

***Towards an infrastructure for knowledge
management within European Ethics:
Methodologies for semantic integration of
heterogeneous resources***

Report for the Feasibility Study

**for an advanced systematic documentation, information
and communication tool in the field of ethical issues in
science, research and technology
(FASTER)**

Work Package 5B: Communication Links



Food and Agricultural Organization
Library and Documentation Systems Division
WAICENT/FAOINFO Dissemination Management Branch

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

Participants Work Package 5 B - Contact Information

Food and Agricultural Organization

Library and Documentation Systems Division

WAICENT/FAOINFO Dissemination Management Branch

viale delle Terme di Caracalla

Rome 00100

ITALY

Johannes Keizer

johannes.keizer@fao.org

Anita Liang

anita.liang@fao.org

Contributors

Gauri Salokhe

Margherita Sini

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

INTRODUCTION	4
SEMANTIC WEB	5
The XML layer: Content of a resource	6
Granularity	6
The RDF layer: Metadata of a resource	7
The Ontology layer: Modelling the domain	8
METHODOLOGIES FOR SEMANTIC INTEGRATION	9
Ethics Application Profile	9
Definition and rationale	9
Methodology for the development of a European Ethics application profile (EEAP)	11
Ethics Application Ontology	15
Definition and rationale	15
Mapping the application profile to the application ontology	16
The ongoing development of the EEAO	17
Methodology for ontology-building	20
SCENARIOS: RELATING ONTOLOGIES AND DATA	25
Mediated (centralized) v. federated multi-host (distributed) architecture	26
Three scenarios	27
The AO and unindexed structured data	27
The SO and unstructured data	28
The AO + one or more thesauri	29
The AO + sub-ontology	30
CONCLUSION	30
REFERENCES	31

Abstract

As with many domains, information retrieval and knowledge management (IR/KM) in ethics suffers from the problems of semantic heterogeneity, making it difficult for providers to disseminate their services effectively and for users to retrieve the information they need. An analysis was conducted of resources provided by two hundred or so information services within the domain of ethics as it applies to issues in science, technology, and research in a European context. Based on this analysis, this paper proposes a) an application profile for dealing with the problem of heterogeneity originating from differences in terminologies, domain coverage, and domain modelling, and b) a root ontology based on the application profile which can serve as a basis for extending knowledge of the domain.

Introduction

Recent surveys (Hopt et al. 2004; Denger 2004a, Denger 2004b) of online and offline information services in the domain of European ethics have identified over 200 information systems and services providing access to a wide range of information resources. Another survey of user behaviours and preferences (Burgess *et al.* 2004) identified the number of frequent and occasional users of fourteen sites, while another 235 sites were identified by respondents themselves as sites that they consulted.

The information resources provided by these services differ in degree of coverage (e.g., some resources, such as AGROVOC, principally concern other domains but contain information related to ethics), sub-discipline (e.g., medical ethics, legal ethics), type (e.g., databases, link collections, images), and content (e.g., journals, institutions, expert information, project descriptions, thesauri). What is immediately evident from these surveys is the distributed, heterogeneous nature of the resources comprising this domain, and that no single search engine could retrieve a comprehensive set of the resources relevant to a user's needs. Users who are looking for information on, say, French researchers working on stem cells, have at their disposal search engines that can go through these files (and indeed, the millions of files on the WWW) at the blink of an eye. Yet, there is little guarantee that what the search engines will find and display will correspond in meaningful ways to the user's query: She may enter the query *stem cell researchers in France* but the resources that might constitute good responses to her query refer to *French researchers* and employ the plural form *stem cells*¹; or, alternatively, it may be the case that relevant information is included within the results displayed, but that she has to manually sift through a dozen pages of irrelevant results (e.g., pages containing her query in the form of a bag of words, i.e., *stem, cell, researchers, and French*, appearing as separate terms) to get to them; or, the page might contain biased or inaccurate statements about the topic.

¹ For example, Google displays 2.49m results for *stem cells* and 3.24m results for *stem cell* with an overlap of 6 links on the first page of results for each query.

Towards an infrastructure for knowledge management within European Ethics

Up until now, the WWW has been a system for the global dissemination of electronic information, meant for human consumption. The ease of publishing content on the Web has led to an explosion in the quantity and heterogeneity of information published that is far beyond what humans are able and willing to process. To handle these huge quantities of information intelligently and efficiently, the WWW needs to transform itself into a system for disseminating knowledge that can be interpreted not only by humans but also and especially by computers. This implies an evolution to a web that is first and foremost meaning based rather than form based. Once the semantics of web content is exposed, the foundation for data interoperability, reuse, and categorization is laid.

An intelligent WWW, that is, one using semantic technologies, could then process the query *stem cell researchers in France*. Based on its “understanding” of the query (e.g., through a process of resolving the query terms into concepts and matching those concepts to an ontology over which reasoning can be performed), it could conduct not only a comprehensive search of all of the above mentioned 200-some services, but also retrieve/suggest related concepts and resources, irrespective of the actual terms and language of the query. In addition, the system could engage the user in a dialog to help her refine her query.

The aim of this paper is to specify the rationale and the methodologies for developing semantic standards in the domain of European ethics in the fields of science, research, and technology. In particular, we propose, on the one hand, a *European Ethics Application Profile* (EEAP) to address the problem of semantic heterogeneity. On the other hand, from the EEAP, we derive a corresponding ontology, the *European Ethics Application Ontology* (EEAO). The application profile upon which the application ontology is based makes explicit the semantics that already exists within the application profile, and may be further enriched with additional semantics through the introduction of schemes, thesauri, and other terminologies. Thus, the semantic richness of the application ontology varies according to the extent to which additional concepts and relations have been incorporated into the ontology. We will refer to this EEAO alternatively as a root ontology, since it serves as a starting point for further semantic extensions.

Developing and applying standards for resource description is a prerequisite for creating the infrastructure for a network of information services that can alleviate the semantic heterogeneity of the diverse and distributed services providing information resources in European ethics. Moreover, this emphasis on meaning over form allows for the development of "smart" applications for areas such as content management (e.g., automatic mark-up of documents), knowledge management (e.g., expert locators, concept-based search), and advisors/recommenders (e.g., mediators).

Semantic Web

In his vision of the Semantic Web, Tim Berners-Lee (Berners-Lee 2000) outlines an architecture for the Web that is multi-layered and machine processable, as depicted in the much-reproduced image in Figure 1. The layers with which we will principally be concerned are the resource description framework layer and the ontology layer. The XML layer will be touched upon insofar as it addresses the issue of content.

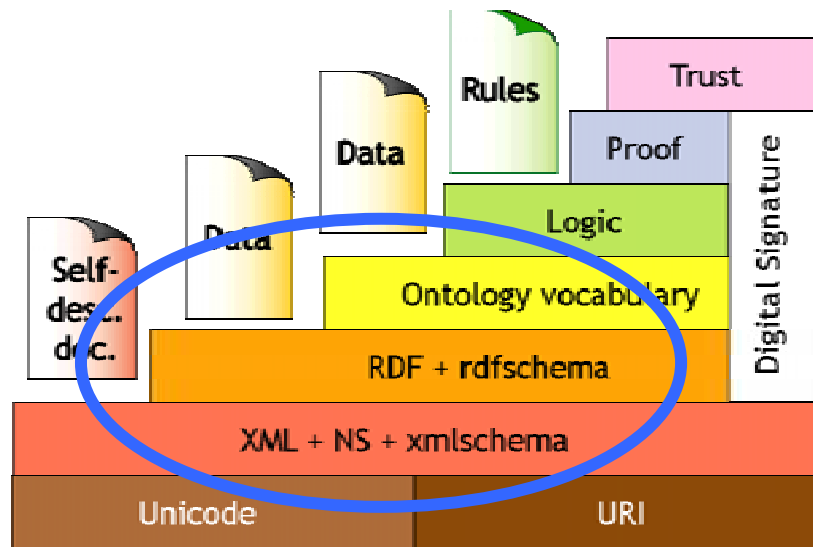


Figure 1. Layers of concern for the development of semantic standards in European Ethics.

The XML layer: Content of a resource

The XML layer is concerned with the description of what a document or resource is about. Inasmuch as data that is proprietary to an application has limited use, the XML layer provides for standardized means of describing content in order to free up that content for use by any number of applications.

Granularity

Domain-specific XML tags can be used to mark up the content of a resource at various levels of granularity ranging from the level of the resource itself (i.e., to describe what the resource is about using descriptor terms or abstracts) down to the level of the section or passage within the document (if it is a text), to the sentence-level, to the level of a single term (i.e., to describe what the term means or refers to). In the case of structured data such as databases, the database itself might be described, or the fields of the database might be semantically indexed.

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

The level of granularity at which data is indexed is directly related to the types of queries the user can ask and the types of results that can be retrieved. If resources are marked up coarsely, such as at the level of the website (or individual pages on a website) or metadata record, then the user's query, normally in the form of one or more keywords, will retrieve a set of links that either contain or are associated with (via, for example, matching strings in the resource's metadata) the user's keywords. This option limits resources to those that can be identified via a URL or metadata records retrieved via a SELECT SQL query. Contrastively, at finer-levels of marking up information, such as the chapter of a book, or a passage, the results retrieved can be more directly targeted to the user's query, more so than matching keywords against metadata describing a document or URL, which may not contain the information needed by the user. For instance, if the user is looking for information on the health hazards to humans of pesticide use in Africa, she might indeed find a document keyworded with (or having significant frequencies of occurrences of) "pesticides", "health hazards", and "Zambia," but the document might be about the removal of these substances, or the amount sold of those types of substances, or regulations about their use, etc., rather than about the ways in which they are hazardous. This is because there is little or no indication of what the relationship is between the terms used to describe a given resource and the resource itself, or among the terms themselves. When information is described below the level of the resource, retrieval results may match user queries more effectively. So an article containing a section on pesticides, health hazards, and Zambia would be indexed differently, and ranked at a higher position than one containing one section on pesticides and health hazards and another on Zambian culture. At the finest level of mark-up, where individual words are indexed, the system understands the meaning of each term in a sentence, as well as its relation to other terms. In such a system, the user can issue a well-formed question as a query, and the result elicited would be in the form of an individual sentence, based on an analysis of the user's query and a search for the best match among the sentences within the resources. For example, it would be possible for the user to input "What are the health hazards of pesticides used in Zambia?" and for a direct response to be in the form of a sentence drawn from resources, e.g., "Pesticide use in Zambia are associated with the following toxic effects."

Within the domain of ethics, content description is initially envisaged at the metadata level, within the RDF layer (cf. next section), using controlled vocabularies. Thus, what can be retrieved are resources such as experts, software, and document-like information objects² or DLIOs (as opposed to individual answers, discussion threads, or text passages). This is mainly for practical reasons, given that there already exist numerous bibliographical databases that contain descriptions of bibliographic metadata using controlled vocabularies. The database structures can be studied to develop the initial version of the EEAP and the EEAO, while the controlled vocabularies can serve as the basis of further developing the EEAO. However, it must be stressed that this is a starting point, and that more sophisticated systems can be developed once the EEAO has been extended using vocabularies containing rich semantics.

The RDF layer: Metadata of a resource

² A DLIO is a unit that is comparable to a paper document. The term is used to indicate resources such as websites, presentation files, photos, etc. but may not cover, for example, organizations or projects.

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

The RDF layer contains information about a resource, viewed externally, that is, from outside the resource, and includes information such as its title, author, and publisher. This information that describes a resource is called metadata. Standardized XML tags can be used to mark up metadata. For resource description, there already exist standards such as the Dublin Core Metadata Element Set³ (DCMES). What distinguishes, however, the lower XML layer, which merely describes a resource, from the RDF layer is that the latter is able to express relations between resources. In contrast to the WWW, where associations, i.e., hyperlinks between resources are meaningful to the extent that they are interpretable by humans (e.g., while a human could understand why a string *President Bush* would be hyperlinked to an image of the American flag, to a computer, such a relation would be indistinguishable from any other text that was hyperlinked to an image). RDF provides a standardized format for uniquely defining resources and a well-defined syntax for making statements about those resources. Figure 2 exemplifies the type of statements that RDF allows about a resource.

Subject	Predicate	Object
http://http://www.gmi.stanford.edu/pubs/SMI_Reports/SMI-97-0685.pdf/	http://purl.org/dc/elements/1.1/title	"A Study of Collaboration Among Medical Informatics Research Laboratories"

Figure 2. [resource] --dc:title--> v[dc:title]

As mentioned, for developing an integrated information service for the domain of European ethics, resources will be described using an *application profile* (AP), metadata elements that are drawn from one or more standardized metadata element sets that may also be extended and customized to the types of resources to be provided by the information service. This will allow resources to be described using standard mark-up languages that, on the one hand, are independent of local platforms and applications and, on the other hand, can enhance the possibility of semantic interoperability of resources within the domain. In section 3, we specify a methodology to develop an AP for European ethics.

The Ontology layer: Modelling the domain

In the RDF layer, resources are defined by virtue of their relationships to other resources. The ontology layer offers, in addition, the possibility of reasoning within the domain through precise specifications of concepts, relations, and rules, thereby creating the possibility of inferring new data from existing data. In other words, an ontology provides a knowledge model of a given domain that can interface with the RDF layer via mappings to its metadata elements. The model is made explicit via a knowledge representation language. Although many such languages exist, we use OWL Web Ontology Language⁴, the W3C standard knowledge representation language that offers rich semantics and is native to the Web (i.e., is serialized in XML).

³ Dublin Core Metadata Initiative <http://www.dublincore.org/>

⁴Cf. the OWL Web Ontology Language Overview (<http://www.w3.org/TR/2004/REC-owl-features-20040210>).

Towards an infrastructure for knowledge management within European Ethics

For the domain of European ethics, we distinguish two levels of knowledge to be represented: (1) One consists of the root ontology, where concepts, relations, and rules corresponding to the resource metadata will be specified for and mapped to the elements comprising the aforementioned application profile. (2) The other consists of all other ontologies derived from knowledge organisation systems such as thesauri and terminologies that can extend the root ontology. These other knowledge organization systems may provide a set of valid metadata values for resource attributes, or they may comprise an entire (sub)ontology in their own right that can extend the root ontology.

Methodologies for semantic integration

In the previous section, we outlined the parts of the Semantic Web architecture with which we will be concerned for the development of a semantic integration solution for European ethics, as we have defined it. In this section, we will define the terminologies we are using and discuss explicitly the rationales and the methodologies for the development of those components.

Ethics Application Profile

Definition and rationale

An application profile (AP) is a flexible, platform- and architecture-independent, information exchange format to facilitate the exchange of information resources via the Web for a given project or application. It consists of data elements (i.e., XML tags), drawn from one or more namespaces (i.e., named collection of elements and attributes), combined together and optimised for a given domain. By reusing elements specified in already-existing metadata standards, such as the DCMES, the AP transcends proprietary systems and organizational boundaries, and thus creates the possibility of improving management of and accessibility to domain-specific information materials.

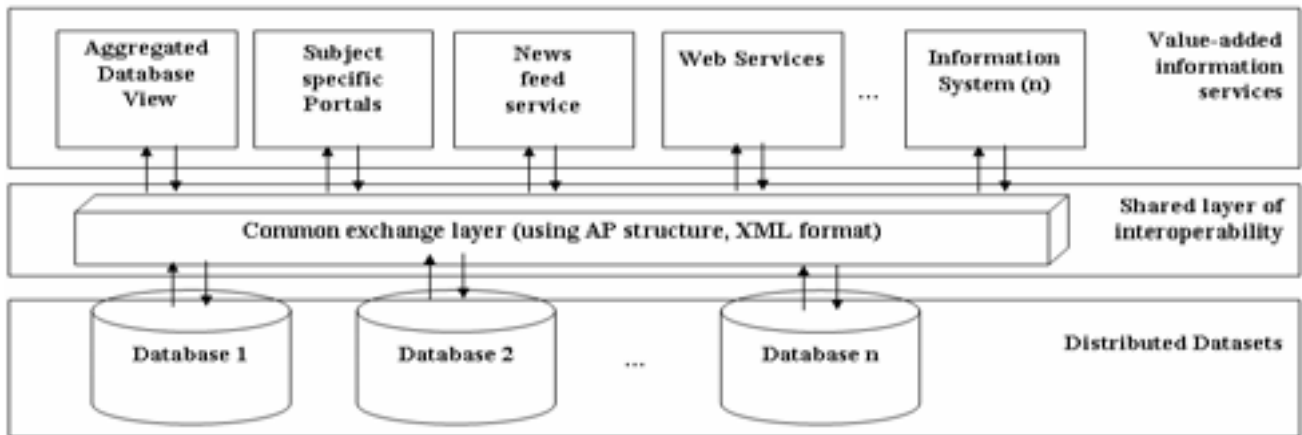
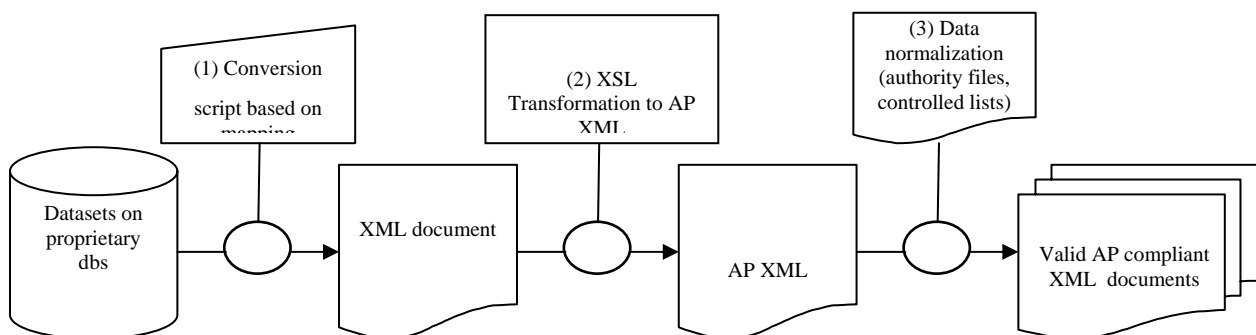


Figure 3: Interoperability between datasets allow for creation of value-added services and systems

Figure 3 shows the use of an AP as a common exchange layer to resolve the heterogeneity among information systems, and as a basis for the development of value-added services.

An AP prescribes the vocabulary, content, and structure rules that can be used to share information between heterogeneous datasets without requiring any change to the local system. With the possibility of using tools such as XSL Transformation (XSLT)⁵, the information extraction and conversion becomes a simple yet extremely important task towards facilitating interoperability. The fact that the resource itself does not have to be attached to the metadata makes it easy to control access rights on it.

The process of converting domain-specific data from a non-XML-native db to an AP is illustrated below, in Figure 4.



⁵ XSL Transformations (XSLT) Version 1.0 <http://www.w3.org/TR/xslt>

Figure 4: Converting from a proprietary dataset to the AP

The first step (1) involves mapping the local data model to the AP, and then either creation of a conversion script or direct export of data into XML (as allowed by some databases). In the next step (2), XSLT is applied to convert the simple XML document into the AP compliant XML format. In the case of native XML database this step might not take place. The last step (3) comprises data normalization, which ensures not only that the XML file validates correctly against the EEAP DTD but also that all the requirements, such as use of controlled vocabularies or declaration of used schemes, are met.

The next section describes the methodology for developing the AP in the domain of European ethics.

Methodology for the development of a European Ethics application profile (EEAP)

The creation of the EEAP will involve several phases.

Phase 0: Definition of the project, its goal, and its scope

The first and foremost task is to specify the goals of the project, especially in terms of its short-term and long-term objectives. This allows the scope of the project to be defined in order to keep the work within its boundaries. In the case of the project to develop an integrative information system for European ethics, the following goals are identified:

1. provide a platform independent exchange format that can alleviate the semantic heterogeneity characterizing the resources provided by the 200-some information systems identified thus far;
2. do the groundwork to enable information service providers of ethics resources to achieve digital information management standards for the next generation Semantic Web.

The solution for the first goal is clearly to provide a format, such as XML, that will not bind data owners to any specific information system yet allow them to share their data, regardless of the platforms and technologies they are using. These resources can remain distributed and can use either Hypertext Transfer Protocol or Web Services to perform multi-host searches, or they can be centralized to a single database. If they remain distributed, Web Services is recommended both for scalability and for automatic discovery of resources, especially for the future, when the adoption and implementation of semantic technologies (presumably) becomes more widespread. The second goal, which depends on the first, as well as on the development of the ethics metadata and subject ontologies (defined above in section 2), provides a means of converting data into machine-processable smart data, a prerequisite for transformation of the WWW into the Semantic Web.

As shown in Phase 1 below, a review of the resource types identified by the initial surveys of the ethics information services indicates that in all likelihood, it will need to be extended to include not only DLIOs but other entities such as institutions, individuals, and projects, and even services provided by collaboration software and texts generated by means of those services.

Phase 1: Assessment of the information objects

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

The major objective of this phase is to specify the range of resource types comprising the ethics domain. Within the domain of European ethics, the following resources have been identified thus far (cf. Hopt et al. 2004 and Denger 2004b)⁶:

- internet portals, link collections, personal web pages, web pages of institutions and organisations
- databases: institutes, experts, literature, press articles, multimedia files, bibliographic data, projects, events
- publications: journals, newsletters, book excerpts, online texts
- collaboration software: discussion fora, calendars, event notification service, etc.

Once the actual resource types to be accounted for has been determined, for example, through user surveys, web logs, etc., each needs to be analyzed to determine the properties characterizing it. Such analyses establish the initial requirements for specifying an ethics application profile. Non-traditional information objects such as institutions and collaboration services may also need to be considered. Note that different types of resources will be described using different criteria. For example, part of an adequate description of a book should include information such as its title and its identifier, which usually is expressed as an ISBN. An adequate description of a journal article should include not only the title of the article and its ISSN but also the title of the serial. In contrast, to describe an individual, information such as the employing institution, the individual's title, research areas, and email address, might be deemed crucial.

Phase 2: Assessment of the existing metadata standards and creation of the AP

As shown above, different types of resources exist within the domain of European ethics: DLIOs; non-DLIOs such as persons, institutions, events, and projects; and services provided by collaboration software.

Because many of the resources are in fact document-like resources found in digital library collections, a natural starting point from which to create an AP is the set of standards proposed by the Dublin Core Metadata Initiative⁷ (DCMI). The DCMI has come up with a set of widely accepted elements, refinements, and schemes, such as the DCMES, mentioned previously, for the description and discovery of resources. The DCMES is clearly defined yet shallowly scoped to serve the aim of wide applicability, i.e., cross-domain description, discovery and retrieval of information objects. It is also extensible in that additional elements, refinements or schemes may be added. However, this extensibility has to be controlled as it can be counterproductive to achieving the aim of interoperability.

Other entity types for which suitable metadata standards must be found or developed. For instance, for describing persons, the suitability of standards such as vCard and FOAF can be assessed.

Phase 3: Developing new properties.

⁶ Later revisions of this paper will account for survey data collected by WP2 on commonly accessed websites in ethics as well as types or resources searched for.

⁷ Dublin Core Metadata Initiative: <http://www.dublincore.org/>

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

Because currently available metadata standards may not be sufficient to cover all of the needs particular to ethics, rather than extend the current standards beyond recognition, a metadata element set specific to the domain, European Ethics Metadata Element Set (EEMES), can be developed to act as an umbrella namespace under which new elements that are deemed necessary for resource description in the domain of ethics can be declared.

The initial step is to determine a set of properties needed to describe the different resources available in ethics independently of any given standard. This process helps to decide if a particular property is really needed to describe the resource. Table 1 outlines the series of questions to be posed for determining the necessary properties.

Table 1. Determining the necessary properties to describe a given type of resource.

Is the elements/refinement/scheme really required for supporting:

- resource description?
- resource discovery?
- interoperability?

Then, with the properties deemed necessary (for description and searching), another sequence of questions, shown in Table 2, concerning each of these elements is posed, again in iterative fashion⁸.

Table 2. Determining the need for a new element.

Once it is determined that the need for a given property exists, then:

- Can the need be solved with a scheme value for an existing DC element? If yes, then create an EEMES scheme for an existing DC element, or else
- Can the need be solved with a refinement for an existing DC element? If yes, then create an EEMES refinement for an existing DC element, or else
- Can the need be solved by a qualifier from an existing non-DC set? If yes, then use that as a qualifier for the DC element, or else
- Can the need be solved by an element from an existing non-DC set? If so, then use that element, or else
- Create a new EEMES element (and, if necessary, a scheme).

This task of trying to match each property to an existing element, refinement, or scheme is meant to avoid reinventing the wheel. One consequence is that all declared elements, refinements and schemes in EEMES may end up looking like a hodgepodge. To make sense of them, they would need to be seen along with their parent. Two further steps are necessary for implementing this phase.

(a) Provide the ISO/IEC 11179 metadata for each element, refinement and scheme in the EEMES namespace

⁸ Adopted from original guidelines from Stuart Sutton (add URL to his presentation)

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

Once the elements, refinements, and schemes are given entry into the EEMES, they should then be carefully described using the ISO/IEC 11179 standard for the description of metadata elements. The use of the ISO/IEC 11179 helps to improve consistency with other communities and augments the scope, consistency, and transparency of the EEMES.

Ten attributes are used for defining the elements, all of which may be deemed suitable for meeting the project's needs.

Attribute Name	Definition
Name	The unique identifier assigned to the data element.
Label	Label assigned to the data element.
Version	The version of the data element.
Registration Authority	The entity authorized to register the data element.
Language	The language in which the data element is specified
Definition	A statement that clearly represents the concept and essential nature of the data element
Obligation	Indicates if the data element is always or only sometimes required (mandatory, optional, conditional)
Data type	Indicates the type of data that can be represented in the value of the data element
Maximum Occurrence	Indicates any limit to the repeatability of the data element.
Comment	A remark concerning the application of the data element.

The terms **Name** and **Label** are not as they appear in ISO/IEC 11179 and were modified to adhere to the terminology currently being used in the XML community. This approach is taken to facilitate the assimilation of this set into the XML and RDF communities.

Additionally, the following two attributes may also be used.

Attribute Name	Definition
Element Refined	The name(s) of element(s) refined.
Scheme	The applicable schemes for encoding the values of the term.

(b) Create the data model of the AP

The next step involves taking each of the terms and defining them in the context of European ethics. APs allow us to provide application specific definitions as long as they do not change the concepts themselves. For each element, we can provide its definition, cardinality, and data type information by giving some examples of best practice guidelines. These guidelines should try to cover as many scenarios as possible but cannot be exhaustive for practical reasons. They should suggest the use of schemes whenever possible; for example, the ISO639-2 scheme should be used to indicate the language, when necessary.

This process should be applied to all the elements and refinements.

Phase 4: Create an XML DTD or schema

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

The guidelines should then be converted into an XML DTD which would eventually be used to validate all the XML-based inputs to the European Ethics network. The XML DTD provides the following:

Logical structure of the record	The sequence and/or nesting of elements
Obligation	If a term is mandatory or optional
Cardinality	How often can this term appear in one record (0, 1 or more times)

Phase 4: Test the XML DTD using real data

The application profile should then be made available as both a document and also as an XML DTD, which is necessary for validating XML inputs. The guidelines should ideally be applied by an information provider for subsequent refinements of both the document and the DTD.

Once the first release of the AP is established, it will undoubtedly undergo revision. The technical implementers, i.e., those who would be responsible for converting their proprietary databases to the AP format, should be provided with documentation on how to handle the conversion. Each implementer would be given one-to-one feedback to help them successfully implement the exchange standard. In some cases, this could involve the creation of crosswalks or mappings between the local systems and the European Ethics AP.

Ethics Application Ontology

Definition and rationale

An ontology is a shared model of a given domain whose basic components consist of a vocabulary of terms, a precise specification of those terms, and the relations between them. Although an ontology has a structure similar to that of a taxonomy, the real power of an ontology comes from the ability to go beyond the information encoded in the structure to generate new information through inferencing. Using an ontology creates a separate knowledge layer distinct from any local information technology, information architecture, or application. It is more scalable than traditional methods of integration, where fields from separate data sources are mapped to each other. In traditional methods, the addition of a single database to be mapped to n databases requires n mappings from each field in the new database to each corresponding field(s) in the other n databases. Moreover, drawing the correspondences between fields from the new database to those in the others requires an understanding of the semantics of each field in each database. Thus, the task of integrating every new database to the system, or indeed, making a change to any one of the databases, becomes more and more unwieldy, increasing by an order of magnitude the number of mappings to be carried out.⁹ However, when the knowledge layer is abstracted away from the details of a specific application, each new system has only to perform a single mapping in order to communicate with the other systems. This facilitates management of and communication among otherwise heterogeneous systems.

⁹ $n! = n \cdot (n - 1) \cdot (n - 2) \dots 1$ Where n = number of databases that want to share information with each other.

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

Mapping the application profile to the application ontology

By definition, a standardized metadata element set consists of uniquely defined concepts that are in specific relations to each other. Whether explicit, as in the relationship between translations of corresponding resources, e.g., the **ags:isTranslationOf** element, or implicit, as in the relationship between a resource and its file name, the semantics of those elements can be expressed via an ontology.

The European Ethics Application Ontology (EEAO), is the root ontology of the system. The representation of resource metadata elements as an ontology is motivated by the recognition that, as far as a resource metadata is concerned, the normally underexploited semantics existing between extrinsic descriptors of resources could be used to enhance the user's information retrieval/knowledge acquisition experience. For example, nearly all bibliographic metadata contain the following assertions.

[resource] dc:creator v[dc:creator]

[resource] dc:subject v[dc:subject]

where *v* represents the *value* of the property.

A simple but useful inference that can be drawn from these assertions is

v[dc:creator] hasWrittenOn/hasPapersAbout dc:subject




Figure 5. v[dc:creator] hasWrittenOn/hasPapersAbout dc:subject

as shown in Figure 5.

An application such as a search engine could make use of such meanings not asserted by the metadata or the resource (e.g., to make suggestions to the user, to enhance the user's learning experience, etc.). Yet, rarely do bibliographical information retrieval systems take advantage of the ability to make these kinds of inference.

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

Other metadata standards describing other types of resources (e.g., events, experts, etc.) are treated analogously.

Figure 6 depicts the three-tiered organization outlining the relationships between the resource, the metadata elements from, in this case, the AP, and then the application ontology.



Figure 6. The three-tiered relation between Resource, Metadata, and Ontology.

The resource metadata is marked up using the AP. The AP elements and the relations between them correspond to concepts and relations in the application or root ontology that make it possible to generate new information through the extraction of inferences. To take a simple example, a search for any publication associated with a given organization could also yield, by inference, the email address of the institution responsible for that publication. Indeed, in the future, with the use of the W3C standard for OWL, this inference could be made even if the information about the email address and the publication were on different websites.

Query: “Kennedy Institute of Ethics” “publication”

Result: website containing information about the Kennedy Institute of Ethics Journal

Result based on inferencing: email address of Kennedy Institute of Ethics

The ongoing development of the EEAO

FASTER

Work Package 5 B: Communication Links

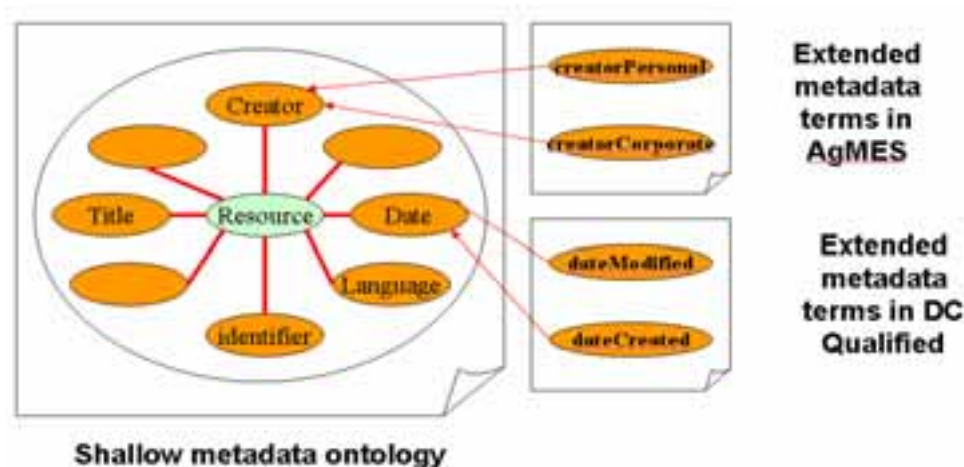
Towards an infrastructure for knowledge management within European Ethics

Using the EEAO as a root ontology to express the semantics of metadata descriptors (e.g., title, creator, publisher) considerably enhances the value of the resources that are described, even without the use of subject keywords to describe their contents. As mentioned above, formalizing the semantics of the metadata of European ethics resources already expands the knowledge provided by the ontology, via inferences, concerning resources. That is, by virtue of explicitly asserting that resources have certain attributes like author, publication year, and keyword, other information, such as number of articles written about a given topic within a range of years, can be extracted through reasoning across those assertions.

However, the ontology can be further extended beyond basic bibliographical metadata. On the one hand, subtypes of concepts already existing in the root ontology can be added. For instance, the **dc:creator** concept subsumes the sub-concept **ags:creatorPersonal**¹⁰, from which more specific concepts such as main author and alternate author can be derived. On the other hand, the root ontology can also be extended through the incorporation of controlled vocabularies. These vocabularies may simply consist of a flat list of terms, such as language codes. When the controlled vocabulary has some explicit semantics, as does a thesaurus, it lends itself to realization as a sub-ontology. These extension types are elaborated in the next two sections.

Addition of sub-concepts

As mentioned previously, Dublin Core, whose semantics is being used as the basis of the EEAO root ontology, was deliberately designed to be shallow. This shallowness allows for flexibility in its applicability to the specific needs of a given domain or application. For instance, a need may arise within European ethics to distinguish different kinds of authors (e.g., personal authors are often also experts and should thus be considered resources in their own right while corporate authors might be associated with institutions that can provide other information, e.g., events, etc., relevant to the domain). Consequently, the concept **dc:creator** can be further refined to the more specific concepts **ags:creatorPersonal** and **ags:creatorCorporate**.



¹⁰ See the examples of extending the DCMES at <http://www.fao.org/agris/agmes/Documents/Elements.html>

Figure 7. The shallow metadata ontology can be extended with terms

Controlled vocabularies

Controlled vocabularies are terminologies consisting of a set of terms and associated meanings that have been standardized for searching for and describing resources. They often represent the intellectual work of experts and/or standards bodies that can and should be reused both to avoid reinventing the wheel and also to increase the possibilities for interoperability. In general, we will be concerned with, on the one hand, the controlled vocabularies whose terms occur as flat lists without identifiable relations among them and can be used as valid metadata values (e.g., language codes), and those, on the other hand, with rich(er) semantics whose terms and relations can be used for the development of sub-ontologies (e.g., domain specific glossaries).

Valid attribute values

Their main distinguishing features are that they tend to occur as a flat file and/or as short lists without any systematic semantic structure among their elements. Examples include language codes, identification types for bibliographical resources, and abbreviations. These kinds of controlled vocabularies supply a set of valid values for specifying resource attributes. For instance, the values of the **scheme** attribute of the **dc:identifier** element identifies the type of number, e.g., ISBN, ISSN, etc., which indicates the type of resource, i.e., whether it is a book, a journal, and so on. Within the Agris AP, the allowed values of **scheme** have been defined to be **ags:isbn**, **ags:issn**, **ags:arn**, etc.

```
<dc:identifier scheme="ags:isbn">1234567</dc:identifier>
```

Using controlled values is another way of unambiguously defining the semantics of a concept, in this case, one of its properties. It enhances its interpretability by both humans and machines.

Sub-ontologies

Vocabularies that contain rich semantics occur with some indication of the meanings of their terms, whether accompanied by prose definitions where the semantics is interpretable exclusively by humans, as in a glossary or dictionary, or situated within a more or less machine interpretable context of terms linked through explicit relations, as in a taxonomy. These kinds of vocabularies tend to describe the concepts and relations that make up a given domain. In terms of the EEAO, the vocabularies in question describe the content of resources, as opposed to their external characteristics such as author, publisher, and so on. Vocabularies such as thesauri are a good starting point for ontology development because they already are to some degree machine readable. Vocabularies that are primarily meant for humans, such as glossaries, may be used to extract relations to other terms, explicit relations, and other information for incorporation into the ontology. With regard to the root ontology, these domain vocabularies can be developed into a sub-ontology attaching to the root via the **dc:subject** concept.

Towards an infrastructure for knowledge management within European Ethics

Within any given domain, especially one that is as interdisciplinary as ethics, several vocabularies are likely to exist and to overlap in their coverage of the domain. They may use different terms for the same concept. Moreover, these vocabularies may have slightly different understandings of a given concept or set of concepts. That is, one provider might use *gmo crop* while another might use *novel food* to refer to the same concept. One provider may view a *gmo crop* as a type of plant while another might view it as both a type of plant and a biotechnological invention. The sub-ontology can act as a mediating structure for multiple vocabularies within the same or overlapping domains. Because an ontology is concept- and not string-based, terminologies can map their specific terms to the corresponding concepts within the ontology. Further, with the help of domain experts, relations can be drawn between each uniquely defined concept. Consequently, providers can maintain the use of their terminologies while also being semantically interoperable with other vocabularies by integrating them based on a common semantic structure that can specify both terminological relationships (such as synonymy) and taxonomic and other semantic relationships (such as part-of).

Methodology for ontology-building

European Ethics Application Ontology

The EEAO is based on elements constituting the EEAP. Most of these elements are derived from already existing standards whose definitions and relations are already well-understood. Indeed, an ontology already exists for the Dublin Core Metadata Element Set (cf. Kamel-Boulos *et al.* 2001) that can be used both as a model and as a starting point for the construction of the EEAO. Therefore, the construction of this ontology should be relatively straightforward. Figure 8 shows the correspondences among the concepts derived from resource, the EEAP, and the EEAO.

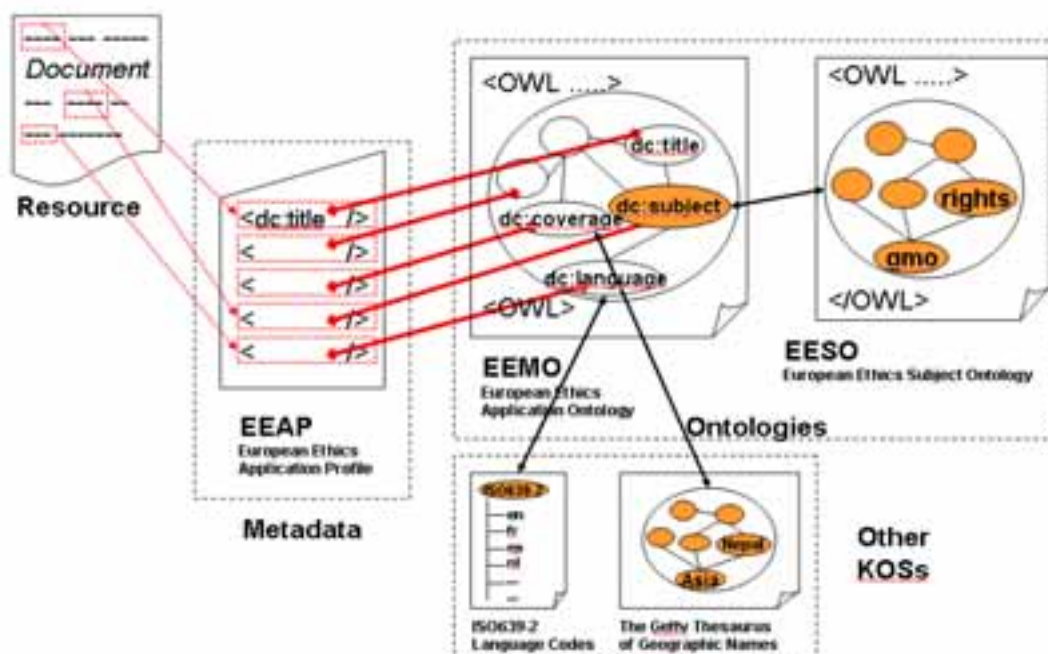


Figure 8. Corresponding concepts in the resource, application profile, and application ontology and the use of different knowledge organization systems

Extensions to the AO

Once the root EEAO has been developed, there may be additional extensions through the abovementioned subclassing of root concepts and the development of sub-ontologies. This section outlines the development of sub-ontologies using semantically rich controlled vocabularies. It considers the problem of integrating multiple controlled vocabularies covering the same domain.

Phase 1: Gather and characterize existing terminological resources in the domain.

In keeping with the principle of reuse (and in the service of interoperability), the first step is to identify the lexical resources that can furnish the raw materials, i.e., terms and meanings, from which to build the ontology. These lexical resources may involve semantics of varying degrees of explicitness (e.g., a word list only identifies concepts without definitions or relations; a taxonomy has some semantics expressed through terms connected via a hierarchy), that may or may not be machine-interpretable (e.g., a glossary is intended for human interpretation; a database scheme can be “understood” and used by a computer). They include glossaries, wordlists, thesauri, taxonomies, subject classifications, XML DTDs, and database schemes as well as ontologies. Figure 9 shows how these resources fall along a continuum, according to the explicitness of their semantics and their amenability to machine interpretation.




Figure 9. Terminological resources on a continuum of semantic explicitness. (Based on McGuinness 1999.)

Further, the degree to which the resource covers the domain in question, as well as the sub-domains covered, should also be assessed. For instance, a dedicated Ethics Thesaurus would obviously be relatively coextensive with the domain whereas a thesaurus such as AGROVOC, which is a general agricultural thesaurus, would contain only parts, scattered throughout the thesaurus, that were relevant. Other information that should be identified are:

- Concepts (and their equivalents in other terminological resources)
- Semantic relations, hierarchical and associative (e.g., RT)
- Top-level terms for those terminologies with hierarchical structure
- Depth of trees
- Classes v. individuals
- Annotations

Thus far, in ethics, about 20 terminological resources, whose content is of varying degrees of relevance to the domain, have been identified:

- Bioethics Thesaurus, published by the Kennedy Institute of Ethics
- Euroethics Thesaurus, the multilingual thesaurus in Ethics

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

- INSERM Thesaurus
- Parts of MeSH, published by the NLM
- Parts of CSA Life Sciences Thesaurus
- Parts of Biosis Controlled Vocabulary
- Etc.

With the help of subject matter experts (SMEs), those parts that are relevant to the domain and need to be incorporated into the ontology would need to be identified.

Based on the findings in phase 1, the next phase can be implemented.

Phase 2: Determine correspondences between individual terminologies and ontology model.

This phase requires analysis of each of the individual terminologies to establish the correspondences to be made to the data model of the ontology. Thus, for example, for multilingual thesauri, the following correspondences are assumed:

- Terms are treated at a level distinct from the conceptual one to distinguish between meaning and form and to allow for lexical relations such as synonymy and translation.
- Terms are treated differently from strings to allow for within-term relations such as acronymy, spelling variants, etc.
- Concepts correspond to classes.
- BT/NT are converted to superclass/subclass relations.
- RT is reserved for non-hierarchical relations.
- An individual is distinct from and a member of a class.
- USE/UF may or may not correspond to synonymy relations (e.g., many of the USE/UF relations in AGROVOC actually correspond to super/subclass relations, antonyms, semantically related but not synonymous relations, indeed, anything but synonymy).

Note that other terminologies may have other correspondences, e.g., in an XML DTD, elements may correspond to concepts; in a glossary, each term might correspond to a concept while relations to other terms or concepts might be derived from informal definitions.

Phase 3: Convert terminologies, following the correspondences made in Phase 2, into OWL.

This stage involves the transformation of the correspondences made in Phase 2 to a knowledge representation language. Each transformation should retain information about the source terminology. The following exemplifies the concept-to-class correspondence.

Bioethics Thesaurus:

human rights ==> <owl:class rdf:id="HumanRights@BioThes">

Thesaurus di bioetika:

human rights ==> <owl:class rdf:id="HumanRights@ThesdiBio">

Phase 4: The core subject ontology: Subject matter experts (SMEs)

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

The core subject ontology consists of the fundamental concepts, relations, and assumptions for a given domain. In this phase, SMEs are provided with a set of questions to identify these basic notions for building the ontology. Within European Ethics, a dual-pronged approach to define the domain was suggested (cf. WP1) combining the problem-oriented perspective (e.g., laws governing the growth of crops with genetically-modified organisms) with the real-world one (e.g., genetically-modified organisms as a product of biotechnology). Thus, two sets of question types would be drawn up, e.g.,

Problem/issue oriented:

What issues exist that are of concern to anyone, etc.?

Why are they of concern?

To whom are they of concern?

Whom do they effect?

Real-world oriented:

What is the purpose/function of X?

What does it consist of?

The objective of this stage is to specify the domain-specific concepts and relations at the highest level of abstraction. The ontology that is developed at this stage then can serve as the foundation for the hierarchies identified and extracted in the next phase.

Phase 5: Identify hierarchies within terminologies.

Once the resources have been identified, they need to be classified according to the degree of explicit structure contained in the resource. Those terminologies containing hierarchical structures can be (re)used to build the structure of the ontology while those with semantics meant for human interpretation such as glossaries can serve to provide synonyms and annotations.

Phase 6: Alignment.

These top terms along with their hierarchies are then aligned to the core domain ontology created in Phase 4. If a corresponding class does not exist, and the term is pertinent to the domain, the relevant part of the core domain ontology is enriched to create a place for alignment. At least some parts of this procedure can be done automatically based on tools using a string-matching algorithm. Figure 10 shows a graphic of how alignment is done using the example given in Phase 3..

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

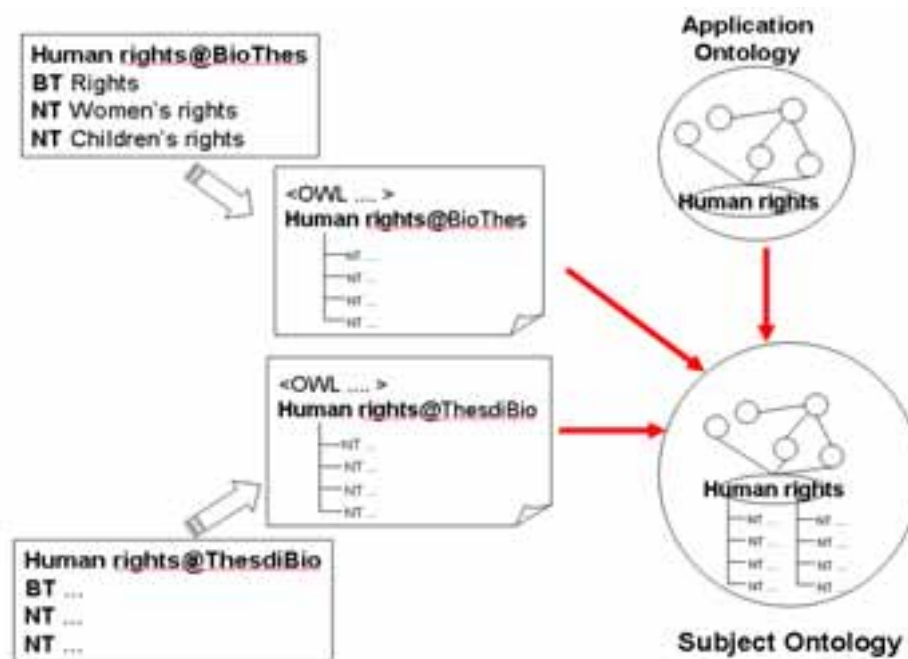


Figure 10. Aligning hierarchies from different terminologies to the core subject (i.e., domain) ontology.

Phase 7: Merging.

Merging is the process of integrating corresponding concepts from the source terminologies. In some cases, this is trivial. For example, the concept { stem cell } is homonymous in the Bioethics Thesaurus, Euroethics thesaurus, and AGROVOC, the concept { stem cells }, i.e.,

Bioethics Thesaurus: stem cell

Euroethics: stem cell

AGROVOC: stem cells

→ core: stem cell

Thus, the three sources are realized in the core domain ontology as lexicalizations of the same concept. Tools¹¹ are available to help SMEs with this process.

In other cases, homonymous terms may refer to different concepts. In AGROVOC, *euthanasia* refers to putting animals to death, while in the Thesaurus di Bioetika, the context in which it occurs suggests that it refers to the putting a human to death as a treatment for certain medical conditions.

AGROVOC:

Euthanasia

USE Destruction of animals

¹¹ Prompt and (<http://protege.stanford.edu/plugins/prompt/prompt.html>) and Chimaera (<http://www.ksl.stanford.edu/software/chimaera>) from Stanford are among those tools developed for managing multiple ontologies.

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

Thesaurus di bioetika:

Euthanasia

BT Killing

NT Active euthanasia

NT Involuntary euthanasia

NT Allowing to die

NT Voluntary euthanasia

RT Extraordinary treatments

RT Coma

RT Right to die

RT Terminally ill

RT Treatment refusal

RT Living will

For terms referring to the same concept, e.g., *GMO crop* and *novel food*, a SME is required to make the determination of synonymy.

Phase 8: Enrich through annotations.

Comments (e.g., glosses, contexts, etc.) and identifier numbers can be mapped to the concepts. Synonyms and translations are also treated as annotations.

Phase 9: Post-processing.

Once alignment and merging have taken place, the resulting structure should be checked for inconsistencies. The exposure of such inconsistencies is facilitated using tools¹² built for such purposes.

XThes: DisjointClasses(a:female a:male)

XThes: Class(a:Sam partial a:female)

XThes: Class(a:Sam partial a:male)

→ **inconsistent**

Scenarios: Relating ontologies and data

Thus far, we have been placing emphasis on the use of the AO for retrieving information from structured metadata repositories. In this section, we broaden the discussion to the relationship between ontologies and data in general. That is, we show that it is possible

(1) to make use of the AO to search structured data that has not been indexed;

¹² For instance, the OWL plug-in in Stanford's Protégé tool is able to highlight logical inconsistencies in an ontological structure.

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

- (2) to make use of an ethics domain ontology to search unstructured data;
- (3) to make use of a semantically enriched AO at various stages of complexity.

The continuum depicted in Figure 9 shows the kinds of resources that can be incorporated into the AO to enhance the latter's semantic richness. The scenarios will be described using vocabularies taken from different locations on the continuum.

We first outline the main architectures within which the AP and AO may be contextualized. These vary along two dimensions, namely, the centralization/distributedness of the AO, on the one hand, and of the data repositories described by the AO, on the other. Then we show that while the development of semantic structures is an ongoing process, each stage can be used as the basis for solving many of the problems currently plaguing heterogeneous domain-specific information retrieval and knowledge management systems. The richer the semantics, the more possible it becomes to develop “smart” applications that not only increase the effectiveness of IR/KM but also enhance the user's experience in the process of search and retrieval.

The overall architecture integrating the AP and the AO is depicted in Figure 11. However, as shown, the system works with both structured and unstructured data.

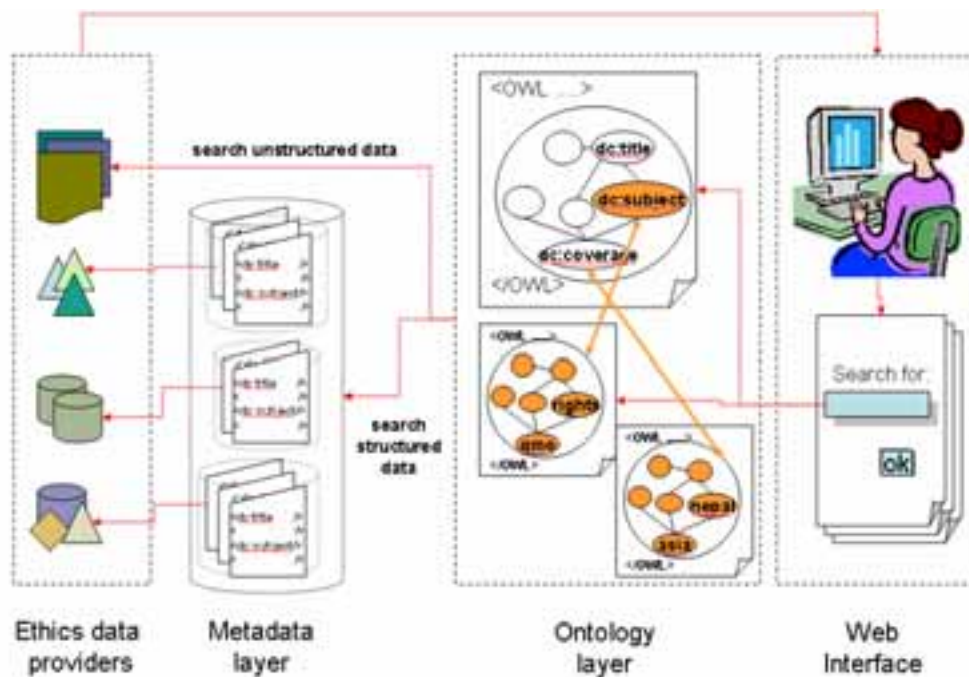


Figure 11. Integrating the ontology layers with the metadata layers.

Mediated (centralized) v. federated multi-host (distributed) architecture

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

In Figure 3, we showed graphically the role of the AP as a layer mediating, on the one hand, between heterogeneous, distributed datasets and, on the other, applications making use of the information contained in those datasets. Once owners of distributed databases have mapped their records to and exposed their content in the AP format, this mediating layer may be realized in one of two ways. Either the resulting metadata is sent to a centralized database or it is made available in individual databases that are accessible to multi-host searching. In either case, the corresponding AO will be made available in a pre-defined location so that the metadata vocabulary can be interpreted (i.e., ascribed meaning). In the future, the AO in its entirety may be stored in a registry of ontologies and accessed dynamically via Web Services technologies. However, the use of OWL to describe the AO will allow distributed storage, maintenance, and enhancement of the ontology. Applications based on the development of this AO will see it as a single ontology.

In the mediated approach, all metadata is stored in a single centralized location. Queries are first pre-processed (e.g., parsed, spellchecked, normalized for singular/plural forms) and then interpreted via the AO, where they are resolved to concepts or instances. These concepts or instances are then matched to the relevant fields within the central database. Note that, as mentioned, the AO itself may be centralized or distributed.

In contrast, in a federated architecture, databases are stored locally and made available for WWW access. As in the mediated approach, queries would undergo pre-processing and interpretation via the AO. But rather than conducting a search on a single database, the interpreted query would be sent via web services to databases hosted on multiple distributed servers and a search executed on each of those databases.

Thus, there are two dimensions to the building of the architecture: the centralization (or not) of the data, i.e., resources, and the centralization (or not) of the domain knowledge that describes those resources.

Three scenarios

The AO and unindexed structured data

In the first scenario, the AO is exclusively based on the AP, and consequently, consists only of concepts describing resources. That is, it is the root ontology without any further extensions. In this case, relatively little analysis is required for the development of the ontological structures involved.

The concepts concern only those used to describe the extrinsic properties of resources. Controlled vocabularies or lists may provide values for the attributes of the resource, e.g., language, keyword, etc., and in such cases, should be specified; however, they are not essential for this scenario.

In and of itself, the AP already resolves, in a substantial way, the issue of semantic heterogeneity and the associated problems of maintaining interoperability among multiple distributed data repositories. However, further explicating the semantics of the AP in the form of an ontology provides the foundation for the development of semantically driven applications.

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

For instance, an information retrieval application based on the AO could have the following features.

- The user could search resources by resource type, author, year, publisher, organization, etc.
- The user could issue complex queries, such as checking to see if a given author wrote any articles written within a given time period, or if there were any current projects ongoing.
- The system could present the user with information related to her query, such as titles associated with a particular author, co-authors with which a given author has written, and so on.
- The user could learn about resources and their properties, e.g., discovering that the same individual who led a project promoting a certain technology also wrote a paper arguing against it, by traversing the links in the ontology.

The SO and unstructured data

The second scenario comprises the use of the ethics subject ontology and one or more unstructured data repositories. In this case, the SO is used primarily to help the user form effective queries using the terms and relations in the ontology for query expansion. The actual query expansion that is realized depends on whether the data repository is domain-specific, or if it is a general one such as the WWW. The richer the vocabulary (synonyms and translations), the more effective the search.

For instance, if the user is looking for information on BSE within domain-specific repositories, and he issues the query term BSE, the query would be expanded to include all of the synonyms for that concept. This ensures the greater recall of resources.

Domain-specific context

User's Query: BSE

Expansion: BSE OR bovine spongiform encephalopathy OR mad cow disease

Because the search context is limited to the domain, the problem of false positives is reduced. That is, in a repository of information on, say, bioethics, the string BSE is more likely to refer to the *bovine spongiform encephalopathy* concept than the *Bombay Stock Exchange* and thus retrieve the correct results.

Within general data repositories, the expansion strategy may be slightly different. In this case, the abovementioned query would be much more likely to retrieve false positives. To minimize ambiguity, the query would then be expanded to include the parent concept, with which it would be combined with the Boolean AND.

General context

User's Query: BSE

Expansion: (BSE OR bovine spongiform encephalopathy OR mad cow disease) AND (disease OR syndrome OR disorder)

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

By expanding the query to include the parent term, which disambiguates the term, false positives arising from the ambiguity of the original query are reduced.

When the contents of a repository are determined (automatically or manually) to occur in a particular language, the terms in the appropriate language can be used to expand the query in the same manner as described above. When the repository is general and multilingual, as is the WWW, a simple OR query consisting of all synonyms and translations corresponding to the user's query could be issued.

Thus, in addition to the following functionalities,

- aided query formulation through query expansion
- multi-lingual search

the SO also enables the user to

- learn the domain vocabulary as well as the domain itself
- search related terms

The AO + one or more thesauri

The third scenario comprises the root ontology with one or more thesauri containing the usual BT/NT/RT/UF relations. In contrast to the previous scenarios where semantic relations exist only between metadata descriptors but not between terms within controlled vocabularies, or only between concepts specific to the ethics domain, in this scenario, the controlled vocabularies supplying values to one or more of the concepts, usually **dc:subject**, contain some explicit semantics. Although the terms and relations within thesauri are often ill-defined and cannot be used for reasoning, they nonetheless contain some machine-readable semantics that can be exploited for the development of intelligent applications.

An information retrieval application based on this extended ontology can have, in addition to the features described above in the previous two scenarios, the following:

- Resources containing terms related to those in the user's query, as well as the terms themselves, could be displayed.
- If the user enters a query, she receives a list of results containing the metadata for the articles associated with those keywords. Clicking on an author's name retrieves all the resources containing those same keywords.
- The system can help the user find the information she is searching through a series of questions that filter through the information, e.g.,
 - what resource type? (e.g., author)
 - wrote what publication type? (e.g., introductory text)
 - about what? (e.g., keywords k1, k2, and k3)

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

- when? (e.g., between the years y1 and y2)
- in which language? (e.g. language l)

The AO + sub-ontology

The third scenario is the most complex. It consists of the root AO containing concepts corresponding to all the resource types. Where a given concept is associated with multiple controlled vocabularies, those vocabularies are integrated. If they consist of flat lists, a list akin to an authority file can be incorporated into the ontology. If they contain a more complex structure, a sub-ontology can be developed that integrates the different terminologies (cf. Section 3.2.4.1.1).

Moreover, in this scenario, the knowledge itself may be distributed. Thus, parts of the root and sub-ontologies may be stored on distributed servers. Access to these distributed parts may be made via URL references or through web services.

All functionalities described in the previous scenarios are possible. Because the semantic structures are truly concept-based, other features such as

- cross-language information retrieval,
- terminology brokering across multiple databases using different terminologies,
- intelligent query expansion

become possible.

In addition, more sophisticated applications can be developed. For instance, an customizable information delivery system can filter information for people needing to monitor and assess large volumes of information. The volume of targeted information is reduced based on its relevance according to the user's "need to know." In a real-time monitoring system consisting of online RSS news feeds, the user could enter parameters of interest (i.e., concepts). A change or update in information that conforms to those parameters (i.e., that contain those or child concepts) could trigger an alert. In a well-designed robust ontology, information can be filtered independently of language or specific terms used.

A robust ontology also serves as the basis for automatic indexing of texts at multiple levels of granularity. Based on the usual statistical analyses of term frequencies, terms can be resolved to corresponding concepts in the AO (and indeed to those in corresponding vocabularies). Documents can then be tagged with those concepts. At the content level, semantic tags can be provided that allow a resource to be "better known" by one or more systems so that search, integration, or invocation of other applications becomes more effective. Tags are automatically inserted based on natural language analyses of texts.

Conclusion

FASTER

Work Package 5 B: Communication Links

Towards an infrastructure for knowledge management within European Ethics

We have seen how semantic standards promise to be an effective approach to resolving the problems posed by semantic heterogeneity and how they can be the source of applications that help users find and discover information efficiently and effectively. The development of semantic technologies is an ongoing process, whereby any given stage can be the source of application development. The richer the semantics, the greater the possibilities for developing smart applications.

Contrary to popular belief, it is not the case that large expenditures of time and effort are necessary to develop nor to enjoy the advantages of semantic technologies, nor is it the case that structured indexed data are necessary to realize the benefits. As we have shown, even a small investment in the enhancement of relations between vocabularies, both metadata and domain-specific, yields a relatively large return on investment. We have shown several scenarios of varying complexity that enable information providers within the ethics domain to exploit semantic technologies to provide information effectively and allow their users to access it easily.

References

2000. Berners-Lee, Tim. Semantic Web on XML. Keynote presentation for XML 2000. Slides available at: <http://www.w3.org/2000/Talks/1206-xml2k-tbl/slide1-0.html>. Reporting available at: <http://www.xml.com/pub/a/2000/12/xml2000/timbl.html>
2004. Hopt, Sabine, Angelica Maria Kappel, Dirk Lanzerath, Jens Wolff. Overview of existing information services in the field of ethics. Unpublished manuscript.
- 2004a. Denger, Julia, X, Y. Linkliste mit Thesauri und Datenbanken zum Themenbereich Ethik/Philosophie bzw. Grune Gentechnik. Unpublished manuscript.
- 2004b. Denger, Julia, et al. Database of web sites (WP 1A). <[http://merkur.informatik.rwth-aachen.de/bscw/bscw.cgi/d1206462/Database%20of%20web%20sites%20\(WP%201A\)](http://merkur.informatik.rwth-aachen.de/bscw/bscw.cgi/d1206462/Database%20of%20web%20sites%20(WP%201A))>
2004. Burgess, M. et al. Survey data - 09 July 2004.
2001. Kamel-Boulos, M. N., A. V. Roudsari, and E. R. Carson. Towards a semantic medical web: HealthCyberMap's Dublin Core ontology in Protégé-2000. Presented at the Fifth International Protégé Workshop, SCHIN, Newcastle, UK, July 2001.
2002. McGuinness, Deborah L. Ontologies come of age. In Dieter Fensel, Jim Hendler, Henry Lieberman, and Wolfgang Wahlster (eds.), *Spinning the Semantic Web: Bringing the World Wide Web to Its Full Potential*. MIT Press.
2000. Heery, Rachel and Manjula Patel. "Application profiles: mixing and matching metadata schemas". *Ariadne*, No. 25, September. <http://www.ariadne.ac.uk/issue25/app-profiles/intro.html>