

Semantic learning object repositories

Jesús Soto Carrión

Elisa Garcia Gordo

Languages and Informatic Systems Department,
Pontifical University of Salamanca,
Madrid Campus, Spain
E-mail: {jesus.soto, elisa.garcia}@upsam.net

Salvador Sánchez-Alonso

Information Engineering Research Unit,
University of Alcala, Spain
E-mail: salvador.sanchez@uah.es

1. Introduction

A Semantic Learning Object Repository (SLOR) is a software system that stores educational resources and their metadata (or only the metadata), and provides some type of searching interface to human or to other software systems. Repositories provide access to the collections of educational resources, usually in electronic format, although most of them do not store the educational resources themselves, but only their metadata. Therefore, it is possible to find the same resource from different repositories. The main functionality of a learning object repository is that of allowing for the search of educational resources. From this perspective and, in a broad sense, the following two types of repositories exist:

- 1 Interactive searching interfaces, to be used by humans.
- 2 Querying interfaces, mostly aimed at being used by software agents (for example, through web services).

Sometimes, the same repository includes a searching method that can be used to both usages. Nevertheless, it must be noted that most usual general propose information retrieval mechanisms (Baeza-Yates and Ribeiro-Neto, 1999) (like internet searching) must be complemented with metadata browsing mechanisms. The simplest way to implement this is allowing metadata field searching. However, very often those interfaces are not satisfactory. The research reported in this paper departs from this point, and focuses the potential of ontology-based techniques to enable knowledge use about metadata domains to provide enhanced behaviour and improved performance.

Recent studies have attempted to unify different learning object definitions (McGreal, 2004). These studies show that learning-oriented entities in a repository have a high variability on its characterisations. The non-existence of a common vocabulary, as well as the coexistence of different learning object definitions, point out the need of flexible repositories that can fit all existent conceptualisations.

Formal ontology as a discipline (Welty and Guarino, 2001) is aimed at studying possibilities, so that it can be used to compare learning element representations according to the flexibility of their coverage, and term subsumption properties. In fact, ontological representations can play an important role as support for sound semantic models that fulfil a number of new requirements related to automation, such as search, retrieval or composition of new learning materials from others that already exist. The existence of ontology-based schemas becomes essential when some of the functions are to be delegated to automated or semi-automated systems, following the semantic web vision (Berners-Lee, Hendler and Lassila, 2001).

This paper shows the progress of a research in course whose main goal is the design of an ontology schema capable to both bringing more flexibility to the description of the entities stored in SLOR and allowing automated functions and/or task delegation to agents. The most significant contributions of this work are the following:

1. Semantic description of standard IEEE LOM¹ metadata.
2. Resource description with several ontologies.
3. Design of the SLOR core basic functions that enable a variety of ontological characterisations about the learning object concept.
4. Upgrade searching and browsing functions.
5. Generalisation of the concept of 'semantic repository'.

Those contributions are the starting point to others applications that will eventually make use of the proposed technique. In the rest of this paper, learning object repositories lacks are outlined. Later on, some of the benefits of the so-called semantic repositories are exposed, as well as some problems found during the SLOR prototype development. Finally, conclusions and further work are provided.

2. Lacks of current repositories

Currently, learning object repositories, such as MERLOT or CAREO, describe web-oriented learning resources by storing metadata records linked to them (in a general manner). Thus, they guarantee best search results on the basis of a fixed structure of the knowledge stored. Nevertheless, better searching capabilities are not the only advantage provided by these repositories: others, like cooperative review work about learning objects are definitely remarkable, since the quality of the content can be this way reviewed, analysed and discussed by several users of the repository.

It is expected that these repositories can play an important role about e-learning in the near future. Both humans and software agents (such as Learning Management Systems – LMS) will be capable of querying and searching the information stored in them. However, to execute some reasoning or inference tasks, quality metadata information are extremely necessary. For the purpose of this work, 'quality metadata information' is defined as the information that has accomplished, at least, minimum requisites of formal description and where the data provided are based on a pre-established, uniform and preferably universal formal schema. Current repositories lack a conceptual model establishing what a learning object is and what its metadata descriptors are, as this is often related to each of the many different conceptualisations available. However, without a universal agreement about the metadata model to be used and no certainty about the requirements of an utter formal description, there is an important shortage that complicates the automation of management tasks in these repositories. Currently, the quality of metadata records relies on the following factors (among others):

- The learning object creator's willingness to associate metadata information at the time of adding their materials to the repository.

¹ <http://ltsc.ieee.org/wg12/par1484-12-1.html>

- The editing capabilities or tools provided by the repository for learning object creators to associate metadata information to their materials. These capabilities heavily restrict the kind of information that can be stored and its level of formalisation.
- The level of cognition or instruction of learning object creators on the existing metadata practices and standards.
- The conceptual model of the repository, that is to say, what do the repository creators understand that a learning object is, and what structure the metadata information associated to it should have.

Writing metadata information about the learning objects, in the form of records describing their content, facilitates not only a number of processes such as storage, search and retrieval from distributed repositories but also the composition of new learning materials as an aggregation of others. Accepted metadata specifications and standards make learning objects interoperable and reusable, even though a number of shortcomings regarding current learning object metadata still exist:

- Current standards are purposefully descriptive instead of normative: they are intended to give information about the contents or the format of the learning object, but do not generally entail explicit run-time semantics for LMS that use the learning object (Sánchez-Alonso and Sicilia, 2005). An exception to this is the IMS Simple Sequencing Specification, which allows representing ‘the intended behaviour of an authored learning experience such that any learning technology system can sequence discrete learning activities in a consistent way that includes explicit runtime support’ (IMS, 2003).
- The information in a learning object metadata record is not, as currently defined in international standards such as LOM (LTSC, 2002), machine consumption-oriented. In fact, most metadata in current learning object repositories are no more than an overall content identification and description, thus providing limited value from the viewpoint of delegation. In addition, information in metadata records is mostly in the form of unstructured texts written in natural language, such as faculty member’s peer reviews and comments, a kind of information that software agents would find to be difficult for processing. All this hinders the possibility of programming applications capable of ‘behaving’ according to the information in a learning object metadata record (for example, to process a sequence of actions into an LMS according to object metadata information).

All the shortages indicated prompt a problem when it comes to software agents’ interpretation about the existing elements in a repository, for example, about the knowledge that might be interchanged by those agents. If this knowledge is formalised in cognitive structures (called ontologies), the use of ontologies, which provide formal models of conceptualisation, is thus the only way to process and understand the knowledge holding within metadata records.

Other remarkable problems are related to communication and interoperability. Because of the fact that these repositories do not include any restrictions, and do not define any interchange schema, their only advantage appears to be the web interface provided. For example, knowledge interchange functions between different kinds of repositories cannot be accomplished, since there is neither any agreement nor any recommendation to fit this type of transactions. Furthermore, no software agent could accomplish an autonomous behaviour by itself, as the choice of one learning object or another (all having the same didactical objectives) rely on the pre-conditions stated by the learner. Other examples of current disadvantages are the impossibility of creating different levels of trust for users in the repository, or the inability to add new schemas to repositories with the objective of storing metadata records holding a planning by cost in accordance with students’ knowledge levels (Soto and García, 2005).

3. Towards a flexible definition of learning object

Recent studies about the unification of existing conceptualisations of learning object into a single definition suggest the possibility of coexistence of all the available definitions (McGreal, 2004). According to

McGreal's study on existing learning object characterisations, five definitions – ranging from general to specific – coexist:

1. Anything and everything
2. Anything digital, whether it has an educational purpose or not
3. Anything that has an educational purpose
4. Digital objects that have a formal educational purpose
5. Digital objects that are marked in a specific way for educational purposes.

According to the first definition, the use of an object is what determines whether or not the object becomes a learning object, and thus everything that exists (the universal concept) can be considered a learning object. Nevertheless, in software-based representations, the only objects considered in practice as existing (in the sense of being able to 'talk' about them) are those that are represented. In our case, the scope of representation is that of the different elements of the ontology. Taking OpenCyc as a case of ontological representation, the term Thing, defined in Opencyc as 'the collection which, by definition, contains everything there is', subsumes anything that may eventually be considered a learning object. This definition has the obvious drawback of not adding any defining characterisation to the concept.

Learning can be considered as an Event, defined in Opencyc as 'a dynamic situation in which the state of the world changes'. Accordingly, everything that is linked to the representation of learning activities, or declared to have educational purpose in some way, should be considered as a learning object. In addition, some axioms could automatically classify some things as learning objects. For example, 'every Book classifies as a learning object'. These are the examples of concrete characterisation of classes of learning objects, which can be used for practical applications.

In consequence, the first definition may be interpreted in the following way: '[1] LearningObject-AsAnything: learning objects are things that either have been used in learning events or have been provided with descriptions that specify possible usages in learning'. The latter part of the sentence still requires much clarification, but it can be used provisionally until more detailed clarifications are proposed. An example of learning object that fits this definition would be a textbook, a pen or a printed copy of this paper. Figure 1 depicts both this and the rest of the definitions being discussed.

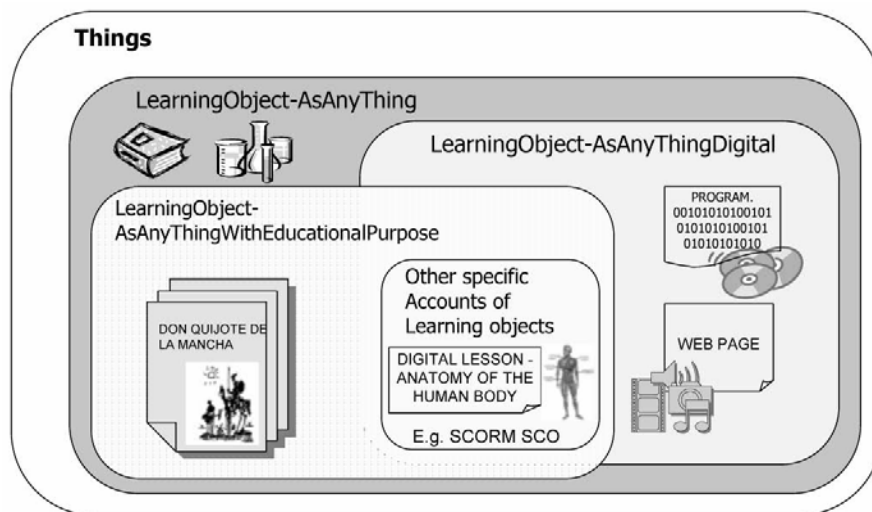


Figure 1 McGreal's study: different conceptualisations of a learning object

The second definition introduces the concept of 'digital object' in an attempt to further specify that learning objects are artefacts. The term `ComputerFileCopy` is defined in `Opencyc` as 'an information bearing thing that is identified as a unit by a unique name, and which is object-like in an important respect'. Examples include individual image files, text files, sound files and executables stored on some `ComputerStorageDevice` (defined as 'the collection of devices used by computers to store information'). Consequently, the term `ComputerFileCopy` can be used as a possible characterisation of the concept, since it requires unique identification, and is not restricted to 'data' but instead subsumes programmes in a general sense. Thus the following definition can be used '[2] `LearningObject-AsAnythingDigital`: learning objects are `LearningObject-AsAnything` instances that are subsumed by `ComputerFileCopy`'. An example of learning object that fits this definition would be the PDF version of this paper.

The third definition introduces a consideration of purpose. In this case, the purpose should be interpreted as something that was present in the act of Designing ('the act of designing something, be it clothing, cars, computer chips or buildings', as defined in `OpenCyc`) the learning object, which entails the associated restriction of learning objects to be Artefacts, i.e. at least partially tangible things intentionally created by an Agent (or a group of Agents working together) to serve some purpose or perform some function, which separates them from 'natural' things. This leads to definition [3]: '`LearningObject-AsAnythingWithEducationalPurpose`: learning objects are `LearningObject-AsAnything` instances that have somewhat a record of the educational purpose put in the object in the act of its Designing'. Here 'record' is used in a generic sense and it may simply include the trace of the one that created it. Further, it can be considered that some tacit traces of an object again exist, as having being created with an educational purpose. An example of learning object that fits this definition would be a digital unit including both a PDF version of this paper purposefully designed to serve a particular educational purpose and a metadata record including specific information on the pedagogical use and goals of the paper in an educational context.

Since the purpose in the design is an intellectual process, a notion of 'record' of it should be introduced. Such purpose may be internal to the learning object, e.g. the 'objectives' section in a web page, but it could also be tacit, i.e. when it takes a form that is easily recognisable as an educational artefact. This may be the case of slide presentations. That notion of 'record' of the purpose is deliberately kept open to divergent interpretations.

4. The SLOR ontology

As all the definitions described in the previous section can fit together, the recognition of such a diversity of conceptualisations and the taxonomisation of all of them could be the basis of a neutral conceptual model. This new model would eventually provide the users with a number of different functionalities, adapted to each particular concept of learning object, and not necessarily restricted by only one of them. The flexibility is dealt with as a key issue for guiding the implementation. Based on a sound semantic model, defined by an ontology schema engineered using the Web Ontology Language (OWL, <http://www.w3.org/TR/owl-features/>), such a model would include all the definitions in McGreal's study, thus supporting different types of learning objects. In short, the repository clients – final users, agents and LMSs – could, among other functionalities, add, retrieve, modify and search for learning objects regardless the definition of learning object used by the learning object creator. For example, client software using a LOM-based model would be able to retrieve learning objects from such repositories, even though the objects that fulfil the requirements of the client system were stored by a different system that was using the SCORM definition of learning object. From the different characterisations of learning object by McGreal, an ontology has been created.

If learning can be considered an Event, any term linked to the representation of learning activities, or declared to have an educational purpose in some way, may be considered a learning object for practical purposes. In Figure 2, all the classes prefixed by 'oc' are classes in the original OpenCyc knowledge base. Therefore, the term learning is represented by the class `oc_Learning`, which represents the definition of learning in OpenCyc. The abstract representation from which all the terms in McGreal's terminology have been derived is defined as the `LearningObject-Generic` class. An instance of this class may be anything used_in `oc_Learning`. The `LearningObject-AsAnything` class encompasses the broad definition of every possible meanings of learning object. The `LearningObject-AsAnythingWithEducationalPurpose` class has been defined to represent the objects for which some declaration of their pedagogical purpose exists, while the class `LearningObject-AsAnythingDigital` represents digital objects. These two latter characterisations are combined in current specifications of learning technology.

In addition, as the current learning object standards and specifications suggest, an individual of the class `LearningObject-AsAnythingDigital` should be linked to at least one `LearningMetadataRecord` by stating that the domain of the property `hasAssociatedMetadataRecord` is the class `LearningMetadataRecord`. `LearningMetadataRecord` is a generic term that can be used to derive specific terms supporting each particular specification such as `LOM_Record` or `SCORM_SCO_Manifest`. Any specification-specific learning object is, by its own nature, something with a declared educational purpose, which is at least tacit in the standardised schema. Figure 2 shows the relationships between all the terms, as described.

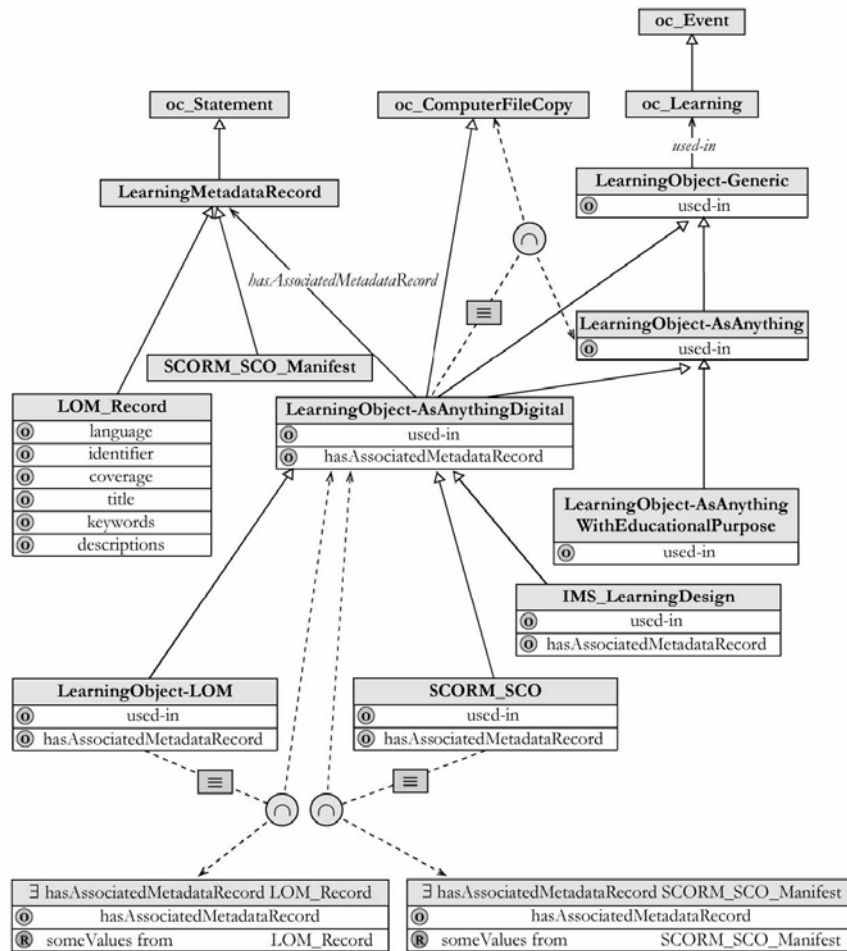


Figure 2 Ontología SLOR

In figure 2, classes in the ontology are shown as solid-outline rectangles, with several compartments separated by horizontal lines. The name of the class is held in the top compartment. The rest of the compartments hold both the properties of the class and the restrictions defined for it. Object properties are shown preceded by a circled 'o', while restrictions are preceded by a circled 'R'. For example, the class LearningObject-Generic is linked to the class oc_Learning by an arrow that represents the object property used_in in the corresponding compartment of the class. Note how the arrow is labelled with the name of the property, used_in in this case. The figure also shows the hierarchical relationship between the classes in the ontology. Classes representing more generic concepts appear in the top part of the figure. The hierarchical relationship is shown as a solid path from the more specific classes (e.g. LOM_Record) to the more generic ones (e.g. LearningMetadataRecord), with a hollow triangle at the end of the path where it meets the generic class. Other links, like those labelled with the intersection and the equivalence symbols, represent the kind of formal definition for a class as stated in the ontology. The class LearningObject-AsAnythingDigital, for example, is linked to an intersection symbol by an arrow labelled with an equivalence symbol. This is to represent the fact that the class LearningObject-AsAnythingDigital is defined as the class of objects that are instances of two classes simultaneously. In this particular example, instances of LearningObject-AsAnythingDigital accomplish the properties in both the class oc_Learning and the class LearningObjectAsAnything.

Some examples:

- LearningObject-AsAnything: books, notes, applets, web pages, a paper with schemas, questionnaires, etc.
- LearningObject-AsAnythingDigital: web pages, applets, digital papers, programs, e-books, etc.
- LearningObject-AsAnythingWithEducationalPurpose: books with educational purpose, practices, exercises, questionnaires, etc.
- LearningObject-LOM or LearningObject-SCORM_SCO: digital lessons with educational purpose – Cisco Course, Moodle Course, ATutor Course, etc.

5. A flexible repository: SLOR

All definitions explained in Section 4 can coexist in a neutral concept model. This new model provides a new set of suitable functionalities on each individual learning object conceptualisation. The flexibility acquired with this new schema provides a new scenario to store whatever type of information (normalised or no) about a learning object. If it wished (it depends on learning object creator's conceptualisation), besides it will be possible including a new specification with educative aim.

According to Figure 3, a 'LO-CREATOR' (learning object author) creates a learning object, which a final user might use after (probably) locating it by means of a repository searching capabilities. Besides, a learning object author can describe a learning object in multiple formats, for example, creating a SCORM-compliant metadata record including a table of contents, the types of the resources in the learning object, the navigation sequence and other information fields. However, it is also possible to create a LOM-compliant metadata record referred to the same LO. Therefore, SLOR can hold different metadata descriptions referred to the same object. This feature provides a new interoperability scenario where the external agents are allowed to search and interoperate with learning object metadata records of different types.

The SLOR prototype has been specifically designed for the metadata learning object creation and management, with the focus on the integration and interoperability with other systems. Compared with other repositories, the SLOR prototype provides new functionalities, most of them derived from the powerful

underlying ontology model and semantic web-oriented middleware tools, which enable the execution of reasoning and inference tasks on the metadata records it stores. The SLOR functionalities are grouped according to scalability principles:

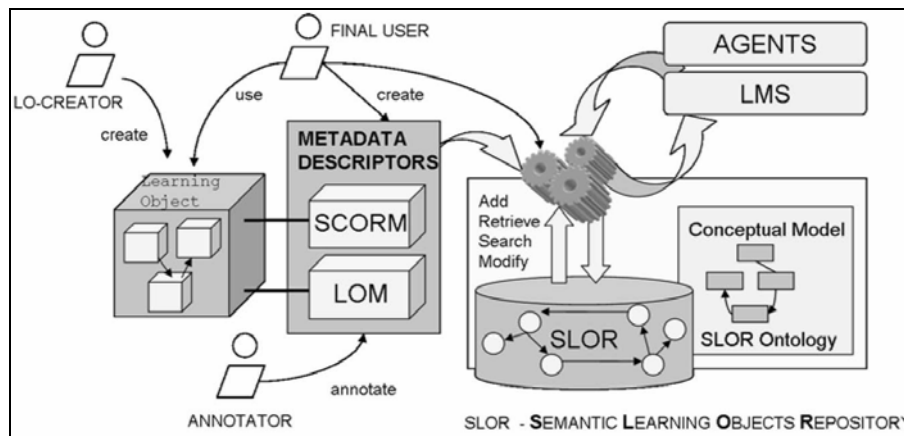


Figure 3 SLOR interactions

5.1 Metadata creation

According to the conceptual model previously established by the metadata creator, the learning object creation function allows inserting a new metadata record through the 'new RLO' interface. The entry fields in the creation form correspond to a given conceptual model of learning object (LearningObject-LOM in the example in Figure 4). However, other models – listed in the left hand panel in Figure 4 – could be used for storing metadata in the repository.

The SLOR basic model allows linking information fields to outward ontology concepts. Figure 5 shows an example of metadata edition where part of the information in the metadata category 'coverage' is set to the value 'Spain', an instance of the TGN-Nation class in the Thesaurus of Geographic Names (TGN) ontology², a structured vocabulary containing around 912,000 records, including 1.1 million names, place types, coordinates and descriptive notes, focusing on the places important for the study of art and architecture. In this thesaurus, the class TGN-Nation evokes the concept of Spain as a country. This scenario gives support to a deep level of search that allows making complex queries, e.g. retrieving elements of the baroque period that are situated in Spain.

5.2 Semantic search

This function allows searching instances of concepts in the ontology model, such as retrieving all learning objects marked as 'digital' or those learning objects that have an educational purpose. Several restrictions can be defined as part of the searching process. Restrictions allow filtering learning objects on different criteria (pedagogical, economic or other), thus providing a set of results that better accomplishes the end-user needs.

5.3 Browsing capabilities

Learning object browsing is implemented as an ontology-based seeking interface. The browser's role in this

² http://www.getty.edu/research/conducting_research/vocabularies/tgn/

model is to allow that any metadata category in LOM (excluding lifecycle and meta-metadata that are not related to the educational purposes of the objects) can be used as top guiding criteria. The ontology terms attached as descriptions are displayed in the browser (Figure 6), and it is finally the user who has the decision on their selection. The result of user selection is a query expression formed by a collection of ontology terms. Queries here are by default interpreted on a contextual basis, i.e. all the requirements selected by the user should be matched in the same context of the leaning object.

Figure 4 New metadata record

Figure 5 Linking external ontology concepts

6. Integrating a commonsense ontology

The IEEE LOM standard (LTSC, 2002) has several elements that cover different metadata information to

describe a learning object. LOM stores all knowledge about a learning object on these fields in natural language. As this feature is human-readable, it provides a good expressiveness but does not enable the reasoning capabilities over metadata records. As an example, the coverage field could link to ontology concepts such as 'Baroque' or 'Renaissance period' to insert new facts that enable reasoning capabilities and provide new room to create powerful search methods. The assertions inserted in a field must be described in a logic language.

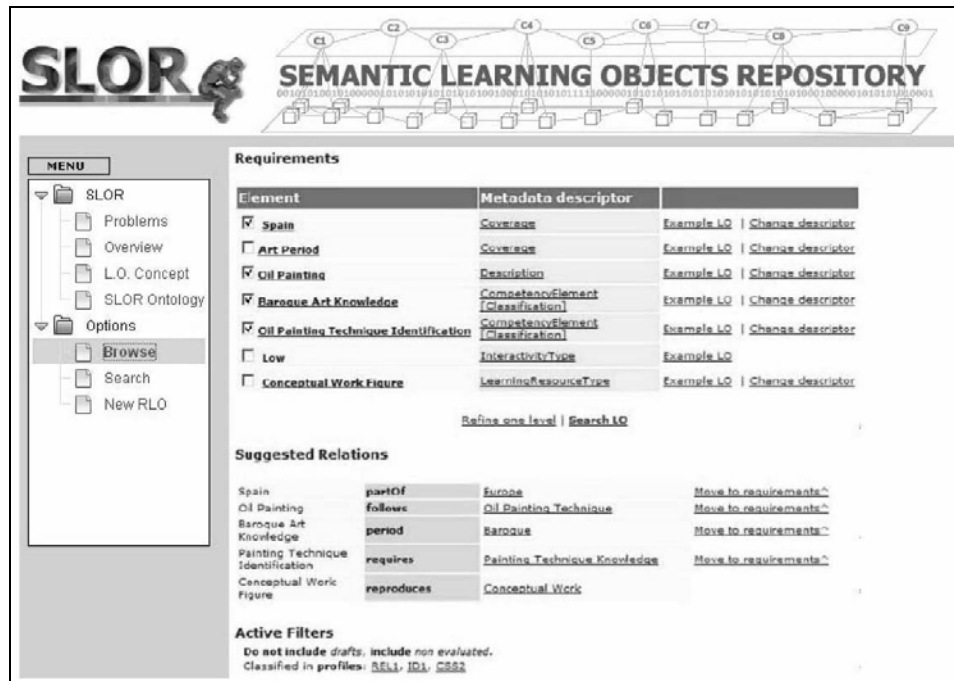


Figure 6 Searching learning objects

Marvin Minsky said: 'People have silly reasons why computers don't really think. The answer is we haven't programmed them right; they just don't have much common sense'. The key is given in the following question jointly with the Marvin Minsky assertion: How to create the common sense on computer reasoning processes? In order to answer this question and to enable the capacity for understanding and reasoning about the world as intimately as people do, Cyc Corporation has been doing a huge research effort during the past few years. Our research explores the integration of the common sense reasoning in SLOR using the OpenCyc knowledge base. In particular, we have studied two ways to integrate OpenCyc in SLOR:

- First approach (using an OpenCyc Server connected through the OpenCyc Java API). In this approach, we chose to write the SLOR ontology in CycL as a MicroTheory inside the OpenCyc Knowledge Base, because of the high performance offered by the OpenCyc engine. In this approach, the links to external concepts and the reasoning capabilities related to this were the only problems that arose. In a similar way to what Reed and Lenat (1995) explain in their study about ontology mappings to OpenCyc, we choose to link to other ontologies concepts despite the fact that this decision could involve the use of too many assertions.

- Second approach (OpenCyc knowledge base stored in an RDBMS System, and management by RDF Frameworks using the semantic technologies features). In this approach, we hypothesise that if all the information was stored in the same format (RDF) the reasoning capabilities would increase, as this would allow to link and create concepts between different ontologies without the necessity of establishing any mapping process. Following this hypothesis, we have serialised the OWL OpenCyc file into a relational database format using a semantic web framework (Section 7 explains this process in-depth). Although this approach solves the problems enunciated in the previous approach (since all knowledge is stored in the same format), the reasoning capabilities are considerably lower. The OWL language only provides a subset of the first-order logic, called description logic (Baader et al., 2003) materialised in the OWL-DL specification. However, this problem is less important since there exist new rule languages [e.g. the SWRL combining OWL and RULE-ML (Horrocks et al., 2002)] that improve the reasoning capabilities providing new inference features. To enable machine-readable descriptions within learning object metadata records, it is necessary to have the knowledge written in a logics-based language. Besides, the techniques used must keep OWL axioms. Therefore, we can conclude that to create OWL relationships between different fields inside metadata records, a sub-language is required.

7. Problems found

Following the different ways mentioned earlier, some problems concerning to SLOR implementation have been discovered:

1. Mapping OpenCyc to OWL. The OWL version of Opencyc lacks the expressiveness needed to satisfy a whole and covered semantic model. This version of OWL does not use description logics, and consequently does not define a 'domain' and a 'range' for ontology properties.
2. Closed model. OpenCyc has a closed model not allowing bindings to ontology concepts written in other languages. Despite the existence of several research projects studying the OpenCyc interaction with other thesaurus, corpus or ontologies schemas to improve the common sense reasoning, the real implementation does not exist yet, in particular that related to the interaction with semantic web schemas.
3. Low performance. The tests run on semantic web frameworks have highlighted problems when handling large amounts of data. Many of these problems are due to the inefficient inference execution over the underlying persistent model. In detail, when a query is dispatched, the reasoner planning the execution tries to retrieve from the database all the data inferred, which often implies hundreds or even thousands of unnecessary operations.
4. Interface problems. The SLOR interface needs improvement in relation to the integration of the semantic concepts and rules within metadata LOM records. In order to solve these problems we have considered new web development trends by making use of AJAX. AJAX (which stands for Asynchronous JavaScript and XML) is an innovative technique that provides a new scenario for the interaction between web client and server. This technique increases the web pages' interactivity, speed and usability. It makes the web pages more dynamic, disabling the post-back actions of the current web technologies (such as ASP.NET or Java Server Faces) and providing, at the same time, a new asynchronous communication based on XML messages between server and client. Hence, web clients do not need to send a dynamic web page and wait to receive it again. A detailed discussion about AJAX is out of the scope of these technical notes; however, interested readers are encouraged to have a look at the work of Garrett (2005) on this subject.
5. Time-consuming task. It is important to mention the research area of automated metadata creation, as one of our aims is to create a usable repository. Each learning object specification provides a considerable number of fields to be filled in order to describe its use in different contexts. To complete a whole description of a learning object is a time-consuming task, even though there exist

several tools to help create metadata records. Studies like the one by Vidal and Meire (2005) show how an automated assistant can be used to auto complete some mechanical and usual fields.

8. Conclusion

Currently, learning object repository models have several disadvantages that hamper a wide use by the community of teachers, learners and instructional designers interested in reusing the existing materials. The Semantic Model Repository (SLOR) described in this paper provides new steps towards the unification of existing learning object conceptualisations. Based on the new formal semantic web specifications, SLOR provides new room for the interchange and reuse of web-oriented educational contents in the form of learning objects. These advantages facilitate the development of new didactic tools with improved searching, integration and building capabilities. The study also remarks different problems found during the development of the prototype, in particular some problems related to the use of some semantic web tools.

References

- Baader, F., Calvanese, D., McGuinness, D., Nardi, D. and Patel-Schneider, P. (2003) 'The description logics handbook: theory, implementation, and applications', New York: Cambridge University Press.
- Baeza-Yates, R. and Ribeiro-Neto, B. (1999) *Modern Information Retrieval*. New York: Addison-Wesley.
- Berners-Lee, T., Hendler, J.Y. and Lassila, O. (2001) 'The semantic web', *Scientific American*, Vol. 284, pp.34–43.
- Garrett, J. (2005) 'Ajax: a new approach to web applications', Available at:
<http://www.adaptivepath.com/publications/essays/archives/000385.php> (last accessed on 10 July 2006).
- Horrocks, I., Patel-Schneider, F., Boley, H., Tabet, S., Grosz, B. and Dean, M. (2004) 'SWRL: a semantic web rule language combining OWL and RuleML', Member submission 21 May 2004, W3C.
- Lenat, D. (1995) 'CYC: a large-scale investment in knowledge infrastructure', *Communications of the ACM*, 38, pp.32–38.
- McGreal, R. (2004) 'Learning objects: a practical definition', *Int. J. Instructional Technology and Distance Learning*, 1.
- Reed, S.L. and Lenat, D.B. (2002) 'Mapping ontologies into cyc', Paper presented at the AAAI workshop on Ontologies and the Semantic Web, Edmonton, Canada. In proceedings.
- Soto, J. and Garcia, E. (2005) 'Failure at school – automated schedules on scholars' study', In proceedings of Sisoft 2005, 1, pp.34–51.
- Vidal, J.P. and Meire, E.M. (2005) 'Bridging the gap between learning management systems and learning object repositories: exploiting learning context information', In proceedings of Telecommunications 2005, Proceedings of ELETE 2005 – eLearning on Telecommunications Workshop, pp.478–483.
- Welty, C. and Guarino, N. (2001) 'Supporting ontological analysis of taxonomic relationships', *Data and Knowledge Engineering*, Vol. 39, pp.51–74.