Ontologies for Information Fusion

Deborah L. McGuinness
Associate Director and Senior Research Scientist
Knowledge Systems Laboratory
Stanford University
Stanford, CA 94305
dlm@ksl.stanford.edu

Abstract - *Ontologies have moved beyond the domains of library science, philosophy, and knowledge representation and are now being used to facilitate mainstream applications in all areas concerning information. Information fusion applications can benefit from access to information contained in ontologies such as controlled vocabularies, term meanings, term relationships, and logical implications of term usage. We will define ontologies by introducing the notion of an ontology spectrum ranging from simple vocabulary information to more complex knowledge structures. We will point to examples of ontologies available at different points along the spectrum. We identify how ontologies at differing points along the spectrum may be used to support improvements in information management solutions ranging from better search to precise semantic matching, noting that all of the benefits can be applied to information fusion. We will discuss the potential impact, benefits, and costs of using ontologies in information fusion applications.*

Keywords: *Ontologies, Semantic Web, Information Merging, Information Mapping, Information Integration, Knowledge Representation, DAML+OIL, OWL.*

1 Introduction

If information fusion is to achieve its goal of using information from multiple sources, it must address the problem of information integration. In this paper, we use the description of information fusion provided on the international society of information fusion page by Dr. Dasarathy[1]: "Information Fusion, in the context of its use by the Society, encompasses the theory, techniques and tools conceived and employed for exploiting the synergy in the information acquired from multiple sources (sensor, databases, information gathered by human, etc.) such that the resulting decision or action is in some sense better than (qualitatively or quantitatively, in terms of accuracy, robustness and etc.) than would be possible if any of these sources were used individually without such synergy exploitation."

If applications expect to exploit synergies amongst different information sources, they will need to have some notion of the meaning of information so that they can determine how information can be combined to enable decisions or actions that are better than those that could have been produced from a single source.

There are many ways that information combination can enable better decisions. If information in a second (authoritative) source confirms information from the initial (authoritative) source, then a decision based on this information may be viewed more credibly. This could be viewed as increasing the confidence of the system answers due to *recognizing redundancy*. In order to present this confirming evidence to an end user, a system must be able to recognize that two or more sources contain the same information. There are a range of techiques that could be used to determine that the information is the same. In this paper, we will focus on declarative techniques that range from matching text to matching an internal representation of the meaning of the text. We will show that ontologies may be used all along that spectrum.

Another way that an additional source of information may add value to the original information source is if a user (human or agent) is able to combine two pieces of information together and deduce additional information. For example, if one information source contained "A" and another information source contained "A => B", then the combination of both sources, could determine that "B" was valid. In order for this deduction to occur however, the combined application would need to recognize that "A" is the same in both information sources (and also if both information sources contain B, that it is the same B). This kind of improvement to the information could be viewed as a refinement of an assessment since it is adding information (in a consistent manner) by *finding information that is implicit and making it explicit*.

Another source of value from information fusion is the *identification of conflict*. For example, one source of information may contain an explicit statement that "A" is true. Another source may contain an explicit statement that "A" is false. Identifying the conflict means that the

[1] http://www.inforfusion.org/mission.htm

application needs to determine that the "A"s are the same in both sources. Additionally, one of the sources could contain information that *implicitly* states that "A" is false (such as a statement that either "B" or "~A" is true and "~B" is true, thus, "~A" must be true). This kind of improvement in information could be viewed as conflict detection and requires both identification of like terms as well as an inference engine capable of identifying information that is implicit in the information stream and capable of identifying contradictions.

Another source of information enhancement from multiple sources involves generalization and aggregation. If one source knows "A" is a subclass of "B" and also knows an instance "A1" of "A", and another source knows "C" is a subclass of "B" and an instance "C1" of "C", then an application that recognizes generalization (as well as recognizing that the class "B" is the same in each source and also that "C1" is distinct from "A1", can know that there are two instances of "B". Using generalization hierarchies to count occurrences can be useful in a number of ways. This can be used to get a sense of how often classes of events happen. It also can be used to abstract information away so that sensitive details are obscured while providing a more general description of the class of events.

The methods of exploiting synergy for information fusion have some properties in common. They rely on some understanding of terms so that they can recognize when terms are identical, distinct, or disjoint. Some of them also rely on the ability to process statements, making implicit information explicit and providing conflict detection (and possibly conflict resolution). This paper explores ways of using ontologies to support varying degrees of term understanding for use in information fusion. We will explore term meaning as captured in ontologies.

2 Ontologies

The notion of an ontology is not new; Merriam-Webster[2] dates it to 1721 and ties it to the branch of metaphysics concerned with the nature of relations of being. For the purposes of this paper, (and the purposes of using ontologies for information fusion and knowledge sharing), we prefer use a less philosophical definition and instead use a more operational definition. We refer to the widely cited definition proposed by Tom Gruber: An ontology is a specification of a conceptualization [1]. Intuitively, users and agents have some notion of the meaning of a term – a conceptualization – and an ontology specifies that concept. Guarino [2], provides an extended discussion of this notion, grounding what a conceptualization is more formally. He also provides a nice collection of fields that have used ontologies providing descriptive examples, including knowledge engineering, knowledge representation, qualitative modeling, database design, information retrieval and extraction, and knowledge management. McGuinness [3] expands on that list and provides more examples including library science, ontology-enhanced search, e-commerce, and configuration. The same paper includes a spectrum that we will expand on here that can be used to view ontologies. Figure 1 is based on the Ontology Spectrum presented in [3][3]. It has been refined as a result of additional work viewing ontologies as spectrums in [4] (which was an extension of [3]), a discussion and evolution of the spectrum in [3] with Joe Rockmore, multiple presentations of the spectrum in ontology talks (e.g., The Semantic Web for Military Users Meeting Series[4]), and [5] (which was done independently but with a similar result).

The spectrum moves from weaker and fewer relationships between terms, (and typically less precise semantics) to stronger and more relationships between terms (and typically more precise semantics). Also, as one moves to the right in the spectrum, there is more possibility for deeper automated reasoning support. The spectrum spans the space that most people consider when they discuss ontologies. In fact, formalists would typically not consider points to the left side of the line to be ontologies however in broad usage, all points on this spectrum have been referred to as ontologies.

The simplest notion displayed on the spectrum is a catalog – which is simply a controlled vocabulary of words or phrases (possibly including an index number along with the phrase name). This includes simple controlled vocabularies and indexing schemes. One useful and long lasting catalog is the Dewey Decimal Classification System[5]. Without the semantics associated with the numbers in the system and the glossary, it can be viewed as a catalog. With the semantics attached to the numbers, it can be viewed as a hierarchy[6]. Adding a little more content, we move into a glossary that captures some notion of meaning, although typically in natural language. The natural language meanings are rarely operational to machines but can be quite useful to humans. In this space we include standard glossaries and data dictionaries. This would also include efforts such as EDI – Electronic Data Interchange [7].

[2] http://www.m-w.com/cgi-bin/dictionary

[3] McGuinness generated the original spectrum from a discussion prior to an Ontologies panel at AAAI 1999 between Gruninger, Lehmann, McGuinness, Ushcold, and Welty.

[4] http://www.daml.org/meetings/

[5] http://www.oclc.org/dewey/about/index.htm

[6] Some may argue that this is a formal taxonomy since the documentation claims that all subtopics are part of all of the broader topics above them , however, we will distinguish instance and subclass relationships from part-of relationships in our forthcoming description.

[7] http://www.itl.nist.gov/fipspubs/fip161-2.htm

Moving a bit more to the right on the spectrum, we include some notion of structure to include thesauri. A thesaurus includes the notion of synonyms (and cross references containing some additional structure between terms). While more structured than a glossary, the precise semantics of synonym is rarely stated to be something that reasoners would expect, i.e., terms that have the identical extension in a set theoretic sense. Synonyms are used to connect two terms that have similar or nearly the same meaning, thus while this information is useful, it is not precise enough to use for certain types of reasoning unless an unambiguous definition is provided for "similar" or "nearly the same". Also thesauri do not typically contain explicit hierarchical information, however when they include narrower and more general statements, simple hierarchies may be deduced.

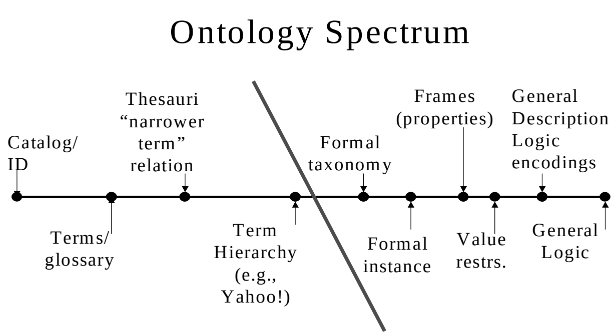

Figure 1: Ontology Spectrum

The next point on the spectrum includes a term hierarchy. This point includes descriptions that include explicit generalization and specialization information. In the previous spectrum, it was called "informal is-a" since the knowledge representation community has referred to the formal subclass relationship as an "is-a" relationship. If a formal subclass relationship holds between class A and class B, then an instance of class B must always be an instance of class A. There are no exceptions to this statement. For example, in the Yahoo shopping hierarchy, it is the case that anything that is an instance of "Girl's Jeans" is also an instance of "Teens" (possibly more appropriately named "Teen's apparel") and that same item is also an instance of "Apparel". In this case, "Apparel" is the superclass, "Teens (Apparel)" is the subclass (and additionally "Girl's Jeans" is a subclass of "Teens"). In the same portion of the hierarchy however, there is a category called "Accessories" which is hierarchically more specific than "Apparel". If this hierarchy supported a formal subclass or is-a relationship, then all instances of "Accessories" would also be instances of "Apparel". Thus, the items under men's accessories including wallets, umbrellas, money clips, etc., would all be instances of Apparel. Yahoo! Shopping and many sites on the web that contain hierarchies contain many subclass relationships but do not strictly enforce subclass relationships[8]. In our spectrum, we refer to these as hierarchies. We also have them to the left of the line since many people want to depend on strict subclass relationship enforcement to be a precondition for ontologies.

The next point on our spectrum includes formal taxonomies with strict subclass enforcement. This point includes class descriptions only. We distinguish two points in the spectrum of formal taxonomies that include only class names and a point that also includes ground instance information as well. Many believe that ontologies only include class descriptions and exclude individual terms, although some believe that individuals may be included in the ontology. One intermediate is to include a small set of "special" individuals in the ontology that are needed for class descriptions (and property descriptions as one moves further to the right on the spectrum). For example, in a wine ontology (e.g., the one used in the OWL Guide[6] and KSL Wine Agent[9]), since one natural way of referring to wines is to distinguish them by color and varietal, those individuals are included in the ontology. Thus, we might include a class for RedWine, which would reference the individual color Red and we may also include varietals such as Chardonnay and Zinfandel. However, all of the individual instances of any particular wine, for example all of the chardonnays produced in California are not part of the ontology but instead are kept in a knowledge base containing instance information. Many languages may be used to encode the taxonomies and even the relatively less expressive markup language XML may be used to encode this point on the spectrum.

As we move further to the right, we include frames that include class descriptions along with properties that may be associated with classes. Properties (also sometimes called slots, attributes, or roles) are used to relate one term to another and are typically two place predicates. For example, the wines ontology may include the class winery and also the property hasLocation that may be used to relate a winery (as well as other things) to a location (such as a wine growing region or a precise address). Many systems may be used to encode this kind of simple frame information, with the open knowledge

[8] In the Yahoo ! Shopping hierarchy apparel when clicked on expands to Apparel, Accessories, and Shoes. This may be in an effort to allow subclass relationships to be supported. However, the travel category also contains as a subcategory accessories, which in turn contains things such as binoculars, pillows, etc. There is no such expansion of the travel category's name, thus if strict subclass relationships were supported, binoculars and pillows would also be instances of Travel, which of course they are not.

[9] www.ksl.stanford.edu/people/dlm/webont/wineAgent/

base connectivity[10] (OKBC) application programming interface being a way to access the information.

The next point on the spectrum includes a simple notion included in essentially all description logics – the ability to put a restriction on a value on a per class basis. Most frame systems allow someone to place a range restriction on a property. For example, in the wines ontology, the hasLocation property would have a range restriction placed on it that anything that fills this property is an instance of a geographicLocation. This is considered a global restriction since any time the property is used, its value must be a geographicLocation. Many times people want to reuse the same property but add additional restrictions on a per class basis. Thus, if one were to define a NapaWinery, one might further restrict the values to be instances of NapaCounty. This point on the spectrum arises as ontologies start to move towards the next point –where we include the full range of standard description logic object (class, property, and instance) descriptions. This point was chosen to represent a movement towards more expressive description logics. OWL Lite [7], for example, contains more than this but does not contain the full expressive power of description logics.

The next point on the spectrum includes the full range of constructors in Description Logics [7]. Description Logics provide a carefully chosen set of expressive features for the representation language that are chosen so that efficient reasoning algorithms may be written to compute the implicit information. The most recent description logic that is seeing wide acceptance is OWL [9] – the ontology web language from the World Wide Web Consortium's Web Ontology Working Group[11] that was based on DAML+OIL [10] – which is the markup language that was the result of the DARPA Agent Markup Language Program[12] along with European Union Semantic Web collaboration. It includes for example the ability to state precise cardinality restrictions, inverse property relationships, disjoint classes, inverse functional properties, etc.

The final point on the spectrum includes general logical axioms. There are some things that can not be expressed in description logics (or any other intermediate point on the spectrum since language designers have limited their language for either understandability or computation properties. Some applications however need to be able to express full first order logic expressions. Many systems that need to do this have chosen KIF[13] – Knowledge Interchange Format – for encoding first order logic. For the purposes of this spectrum, we just stop at general logical statements and do not distinguish first order logic statements from higher order logics.

[10] http://www.ai.sri.com/~okbc/

[11] http://www.w3.org/2001/sw/WebOnt/

[12] http://www.daml.org

[13] http://logic.stanford.edu/kif/

3 Uses of Ontologies for Fusion

In this section we will discuss some ways that ontologies may be used for fusion activities. We begin with applications that can benefit from ontologies further to the left on the spectrum and then move to applications that benefit from ontologies further to the right.

3.1 Uses of Simple Ontologies

The points to the left of the ontology spectrum are easier to either build and maintain internally or to obtain from outside sources. There are many existing controlled vocabularies, glossaries, and hierarchies available for use and reuse. The author in a previous paper [3] identifies a number of sources and also identifies uses for simple ontologies as follows:

Controlled Vocabulary – If authors are encouraged to use terms from a controlled vocabulary, the presentation of material is uniform and more easily searched. For example, if analyst reports used a controlled vocabulary and if output from level 1 fusion efforts used the same vocabulary, consumers of those reports would have an easier time finding things and making sense of them.

Site organization and navigation support – Taxonomies or even hierarchies may be used to organize information in a repository. The organization taxonomy may also be used to set users expectations about what information may be found in a collection or web site. Thus, if a repository is created of reports and incoming data, consumers of the repository may navigate through the taxonomy checking to see the general categories covered by the repository.

Upper Level Organization or Umbrella – A broad taxonomy may be used to help tie different repositories and information together. One could use, for example, the Suggested Upper Merged Ontology (SUMO)[14] or the Upper Cyc Ontology[15] to integrate different sources of information together. Additionally, there are a number of domain specific ontologies covering, for example, weapons of mass destruction. One could reuse one of these ontologies as a general coordination scheme and then extend them in areas as necessary. Different information producing activities could be used to populate and extend portions of interest in these ontologies. The general domain specific ontologies may be used for example, to find lists of subclasses of categories of interest. Also when these mid-level domain ontologies contain property information, they may be used to identify the range of data that is typically collected for a particular class.

[14] http://ontology.teknowledge.com/

[15] http://www.cyc.com/cyc-2-1/cover.html

Browsing support and source of tags for markup – A taxonomy can be used as Yahoo! does to tag content. The tagging may be done by humans as done in the Yahoo! approach or it may be done using (semi-)automated approaches. Once documents are tagged with topic tags chosen from the taxonomy, they may be handled in a manner appropriate to the content.

Ontology-enhanced search – A taxonomy may be used to support query expansion and allow search users to obtain relevant documents that otherwise would have been missed. Query expansion is done by obtaining a search query, checking to see if any of the phrases in the query are in the taxonomy, and if so, adding appropriate terms from the taxonomy to expand the query. For example, if a taxonomy of Weapons of Mass Destruction included a subclass of biological weapons and further included information about specific biological weapons including smallpox, anthrax, and Bacillus anthracis spores, then a query for weapons of mass destruction could be expanded using the taxonomy to also look for biological weapons, smallpox, anthrax, and bacillus anthracis spores.

Enhanced search capabilities is one capability that may be of particular use in fusion. Standard information retrieval techniques, while extremely useful and effective, may not be enough in the face of certain common conditions. For example, if the documents that are to be searched are not long and thus do not contain many phrases that are discussing the same thing, then the chances that the search query matches the exact phrasing of that topic in document is decreased. We include in Figure 2 the architecture we used in FindUR[3] to provide an ontology-enhanced search capability leveraging a standard commercial search engine (Verity) along with background taxonomies built and maintained in a description logic environment. We built a distributed web tool that allowed a broad range of people to update the background taxonomies without knowing anything about the underlying description logic system. Since it was a web deployment maintained in a distributed manner, the front end (called the collaborative topic set tool) included an interface for working in a collaborative manner. While the taxonomies were built and checked by the description logic system, the deployment was done using Verity's topic sets so that it was compatible with the widely accepted commercial grade IT platform at the large company where the work was done. The user interface for search results was also developed to provide a GUI that could be customized according the the context, domain, and user.

This same basic architecture could be used with any application where content (web pages, databases, or anything else that may be searched) contains information that should be retrieved and standard statistical retrieval techniques are not quite enough. It may be used to use customized vocabularies to help refine search queries and results. It also does not need to be maintained by ontology experts or search experts. Once the basic infrastructure is put in place, any domain literate person may be used to do the maintenance. Also, while it utilizes tools such as a description logic system to do some checking, that system is not required for the deployed environment (which was useful in its deployment setting so that IT departments did not need to expand their product array or staff literacy in order to deploy and maintain it.

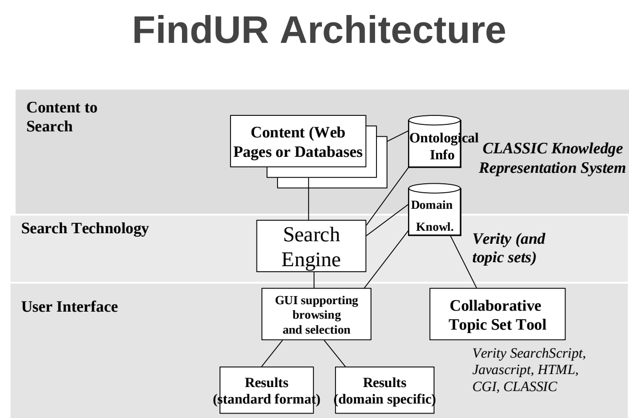

Figure 2: A Knowledge-Enhanced Search Structure

In the simple description above, we only mentioned using taxonomic information and that is largely what FindUR used in most of its deployments. However, some applications such as a medical information retrieval application [11] used more content from the ontologies, typically exploiting property information and value restrictions. While not using the same architecture, many other applications have used similar approaches. For example, TAP's activity-based search[16] and Cycorp's e-Cyc[17] used similar styles and also used more ontological content in their enhancement of search facilities. Additionally, once more structure is used in the search, that same structure can be used for presentation and comparison. We will explore use of more detailed ontologies in the next section.

3.2 Uses of Structured Ontologies

When ontologies contain structured information such as range restrictions (limitations on the types of values that may fill a property) or class disjointness (no object may be an instance of two disjoint classes, e.g., no object is simultaneously a country and a person), then additional processing may be done and this processing may be exploited in applications.

Restrictions captured in the ontology may be used in a number of ways that support fusion. One of the simplest is *consistency checking*. If a range of values is provided

[16] http://tap.stanford.edu/tap/ss.html

[17] http://www.cyc.com/auschron2.html

for a property, e.g., a maximum speed is between 50 and 150 miles per hour, then if a sensor is reporting that the maximum speed is 1500 miles per hour, the application can flag the discrepancy. Consistency checking before information fusion may be useful to help eliminate propagation of misinformation. Consistency checking can be done using all types of information captured in the ontology, thus including more than just value ranges but cardinality restrictions, disjointness information, enumerated sets, etc. Restrictions may also be used when filtering search results. If a search is done over semi-structured or structured information and the enhanced search application knows the range that it should be finding when for example, it is looking for maximum speeds reported for certain classes of objects, then it can prune (or reorder) results according to the restriction.

Ontological information may also be used as a source of *completion*. If descriptions are provided of terms such as "high-performance-x" and that expands to a particular set of characteristics, then whenever an agent or human sees the term, the characteristic set may be expanded and used for processing including consistency checking.

Possibly most importantly ontologies can be used for *interoperability* thus allowing applications to access and integrate multiple information sources. The key to accessing and integrating sources is to understand what is covered in the sources and how the terms are used. Approaches include using ontologies to help applications such as [12] to access and integrate multiple databases and knowledge bases using domain and information source models stored in an ontology. Some view the task as being mediators as in [13] that specify information about accessing information from sources and provide abstraction and transformation information between the sources. Sometimes this information is also viewed as "articulation axioms" that specify how one term in one source is related to a term in another source. We will describe an application architecture below that incorporates these strategies. Also, along the lines of integration and merging, environmental research tools such as Chimaera [14] and Prompt [15] or commercial tools such as the Medius[18] toolkit exist to help users with the phrase merging process.

True support for interoperability includes a number of aspects. Beyond determining how to access data, when terms are the same, and how to encode and maintain controlled ontologies, we may want support for diagnosing ontological information. Ontologies may be used to encode a number of constraints that may be used for integrity checking and for theorem proving. Simple diagnostic environments such as Chimaera provide basic checking capabilities for example to check for value restriction violations, cyclic definitions, etc. More sophisticated domain specific applications have been done called configurators [16] that will take a set of specifications and generate a parts list matching the specifications. The tasks that fall out of a configurator simply are (i) checking incoming constraints to determine if they are incompatible and thus if the problem is over-constrained, (ii) providing completion to the constraints taking any implicit information and making it explicit, thereby providing a more detailed specification of the constraints along with information about how those additions were generated, (iii) generating a solution that meets the constraints, which for a configurator is a parts list but for a general constraint problem could be a plan, (iv) providing an explanation of why any portion of the solution was arrived at (and potentially why it was chosen over other solutions), (v) sometimes providing constraint relaxation suggestions to take an over-constrained problem and make it into a problem that has solutions.

3.3 An integration architecture

Figure 3 provides an architecture[19] that has been used in an application of semantic web technologies, heavily relying on ontologies for providing integrated access to diverse information sources. It is appropriate for use in information fusion activities including multiple information sources with situation and threat assessment needs.

Beginning in the lower left of the figure, there are a number of information sources, including both structured sources such as databases, and documents that may be unstructured. There are also repositories such as keyword indices and category indices. There is an ontology server that contains domain information (in OWL). The query builder tool uses the ontological information to help the consumer via the web portal. Through the portal the user obtains access to the multiple information sources without needing to specify which ones to use and also without needing to know the query languages of the multiple access programs to the information sources. There are mediators that are used by the query processor broker to use the appropriate information source, with the appropriate query string, with the appropriate vocabulary. A query results processor also utilizes the ontology to help with the integration task of presenting results in an integrated format, possibly with language modifications or abstractions based on the ontology. The results may be used to populate a knowledge objects repository that contains results from queries and deductive processing. This knowledge objects repository may be used as a cache for integrated answers, a temporary work space if information is not automatically committed to memory, a store used for viewing information that may be of a sensitive nature and thus need additional filtering based on clearance level or need to know, etc. This knowledge objects repository is part of the knowledge base for the

[18] http://www.sandsoft.com/products.html

[19] Original architecture viewgraph provided by Joe Rockmore. Abstraction modifications done by author.

intelligent application. Also stored in the knowledge base are rules for marking up documents. Some documents come in with limited markup and would be more useful with additional markup and other documents come in with no markup. Stored in the knowledge base, are rules about how markup may be provided for the documents. In a setting where documents come in and markup is to be added, there are semantic markup tool kits that a producer of marked up information may use along with the knowledge base including the ontological information, the markup rules, and the knowledge objects repository in order to generate marked up web documents. These marked up documents may then be used to provide more searchable content as well as content that is more easily integrated and viewed in multiple contexts.

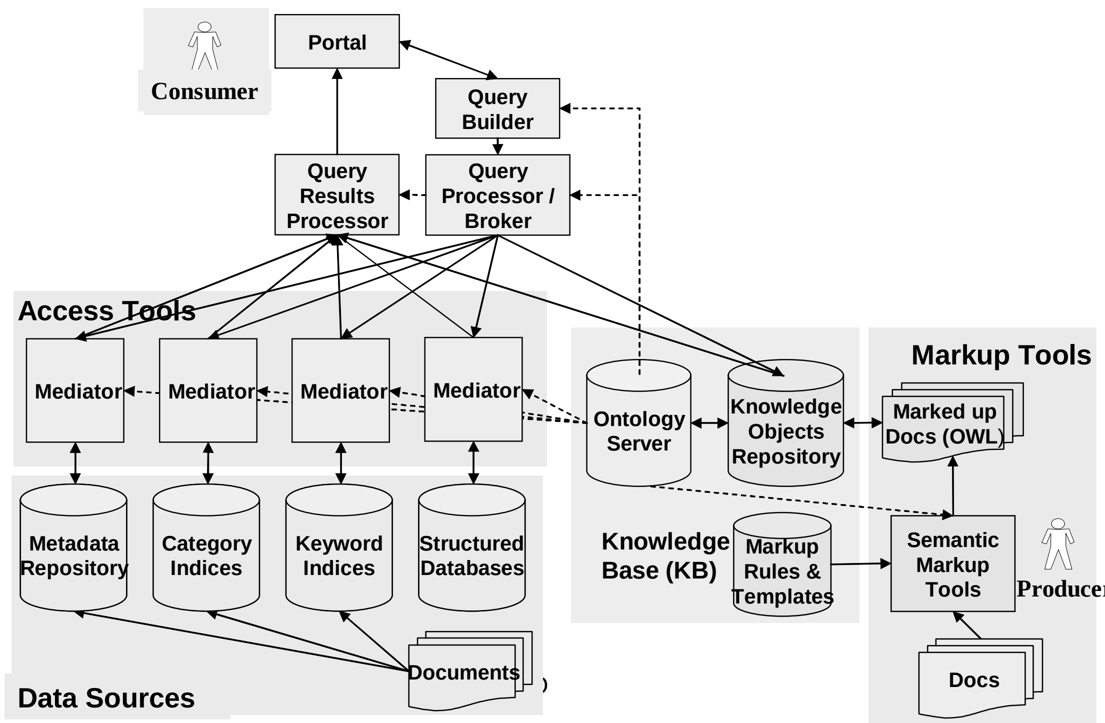

Figure 3: Ontology-enhanced application architecture

4 Conclusions

In this paper, we have described the notion of an ontology spectrum and have discussed points on the spectrum ranging from simple, controlled vocabularies to sets of logical axioms. We have provided descriptions of points ranging along the spectrum and have discussed how they can be used for information fusion problems. We have provided more detail on two kinds of ontologically-enhanced applications: one that uses simple hierarchies in a query-expansion approach to ontology-enhanced search. This may be valuable since the hierarchies the application uses are relatively simple to obtain and maintain and can be integrated into commercial off the shelf readily available products. We also provided more detail on an architecture that can be used with both more detailed ontologies (and simple ones as well if that is all that is available) in order to provide more efficient information fusion applications.

References

[1] T. R. Gruber. A translation approach to portable ontologies. *Knowledge Acquisition*, 5(2):199-220, 1993.

[2] Nicola Guarino. *Formal Ontology and Information Systems*. In Proceedings of FOIS’98, Trento, Italy, June 6-8, 1999. Amsterdam, IOS Press, pp. 3-15.

[3] Deborah L. McGuinness. *Ontologies Come of Age*. In Dieter Fensel, Jim Hendler, Henry Lieberman, and Wolfgang Wahlster, editors. Spinning the Semantic Web: Bringing the World Wide Web to Its Full Potential. MIT Press, 2002. http://www.ksl.stanford.edu/people/dlm/papers/ontologies-come-of-age-abstract.html

[4] Semantic Integration. In Bradshaw, J. M., Boy, G., Durfee, E., Gruninger, M., Hexmoor, H., Suri, N., Tambe, M., Uschold, M., & Vitek, J. (Ed.). (2003). *Software Agents for the Warfighter. ITAC Consortium Report.* Cambridge, MA: AAAI Press/The MIT Press, in preparation.

[5] Michael Daconta, Leo Obrst, and Kevin Smith. The Semantic Web: A Guide to the Future of XML, Web Services, and Knowledge Management. Wiley, 2003.

[6] Michael Smith, Chris Welty, and Deborah McGuinness, editors. OWL Web Ontology Language (OWL) Guide Version 1.0. W3C Last Call Working Draft 31 March 2003. http://www.w3.org/TR/owl-guide/ .

[7] Deborah McGuinness and Frank van Harmelen, editors. Web Ontology Language (OWL): Overview. W3C Last Call Working Draft 31 March 2003. http://www.w3.org/TR/owl-features/ .

[8] Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors, The Description Logic Handbook: Theory, Implementation and Applications, Cambridge University Press, 2003.

[9] Mike Dean and Guus Schreiber, editors. OWL Web Ontology Language Reference. W3C Working Draft 31 March 2003. http://www.w3.org/TR/owl-ref/

[10] Dan Connolly, Frank van Harmelen, Ian Horrocks, Deborah L. McGuinness, Peter F. Patel-Schneider, and Lynn Andrea Stein. DAML+OIL (March 2001) Reference Description. World Wide Web Committee (W3C) Note 18 December 2001. http://www.w3.org/TR/daml+oil-reference.

[11] Deborah L. McGuinness. *Ontology-enhanced Search for Primary Care Medical Literature*. In the Proceedings of the International Medical Informatics Association Working Group 6 – Medical Concept Representation and Natural Language Processing Conference. Phoenix, Arizona, December 16--19, 1999.

[12] Y. Arens, C. Chee, C. Hsu and C. Knoblock. Retrieving and Integrating Data from Multiple Information Sources. In Journal of Intelligent and Cooperative Information Systems, Vol. 2, June 1993.

[13] Gio Wiederhold and Michael Genesereth. *The Conceptual Basis for Mediation Services*. *IEEE Expert, Intelligent Systems and their Applications*, Vol.12 No.5, Sep-Oct 1997.

[14] Deborah L. McGuinness, Richard Fikes, James Rice, and Steve Wilder. *An Environment for Merging and Testing Large Ontologies*. In the Proceedings of the Seventh International Conference on Principles of Knowledge Representation and Reasoning (KR2000), Breckenridge, Colorado, USA. April 12-15, 2000.

[15] Natasha Noy and Mark Musen. *PROMPT: Algorithm and Tool for Automated Ontology Merging and Alignment*. in Seventeenth National Conference on Artificial Intelligence (AAAI 2000).

[16] Deborah L. McGuinness and Jon Wright. ``An Industrial Strength Description Logic-based Configurator Platform''. *IEEE Intelligent Systems*, Vol. 13, No. 4, July/August 1998, pp. 69-77.