Using search engines is appealing for several reasons:

Matching the need with the offer It is easy for the prospect to try to locate companies by just typing in a few keywords in a search engine. The prospect does not have to figure out what categories in the catalog may be relevant from the point of view of his need. For example, if my video camera needs to be repaired I do not have to know in what yellow page categories there may be relevant advertisements for my problem. This mental matching problem is

[6] We investigated the following representative sites: http://www.smartpages.com/, http://www.superpages.com/, http://uk.yell.com/, http://www.brabys.com/, http://web.wlwonline.de/, http://www.keltaisetsivut.fi/, http://www.yritystele.fi/.

[7] http://www.google.com

[8] http://www.altavista.com

especially severe if the prospect is not familiar with the cataloging terminology and business conventions in use. For example, different products/services may be sold by different kind of companies in different countries.

Coverage of the catalog A yellow pages catalog lists only its customer's advertisements, but a general search engine may find also companies not listed in the catalog.

Quantity of information The data on the web sites is more abundant than on the catalog advertisements. For example, full product information with price lists may be readily available there.

Dynamicality Dynamicality of web pages is far beyond the static advertisements used in many catalogs. For example, order forms may be filled and returned electronically, videos clips may be used for showing the products, etc.

Timeliness The data is usually more up-to-date on the WWW than in the catalog that is updated only now and then. For example, there may be special offers for the week, updated price lists etc. on the web.

The keyword-based search methods of search engines are often useful, but have also their shortcomings [10, 6]:

1. A keyword in a document does not necessary mean that the document is relevant. For an extreme example, an advertisement may contain the phrase "We do not sell *stamps*, but ...", which means that a search using the keyword *stamp* will include the page in the hit list.
2. The engines cannot differentiate between *synonyms*. For example, a service selling "personal computers" is not found using the keyword "PC". This lowers the recall rate of information retrieval.
3. The engines do not understand *homonyms*[9]. For example, the keyword "Nokia" would find not only pages related to the Finnish telecommunication company, but also pages related to a city in Finland. This leads to low precision in information retrieval.
4. The engines do not understand general terms or phrases. For example, if your are interested in finding out "accommodation" you do not find advertisements with the words "hotels" or "inns" unless also the word "accommodation" is present.
5. Relevance. It is difficult to evaluate the relevance of a document with respect to a query. A list of 1000 hits is not very useful unless they can be ordered according to their relevance to the user.
6. Implicit information. A textual description is found only if it contains the explicit keyword. For example, one may be interested in companies dealing with "astronomy". A telescope advertisement is not found unless it happens to mention the word "astronomy", which may be too obvious to be mentioned.
7. Images and other non-textual binary documents cannot be matched.

[9] A homonym is a word with several meanings.

A benefit of the catalog is that the data there is structured, which makes the catalog easier to search than a collection of semi-structured WWW pages. Furthermore, yellow page services typically contains lots of small companies, such as barber shops, restaurants etc. that do not (yet) usually have their own web site. Finding such companies with a general purpose search engine like Google is not easy. Still another benefit of yellow page catalogs is that the service provider guarantees the correctness and reliability of the information in the database.

According to the semantic web vision [5], the search for specific information on the (semantic) web will become easier with the help of semantic metadata annotations. However, the need for trusted centralized (semantic) directories and (semantic) search engines stays. Yellow page services can potentially become one of the trusted information providers of the future semantic web.

3 A vision for semantic yellow pages

Service ontology to reflect the user's needs From the viewpoint of the user's need, indexing a company's advertisements according to a business category is not very useful unless the category unambiguously implies the services offered to the user. For example, camera repair service can potentially be offered by an importer company, appliances shop, camera shop, foto shop, optician etc. What kind of companies offer repair services depends, e.g., on the service business at hand, on the thing and brand being repaired, and on local practices. In order to enhance the search capabilities of yellow pages the advertising companies should more clearly state what services they actually offer.

A service cannot be characterized extensively by a simple category label, because services have internal semantical structure. For example, transportation services are characterized by properties such as the cargo (people, furniture, oil, etc.), the instrument used (car, train, ship, airplane, etc.), and the area of transportation (city, country, international, etc.). The internal properties of services can be described by using *ontologies* [10, 3, 6]. Services are often related to each other and make combined services. A classical example is the service needed for making a conference trip, which involves using related services such as buying flight tickets, reserving a hotel room, and registering for the conference. Service combination can be described in terms of ontological structures.

In our view, yellow page online services should be based on ontologies, that define the terminology and concepts to be used by the advertiser and by the end-user[10]:

Ontology-based service annotation From the advertiser's viewpoint, each advertisement in the system is a metadata description of some service that the company offers. A single company may offer many services. It is not enough for a camera shop to put one advertisement under category "camera shops" but, for example, one advertisement under camera "selling services",

[10] This idea has been applied before in other domains, e.g., in photo annotation and retrieval [18, 12]

another under "repair services", and a third one under "rental services" depending on the company at hand. The user's viewpoint is central, and the offerings should be indexed accordingly.

Ontology-based service retrieval To the end-user, the service ontology provides new "content-based" possibilities to finding advertised services. The ontology can be used as a multidimensional semantic navigation directory for browsing the advertisements in the spirit of Topic Maps [17]. For example, services related to the concept "Piano", such as "Piano shops", "Piano transports", "Piano tuners" etc. can be found through the product class "Piano", because this concept is used in different roles in these services.

The ontology also provides a kind of semantic glue that relates the service advertisements of different service classes with each other. If you want buy a flight ticket to Helsinki, you are likely to be interested in booking hotel accommodation there as well. Furthermore, if there are some special events, such as musical festivals taking place near by at the same time, you probably would not mind being informed about them as well. Such semantic relations between service classes can be defined in general terms within the ontology, and be used for generating recommended links for the user. For example, if the user selects the category "flight services" with properties `destination=Helsinki, time=2002-10-18`, then the yellow page system could automatically show links to accommodation services in Helsinki and events taking place there at that time.

4 A prototype of semantic yellow pages

As a first step towards semantic web based yellow page services, we have implemented a prototype called "Keltsi". The goals in the work were 1) to get hands-on experience on the practical issues related to implementing a semantic web based yellow page service, 2) to test semantic web technologies in the yellow page context, and 3) to get experience on how to model the service ontologies and to annotate advertisements.

The Finnish Yellow Pages Ltd provided us with 150 MB of authentical advertisement data, 26452 advertisements in 2574 categories. The database contains advertisements with properties such as the category name ("Restaurants", "Camera shops" etc.), the advertisement text, contact information, and location information (including the municipality and coordinates).

4.1 The knowledge base

The semantic web is based on knowledge presented 1) in the form of *ontologies* and 2) instance level *metadata* [1] conforming to it. An ontology is typically defined as a set of classes interlinked with each other by the subclass relation. For example, the class "HouseCats" could be a subclass of the class "Pets". Classes have *properties* that are called also slots or attributes. Properties may have *constraints* on their usage and other properties that are often called *facets*

[15]. For example, the class "Pets" may have the property "name" whose value must be a literal. Instances are related with corresponding ontological classes with an "instance-of" relation and inherit class properties via the class hierarchy. For example, any particular instance of "HouseCats" would inherit the property "name" from the class "Pets". The combination of ontologies and instances is called *knowledge base*.

To create a knowledge base corresponding to the Yellow Pages database, following major components are needed. 1) A service ontology covering the services provided by advertisements. 2) Ontologies for representing the service properties. 3) Annotation metadata describing the contents of the advertisements.

Given the large size of the database, it was desired to create ontologies and annotate advertisements for only a small fraction of the test data first. In our case study, a service, product, and location ontology was created. The service and product ontologies where based on the Universal Standard Products and Services Code[11] (UNSPSC) classification. The service ontology described the service classes that an advertisement could be connected with. For example, a restaurant advertisement would be connected with the class "Restaurants". The product ontology contained information about the products that an service provides. For example, a camera shop retails products that are of the types "Cameras" and "Photography films". The location ontology was based on the official list of Finnish municipalities[12].

RDF Schema [7] was chosen as the language for representing the ontologies and RDF [13] for representing instance data. This choice was motivated by the possibility to use widely available semantic web tools, such as the HP Lab's Jena toolkit[13] [14] for querying RDF(S) data, and Stanford University's Protégé 2000[14] [15] for creating the ontologies. The number of product and service classes in the ontologies was 432, and 34 different class properties were used. The location ontology contained 461 classes. For test purposes, we annotated 37 advertisements (service instances) in terms of the service, product, and location ontologies in the RDF database. Figure 1 illlustrates the connections between two instance advertisements and ontologies.

4.2 Search functionalities

The Keltsi system provides the end-user with two main functionalities: 1) ontology-based searching and 2) semantic instance browsing.

Ontology-based search The information search process of the user can be divided into the following steps:

1. Formulate the intention (e.g., "My camera is broken. How can I get it fixed?").

[11] http://www.un-spsc.net/, http://www.eccma.org/unspsc/
[12] http://www.intermin.fi/suom/laanit/kunnat.html
[13] http://www.hpl.hp.com/semweb/
[14] http://protege.stanford.edu/

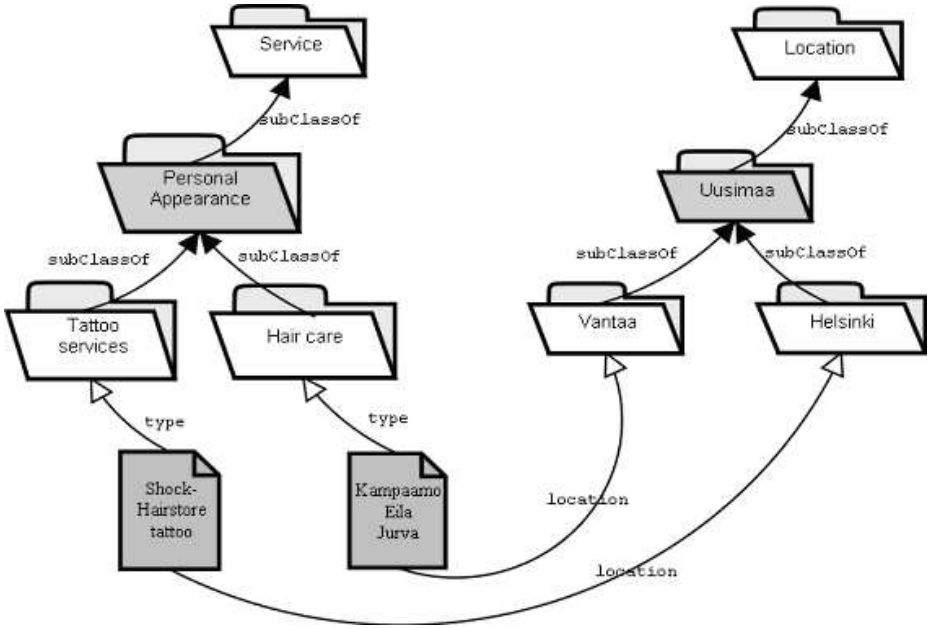

Fig. 1. Two advertisement instances connected to the service and location ontologies.

2. Choose a strategy for solving the problem (e.g., "Look for a solution in the Keltsi system.").
3. Present the intention using the models of the search interface (e.g., by using the user interface, create the following query: `product=Cameras`, `service=Repair Shops`, `location=Helsinki`)
4. Browse the results of the query (the advertisements), until the solution is found. If needed, the user can return to some of the previous steps. For example, one can reformulate the intention ("The repairing is too expensive compared to buying a new camera. Lets buy a new one."), change the strategy ("No solution in Keltsi, let's use Google."), or redefine the query.

Ontology-based searching means that the user can query the advertisement collection by using the classes and properties of the the underlying ontologies.[15]. For example, when looking for a Chinese restaurant located in Helsinki, one could select the "Restaurants" class and its properties `type-of-restaurant=Chinese` and `location=Helsinki`. The properties usually have a class value constraint, which means that the value must be selected from the corresponding ontology subtree, such as `location` and `type-of-restaurant`.

An interface supporting such functionality is shown in figure 2. The query is created by traveling the ontology three from general to more specific service

[15] This approach has been used, e.g., in [18].

Fig. 2. The user interface is divided into two parts: on the left hand side is the ontology-based searching tool and on the right is the semantic instance browser.

classes until the service class that best describes user's need is found. The class can be a very specific class, such as "Restaurants", or a more general class, such as "Eating and drinking establishments", which contains also bars, fast food services, etc. in addition to restaurants. The ontology class can alternatively be selected by writing the name of the class in the shortcut field, which makes the tree open up from the matching part. On the right of every ontology class, the number of corresponding service instances (advertisements) is shown. For example, the class "Services" has 37 instances connected to it and the class "Restaurants" has 10 instances.

The user can set property constraints that the searched service instances must fulfill. The properties and their possible values depend on the selected class. Therefore, if the selected class is "Restaurants", only properties related to "Restaurants" and its superclasses are shown. The selection of the value for a property is either based on traversing through an ontology hierarchy, such as the "location" of the service, by selecting the correct value from a list, or by writing the value in a text field.

The result of a search is a list of matching services where the type of the service and the name of the service is shown as a link. By clicking on the link, the information about the service and the advertisement is shown.

The semantic instance browsing The semantic instance browser shows the following information about the currently selected advertisement instance to the

user: the literal properties of the service (such as the topic and contact information), the graphical advertisement (if available), the relations to the ontologies (such as to the type of service), and the recommended links.

The idea of the recommended links is to provide links to related services that are supposed to be of interest to the user. For example, if when looking at an advertisement of a hotel service, one could also be interested in activities near the hotel, such as recommended restaurants and entertainment services.

In this initial version of the Keltsi system the recommended links are provided explicitly by the advertisers. In future versions, the creation of recommendations will be automatic, as well, and based on the ontologies and the instance data. An application of this idea is presented in [12].

The technology The Keltsi system was build as a three tier web application. The top layer provides the user interface with the help of Java Server Page technologies [11], including JSP Tag libraries [19], and the Apache Tomcat servlet engine[16]. The middle layer handles the business logic of the service. The undermost layer provides the functionality for querying the RDF database. In the implementation HP Lab's Jena toolkit [14] was used. Keltsi is used by an ordinary web browser.

5 Discussion

This paper proposed that service ontologies provide a promising new way to organizing and annotating advertisements in yellow page directories on the web. The underlying ontology can be used as a navigational device by which the end-user can map her needs onto relevant offerings (advertisements) in the catalog. After finding a suitable service and provider, the ontological structure relating services and their instances, the advertisements, provide a means by which recommendations to related advertisements of interest can be generated. By annotating offerings semantically according to the ontologies, more accurate semantic-based search facilities can be provided to the end-user.

The work presented is related research and standardization work on service description and matching on the web, such as [16] where DAML-S [3] is proposed as an ontology for annotating service capabilities. However, the main focus there is on modeling service capabilities in terms of the input and output of a web service [8] for a software agent. In our work, the problem of matching human user needs with advertisements is considered.

A demonstrational web-based system based on RDF Schema ontologies and RDF instance advertisements was described. This system is, however, only a first small step towards the vision of semantic yellow pages presented in this paper. Several advances are needed before an actual semantic yellow page service could be deployed. Firstly, full blown service and other ontologies need to be designed and tested on real customer data. Secondly, the advertisements should

[16] http://www.apache.org/

be annotated accordingly. This is more demanding and time consuming than assigning advertisements to isolated categories or in relational tables. Annotation work is the price to be paid for better accuracy in information retrieval and for better servicing the end-user. We feel, that the advertising companies and the yellow page operators should be interested in providing such annotations since the matchmaking facility is the key point of the whole directory service business in the first place.

An important organizational question here is whether the task of annotating the services should be done by the yellow pages operator or by the advertiser. In our view, this could be done in a distributed fashion by the advertisers active on the web. The advertisements could be published on the customer's public web pages in RDF, and collected by a web crawled for a global repository. The advertisers should be entitled, on desire, to do the job because they have the best knowledge of their services. The advertiser also has the interest in maintaining the advertisements up-to-date, and possibly add new temporal advertisements, such as offers of the week, as time goes by. For this task, instance metadata editor is needed to guarantee that the metadata given really conforms to the ontologies in use.

Distributing metadata and annotation work in this way over the web changes not only the way in which yellow pages are maintained but has other consequences, too. For example, how to prevent the advertisers from corrupting the yellow pages service by generic service descriptions matching practically with any need? This problem was encountered, e.g., on the WWW when providing HTML with the META-tag facility for describing web page metadata. The information provided by the page publishers turned out to be quite unreliable and current search engines nowadays tend to neglect META-descriptions. Providing service and product metadata in this way also changes the ways in which advertisement space is sold to the companies. For example, how to set prices for different kind of service and product advertisements?

Acknowlegdements

Arno Aalto, Merja Jalava, Tuomo Kajava, and Petri Lindgren participated in the team that implemented the Keltsi system. Irmeli Rämö and Pekka Raitala of Finnish Yellow Paged Ltd provided us with the actual case database and helped our work with discussions concerning yellow pages services. Thanks to Vilho Raatikka for fruitful discussions.

References

1. Kal Ahmed, Danny Ayers, Mark Birbeck, Jay Cousins, David Dodds, Joshua Lubell, Miloslav Nic Daniel Rivers-Moore, Andrew Watt, Robert Worden, and Ann Wrightson. *Professional XML Meta Data*. Wrox Press Inc., 2001.
2. Blair Alexander and Karen V. Fernandez. Are publishers putting theory into practice? In *proceedings of Australian and New Zealand Marketing Academy Conference (ANZMAC)*, December 1-5 2001. ISBN 0-473-08206-3.

3. A. Ankolekar, M. Burstein, J. Hobbs, O. Lassila, D. Martin, D. McDermott, S. McIlraith, S. Narayanan, M. Paolucci, T. Payne, and K. Sycara. DAML-S: web service description for the semantic web. In I. Horrocks and J. Hendler, editors, *The Semantic Web – ISWC 2002. First international semantic web conference*, pages 348–363. Springer–Verlag, Berlin, 2002.
4. R. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison-Wesley, New York, 1999.
5. T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Scientific American*, 284(5):34–43, May 2001.
6. A. Bernstein and M. Klein. Towards high-precision service retrieval. In I. Horrocks and J. Hendler, editors, *The Semantic Web – ISWC 2002. First international semantic web conference*, pages 84–101. Springer–Verlag, Berlin, 2002.
7. D. Brickley and R. V. Guha. *Resource Description Framework (RDF) Schema Specification 1.0, W3C Candidate Recommendation 2000-03-27*, February 2000. http://www.w3.org/TR/2000/CR-rdf-schema-20000327/.
8. P. Cauldwell, R. Chawla, V. Chopra, G. Damschen, C. Dix, T. Hong, F. Norton, U.Ogbuji, G. Olander, A. Richman, K. Saunders, and Z. Zaev. *Professional XML web services*. WROX Press, Birminghan, UK, 2001.
9. D. Fensel (ed.). The semantic web and its languages. *IEEE Intelligence Systems*, Nov/Dec 2000.
10. D. Fensel. *Ontologies: Silver bullet for knowledge management and electronic commerce*. Springer-Verlag, 2001.
11. D. K. Fields, M. A. Kolb, and S. Bayern. *Java Server Pages*. Manning Publications Co., 2002.
12. E. Hyvönen, A. Styrman, and S. Saarela. Ontology-based image retrieval. Number 2002-03 in HIIT Publications, pages 43–45. Helsinki Institute for Information Technology (HIIT), Helsinki, Finland, 2002. http://www.hiit.fi.
13. O. Lassila and R. R. Swick (editors). Resource description framework (RDF): Model and syntax specification. Technical report, W3C, February 1999. W3C Recommendation 1999-02-22, http://www.w3.org/TR/REC-rdf-syntax/.
14. B. McBride, A. Seaborne, and J. Carroll. Jena tutorial for release 1.4.0. Technical report, Hewlett-Packard Laboratories, Bristol, UK, April 2002. http://www.hpl.hp.com/semweb/doc/tutorial/index.html.
15. N. F. Noy, M. Sintek, S. Decker, M. Crubezy, R. W. Fergerson, and M. A. Musen. Creating semantic web contents with Protege-2000. *IEEE Intelligent Systems*, 16(2):60–71, 2001.
16. M. Paolucci, T. Kawamura, T. Payne, and K. Sycara. Semantic matching of web service capabilities. In I. Horrocks and J. Hendler, editors, *The Semantic Web – ISWC 2002. First international semantic web conference*, pages 333–347. Springer–Verlag, Berlin, 2002.
17. Steve Pepper. The TAO of Topic Maps. In *Proceedings of XML Europe 2000, Paris, France*, 2000. http://www.ontopia.net/topicmaps/materials/rdf.html.
18. A. T. Schreiber, B. Dubbeldam, J. Wielemaker, and B. J. Wielinga. Ontology-based photo annotation. *IEEE Intelligent Systems*, 16:66–74, May/June 2001.
19. G. Shachor, A. Chase, and Magnus Rydin. *JSP Tag Libraries*. Manning Publications Co., 2001.