

A LOM Research Agenda

Erik Duval
Dept. Computerwetenschappen, K.U.Leuven
Celestijnenlaan 200A
B-3001 Leuven, Belgium
erik.duval@cs.kuleuven.ac.be

Wayne Hodgins
Autodesk Inc. & Learnativity
258 Eucalyptus Rd.
Petaluma, CA 94952, USA
wayne.hodgins@autodesk.com,
wayne@learnativity.com

ABSTRACT

This paper presents a research agenda on Learning Objects. The main intent is to elaborate on what the authors consider important issues for research on learning objects and their use in education and training. The paper focuses somewhat on metadata related issues, but does not restrict itself to only those aspects that have a direct relationship with metadata.

Categories and Subject Descriptors

K.3 [Computing Milieux]: Computers and Education; H.3.5 [Information Storage and Retrieval]: On-Line Information Systems—Web-based services; H.4.m [Information Systems]: Miscellaneous

General Terms

learning objects

Keywords

metadata, learning, training, education, knowledge management, eLearning, library and information science, information management, content management, adaptive hypermedia, learning technology standardization

1. INTRODUCTION

In a very general sense, it has repeatedly been observed [14] that the actual impact of ICT on education and training is rather limited. As an illustration, the "Grand Challenges" conference in 2002 identified "A Teacher for every Learner: Scaleable Learner-Centered Education" as one of the grand research challenges in computer science. The panel envisioned "building the technological infrastructure to support dynamic, ad-hoc communities of lifelong learners who interact within an environment of learning objects through a creative blend of advanced computing technologies, high performance networks, authoring and collaboration tools" [29]. It was estimated that a "Manhattan project" approach, with sustained major funding over a decade or longer, would be needed to finally realize this long standing dream.

Since recent years, much of the research in this area has focused on the notion of reusable multimedia content components, referred to as "learning objects". The driving force stems from the notion that reuse of such components can lead to important savings in time and money, and enhance the quality of digital learning experiences: the end result would be faster, cheaper and better learning.

In fact, different kinds of reuse can be distinguished:

Copyright is held by the author/owner(s).
WWW2003, May 20–24, 2003, Budapest, Hungary.
ISBN 963-311-355-5.

- multiple output/distribution formats, media, outputs;
- multiple purposes (training, performance support, documentation, marketing, sales);
- multiple delivery;
- multiple "disciplines" or market segments.

A somewhat more specific term is *repurposing*, which can be thought of as the ability to use, without any (significant) changes, the same piece of content for a purpose significantly different than that it was originally intended for when created.

Early work in this area included the "Educational Object Economy" [22, 4] and Ariadne [16, 1]. Since then, and spurred by the development of the "Learning Object Metadata" standard by the IEEE Learning Technology Standards Committee (IEEE LTSC LOM), numerous initiatives have been launched in academic and corporate contexts - though these two worlds remain more isolated from each other than the authors consider necessary or healthy [21].

This paper tries to identify the main research issues in this area, grouping related issues together, and indicating the main barriers for widespread learning object reuse at this moment. As such, this proposal is part of a long tradition of papers such as [17] and [25].

In the following sections, we first deal with the notion of Learning Objects (LO's) and Learning Object Metadata (or LOM). In section 4, we deal with authoring aspects of LO's and LOM. Section 5 deals with access to LO's. In order to enable widespread re-use, interoperability issues are extremely important. They are dealt with in section 6. Sections 7.1 and 7.2 deal with issues that have a high impact on the actual adoption in practice of a LO based approach: the development of powerful tools and appropriate business models.

2. BEYOND DOCUMENTS

2.1 Introduction

There is a widespread tendency to equate LO's with "documents", typically represented as a file or a set of files. The authors believe that it is quite appropriate to consider documents and files as a form of output and delivery of LO's. However, it is very restraining to *only* think of LO's in this way. In fact, the more flexible and advanced applications of LO's go well beyond the simple document paradigm.

A simple kind of extension is to think of more sophisticated LO's than simple, static documents. A simulation, for instance, allows for dynamic adaptation to user interaction. The early work on the Educational Object Economy focused on the use of javabeans as reusable components that could interact with each other (see also

section 6.2). This goes well beyond current practice, where the unit of reuse is typically a complete simulation, rather than a component thereof. The use of javabeans or other component technologies aligns the concept of LO's much more with that of the notion of "objects" in the object oriented programming paradigm.

2.2 Research Issue 1 - A Learning Object Taxonomy

According to the Learning Object Metadata (LOM) standard, a learning object is 'any entity, digital or non-digital, that may be used for learning, education or training' [9]. This definition allows for an extremely wide variety of granularities. This means that a learning object could be a picture of the Mona Lisa, a document on the Mona Lisa (that includes the picture), a course module on da Vinci, a complete course on art history, or even a 4 year master curriculum on western culture.

In one sense, this is appropriate, as there are a number of common themes to content components of all sizes. In another sense though, this vagueness is problematic, as it is clear that authoring, deployment and repurposing are affected by the granularity of the learning object.

In order to address this problem, a learning object taxonomy or a set of such taxonomies should be developed to identify different kinds of learning objects and their components. We have developed a first starting point for such a taxonomy. It is important to note that this taxonomy applies to multiple applications. The first two levels are application domain independent, and can for instance also be deployed in the field of technical documentation. Only the third and fourth level are specific to the field of learning.

1. *Raw Media Elements* are the smallest level in this model: these elements reside at a pure data level. Examples include a single sentence or paragraph, illustration, animation, etc. A further specialization of this level (or complementary taxonomy?) will need to take into account the different characteristics of time-based media (audio, video, animation) and static media (photo, text, etc.).
2. *Information Objects* are sets of raw media elements. Such objects could be based on the "information block" model developed by Horn [18]. While Horn's model refers to text and illustrations (as it is based on pioneering work in the mid 1960's!), the plan is to generalize the concepts to deal with more advanced and innovative content.
3. Based on a single objective, information objects are then selected and assembled into the third level of Application Specific Objects. At this level reside *learning objects* in a more restricted sense than the aforementioned definition of the LOM standard suggests.
4. The fourth level refers to *Aggregate Assemblies* that deal with larger (terminal) objectives. This level corresponds with lessons or chapters, which can in turn be assembled into larger collections, like courses and whole curricula.

Clearly, information objects contain raw media elements. Learning objects contain information objects. Aggregate assemblies contain learning objects and other aggregate assemblies.

The smaller level of granularity in this taxonomy is essential, as we believe that repurposing can only be accommodated by explicitly identifying the information objects and the raw media elements they contain.

2.3 Research Issue 2: Learning Object Component Architecture

In order to realize the full potential of dynamic composition of learning object components, it is necessary to develop a flexible architecture that enables:

- *structuring* of learning objects and their components: An important general principle in hypermedia systems is the separation of content, structure and presentation. Structural aspects of aggregate learning objects can be based on the IMS Content Packaging specification [5], the Educational Modeling Language (EML) [19], the Synchronized Multimedia Integration Language (SMIL) [8] and DocBook [27].

Innovative approaches to learning object structuring should include dynamically generated components, based on some processing of (semi-)structured data. A simple example is a learning object on weather forecasting that makes use of a set of current satellite images relevant to the location of the learner, accessing simple data about the learner and the time that the learning is to take place to retrieve a relevant set of images, so as to realise an authentic and thus an effective and efficient experience. This kind of context points to the need for structures that identify relevant components through search criteria, rather than just listing the identities of the components.

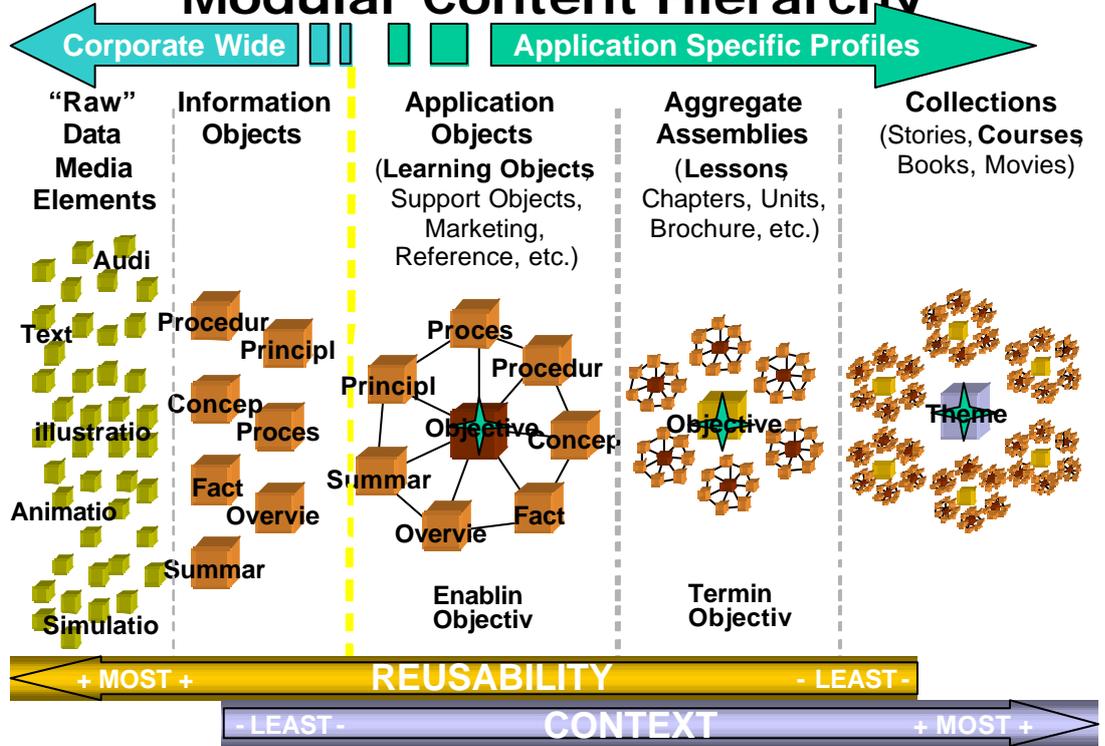
A more complex example would access information about the role of the learner in the organisation, his personal goals and those of the organisation, his agenda, etc. to generate a highly customized learning object that would be relevant for the task at hand, and that would take into account the constraints (time, language, cost, location, etc.) that influence the particular context in which the learning is to take place. This approach to personalized learning is closely related with the field of adaptive (educational) hypermedia, intelligent tutoring systems, etc. [11]. There are some successes in relatively constrained domains [20], but it is clear that much research is needed before generally applicable, large-scale implementations can be developed.

- *interactions* between learning objects and their components: The original EOE approach focused on JavaBeans, where such interactions were based on a common API [22]. The use of javabeans or other component technologies aligns the concept of learning objects much more with the equivalent notion of "objects" in the object oriented programming paradigm.

Consequently, this kind of behaviour can be thought of as a specific application of general software engineering component based approaches. In this specific context, there are many parallels between reuse of software artifacts in general, and learning object reuse in particular. Even though this subject has a long research history in software engineering, large-scale applications in practice remain relatively few. Moreover, they are often constrained to particular domains or to particular technological approaches.

Moreover, not only has that approach not been adopted so widely in the context of learning objects, deeper interoperability also requires further considerations on interactions between for instance the models that underly the behaviour of learning objects. An example that illustrates this requirement is that of learning objects that simulate different parts of the human body, and that need to exchange data between them [14].

Modular Content Hierarchy



3. LEARNING OBJECT METADATA

3.1 Introduction

In the Learning Object Metadata standard, a hierarchical structure is defined of 9 categories. Each of the categories groups related data elements that cover specific aspects, such as technical or educational characteristics [9]. It is important to note that, in LOM, all data elements are optional and that the LOM structure can be extended (see also section 3.3).

The two main research issues we deal with in this section relate to the need for empirical analysis of actual metadata usage (now that LOM begins to be deployed on a large scale), and to the need for guidance on the development and use of application profiles that adapt LOM to the needs of a specific community.

3.2 Research Issue 3: Empirical Analysis

There is a high degree of subjectivity in many of the LOM data elements. This is sometimes perceived as a problem, but it is the opinion of the authors that

- there is no need to avoid subjectivity in the descriptions of LO's;
- in fact, the subjective metadata are often among the more interesting ones.

In other words, the subjectivity of many metadata elements "is a feature, not a bug". The subjective metadata are often the most valuable for increasing the effectiveness of searches and identifying the most relevant content for a given individual or situation. As a simple illustration, it may be more relevant to know that a particular person recommends a particular LO (highly subjective metadata) than to know the exact size of that LO in kiloBytes (highly objective metadata).

However, organisations often provide good practice documents, so as to give guidance to their community of users. Typical examples include CANCORE [2] and SINGCORE [7]. As such documents start to be developed and applied, analysis is needed of the similarities and differences between such guidelines, and their rationale (in reference to the requirements of the intended communities, or otherwise).

Moreover, it would be extremely useful to collect and analyse empirical data about the actual use that is made of metadata, by different classes of end users, including:

- indexers that generate metadata, describing LO's;
- searchers for relevant LO's, who typically define search criteria over metadata elements.

Such analysis could further our understanding of actual use cases. Useful approaches include log analysis of LO Repositories (see also section 6.4), usability studies of LO(M) tools (see also section 7.1), analysis of the actual content of LO repositories, such as the kind of LO's, the actual metadata, the actual amount of reuse, the actual annotations by users who provide feedback on their reuse of LO's, etc. A rare example of a simple initial analysis along these lines can be found in [15].

Another approach would be to analyse the differences and similarities between metadata authored by independent indexers. In principle, as there can be more than one LOM instance for a LO, this kind of analysis can be carried out a posteriori. In practice, it seems like there are not many LO's with multiple LOM descriptions. Consequently, it seems that, at this moment, this approach will have to be based on specifically set up experiments. Such experiments are rather similar to experiments on peer reviews of for instance paper submissions [26].

3.3 Research Issue 4: Application Profiles

It was noted in section 3.1 that LOM is extensible. In essence, new data elements can be introduced anywhere in the LOM hierarchy. An alternative extension mechanism is the classification category for listing of sourced taxonomic stairways in arbitrary classifications.

An alternative mechanism to adapt LOM is that of "application profiles" that enable increased semantic interoperability in one community, in a way that preserves full compatibility with the larger LOM context. The fundamental techniques for the definition of application profiles include:

- *giving elements a mandatory status*: There is in fact a rather large degree in the selection of mandatory elements in different application profiles, such as ARIADNE [1], CANCORE [2], SCORM [6], and others.
- *restricting the value space of data elements*: The LOM standard is intended to be broadly applicable. This means that some of the available options may not be relevant in specific contexts. A typical example is the data element 2.3.1:LifeCycle.Contribute.Role that includes values like "educational validator". In many contexts, that role can be omitted. (One could argue that such irrelevant values can just as well be retained in the application profile. In fact, because many of the tools remain so close to the specification they support - see also section 7.1 - these irrelevant values are often exposed to end users, and thus compromise the usability of LOM.)
- *imposing relationships between elements*: In the context of particular communities, it may make sense to impose additional relationships between LOM data elements. As an example, the ARIADNE application profile mandates that, if 6.1:Rights.Cost equals "yes", then the set of data elements under 2.3:LifeCycle.Contribute with 2.3.1:LifeCycle.Contribute.Role equal to "publisher" or "author" becomes mandatory. Clearly, the idea is that, if payments are involved, the rights holder (the author or publisher in the ARIADNE context) should be known.
- *not including some LOM elements*: As all data elements are optional, a community may decide not to make use of some elements. Alternatively, the value of a particular data element may be fixed for a particular community: as an example, the educational context will always be that of a university for a community that only deals with that context. In that case, the data element need not be presented to the end user for whom the context of use is thus restricted.
- *identifying taxonomies and classifications*: LOM includes a category that supports the use of arbitrary taxonomic stairways, in general classification structures, for a purpose that can be explicitly defined. In the ARIADNE application profile, for instance, a dynamically updated classification structure for the structuring of human knowledge is references [1].

Developers can generate considerable added value for their communities through application profiles: the result can be a more easily applicable version of LOM, with higher semantic interoperability within a community, without compromising overall exchange of LOM instance with those who do not belong to the community.

However, there are important questions about the proper ways to define application profiles, about how to translate from one such profile to another one (through a LOM common denominator, or directly?). Also, it often remains unclear what kind of community

an application profile tries to serve, and how the characteristics and requirements of that community have influenced the definition of the profile. Clear guidelines on how to proceed in this process of application profile definition would be most useful and relevant.

4. LO(M) AUTHORING

4.1 Introduction

A number of issues need to be better understood if large scale LO (re-)use is to become a reality. In this section, we focus on aggregation (section 4.2) and decomposition (section 4.4) in LO authoring, and on the notion of design for reuse (section 4.3). We also look at the related issue of LOM authoring (section 4.5).

4.2 Research Issue 5 - Authoring By Aggregation

Traditionally, authoring tools mainly support the process of authoring from three points of departure:

- a blank document that needs to be "filled" with content, where the structure of the LO is defined during the elaboration of that content;
- a template that needs to be instantiated, where the structure of the LO is defined a priori;
- an existing LO that is edited and modified in the process of authoring, and then typically saved as a new LO.

The main idea in our view however is that LO's are created by selecting content/information objects from a repository, usually with the significant assistance of metadata and profiles to do so. These LO's can then be assembled into a new LO. This can be referred to as authoring-by-aggregation.

The new LO, as it provides new context for the components, may need to provide "glue" that takes the learner from one component to another. A simple example of this kind of facility is the way that presentation authoring tools (like Microsoft Powerpoint, or SliTeX, etc.) allow for existing slides to be included in new presentations and then add automatically "next" and "previous" transitions between those slides. More sophisticated "glue" would enable the author of the aggregated content to include transitional material (like "In this chapter, we will illustrate the concept of relativity that was introduced in chapter V"), so as to give guidance to the learner on how the components fit together in the aggregate.

This kind of "glue" is dealt with by "sequencing" specifications, that enable the definition of learning paths. These learning paths are themselves discrete components or objects and as such can be stored separately, modified independent of the content, reused AND of course also have their own associated metadata to aid with discovery, search and retrieval.

This can also be illustrated by reference to the architecture model presented in section 2.2: the actual content is always authored at the object or component level and from there everything is an "assembly" of these components. In this context, creating a LO is an "assembly" process which includes initially the design or planning of the assembly (blueprint, learning path, etc.) and then the creation of the actual assembly itself, formatting, delivery, etc.

4.3 Research Issue 6 - Design for Content Reuse

We mentioned in section 4.2 that content authors should avoid explicit references to other components, as these other components may not be available in the context of reuse. Continuing the example from that section, a reference to "chapter V" makes little sense

if the context in which a component is reused doesn't include a "chapter V". We already mentioned that the aggregate LO should provide this kind of "glue" information. This is one example of an issue that needs to be considered when content is designed to be reusable.

There are many more issues that need to be considered when "designing for reuse":

- It should be easy to modify *mathematical symbols*: the symbols used often depend on the context, and they should be consistent within a given context. So, if one content component refers to an angle in a diagram as α , and another has a visualization that refers to the same angle as θ , then the learner will have significant cognitive challenges to map α to θ . Whereas an expert in the domain may be expected to be able to do this kind of mapping without problems, the same cannot be expected of a learner. That is why content should be designed in such a way that it becomes easy to transform them from one symbolism to another.
- A related issue is that of *textual labels* in visual material: it should be simple to replace such labels with alternatives, for instance in a different language, or using an alternative vocabulary. This is a problem that has been studied in the context of multi-lingual user interfaces. There are some simple approaches that basically rely on separation of the labels from graphical material, but it remains quite difficult to deal with some of the more subtle issues, such as the differences in lengths of textual labels or the differences in reading direction (left-to-right versus right-to-left) in different languages. It should be noted that such labels can themselves be considered metadata, that get "added" to the actual content when the latter is visualized in a particular context.
- Generalizing further, methods need to be developed for adapting the *look and feel* of content. When different components are aggregated together, the result should not look like a collection of components from different origins. One could think of aggregation tools that allow the author to apply a "design template" to impose a specific look and feel on the resulting aggregate.

In a general sense, then, designing content to be reusable is largely a discipline of not including any more context than absolutely necessary in the content itself and instead adding in context by means of the design of the LO, the overall "assembly" and the Learning Path. Or, context can also be provided through the use of very context specific content such as examples or exercises which could reflect a very specific set of data, application etc.

The above illustrates once more why we need to have a content model such as that of section 2.2, where the reusable components are below the LO level and are these "information objects" or "content objects" which are "pure" or "raw" data in that they have little to no context, no formatting and no "style". This "style" and context is added by combinations of the context of the design, the learning paths and/or the presentation layers via typical style sheets.

Basic technologies that deal with some of the issues mentioned above do exist. In fact, these issues are not that specific to the domain of education and training, and approaches that support the separation of content, presentation and navigation in general can be applied to this context as well. An important problem is that current authoring tools often encourage users to produce content that mixes these aspects in a way that makes it difficult to separate the different layers of information afterwards.

4.4 Research Issue 7 - LO Decomposition

It is important to realise that there is a trade-off between reusability and added value of LO's in terms of granularity: smaller LO's are more easily reusable, yet the added value one derives from their reuse is lower than that of larger LO's that tend to be less easily reusable. That is why we need to accommodate both larger (because adding more value) and smaller (because more easily reusable) LO's.

However, authors are typically not so comfortable with authoring really small or large LO's. The exception for the smaller LO's is that of photographers that would typically focus on authoring single pictures. In a more general sense, many authors find it difficult to produce small content units with a well-defined, restricted scope. There are some examples of tools that allow authors to work on a level of granularity that they are comfortable with, and then decompose that object into components semi-automatically. In the case of video for instance, scene cuts can be detected automatically, to suggest appropriate boundaries to separate components. Text segmentation tools support the transformation of text documents into pedagogical hypertext [28]. Another example would take an HTML file with embedded images and extract from that distinct components for the original HTML file and the images on the one hand, as well as for all of the components on the other hand.

The main point is that we need better support for authoring of LO's by aggregation (see section 4.2) and for automatic decomposition of LO's to extract the components of a LO that was originally produced as an aggregate.

4.5 Research Issue 8 - LOM Generation

One of the most often heard critical remarks about LOM is that content authors are not willing to spend the extra amount of effort to add metadata to their LO. While this may be less true for LO's that involve considerable investment, as the added effort to add metadata then becomes almost negligible in the larger context, it does present a serious problem in more "artisanal" settings.

We believe that this problem can be overcome by techniques for (semi-)automatic generation of metadata. Several sources of metadata can be harvested with this intent:

- The *LO itself* can be processed for metadata. For instance, it is a relatively simple process to extract from an HTML LO the title, the language used, references to other LO's, the name of the author (often included in the metadata that are inserted by the authoring tool), etc. Moreover, search engines like Google illustrate that existing harvesting techniques can be quite powerful. More research on mapping the results of these techniques to the different metadata elements in structures such as LOM is needed.
- *Related LO's* cannot only be described as such in the metadata of the LO to be described (using the relation category in LOM). Further research is required on how metadata can propagate from one LO to a related one. It seems obvious, for instance, that the language of a textual component will most probably be the language of the composite to which it belongs. Likewise, if there is metadata available about the cost involved with component LO's, then it seems safe to suggest that the cost involved with the composite is at least the aggregated cost of all the components involved. Also, existing techniques for document clustering and text mining need to be reconsidered in the context of explicit metadata, so that for instance the metadata of related and already described LO's can be used to generate metadata for new LO's.

- The *author(s)* of the LO is often the source of additional metadata for the LO, as authors mostly produce LO's in one language only, or in one content domain, or for one kind of audience (say university level), etc. Even for authors that produce more diverse LO's, the range of relevant values for many elements can be reduced significantly when characteristics of the author are taken into account.
- The *context* of LO authoring can also significantly reduce the number of relevant options. If the metadata authoring tool is launched within the context of a course for instance, then it seems reasonable to suggest metadata of the course as starting values for the metadata of the LO.

In addition, *templates* of reusable metadata can be created, where many of the relevant fields can be prefilled. Often, instantiating the template will involve little more than simple selection between a small number of relevant values for a few remaining fields.

It is important to note that the "psychological" effect of presenting an automatically generated metadata instance and asking the end user to verify that this description is correct. This is a much less intimidating proposal than being presented with an empty form that includes a large number of empty text boxes to be filled in, as well as many long lists of values to be selected from.

5. ACCESS

5.1 Introduction

Assuming that a critical mass of LO's and associated LOM instances becomes available, there is an obvious need to improve access to the LO's. Most of the tools at this moment are based on an electronic form that enables end users to compose boolean combinations of search criteria (but we know from digital library research that this approach has serious usability problems) or on a simple text box (that fails to take advantage of the structured nature of LOM). What is needed is more research on how to provide flexible, efficient and effective access to a large repository of LO's, so that end users can quickly zoom in on those LO's that are relevant to them.

The responsibility for this kind of selection can reside in different places:

- the *teacher, trainer, professor*, etc. may select the LO's. This selection typically takes place a priori, i.e. in advance of the actual learning experience;
- the *learner* may select LO's, typically a priori in the case of larger LO's (courses) or at the time that the actual learning takes place, in the case of smaller LO's (a document in a module for instance);
- the learning management *system* may process rules that match characteristics of the learner (style, intended learning outcome, already acquired knowledge and skills) against the LOM descriptions, in order to select the most appropriate LO's.

5.2 Research Issue 9: Novel access paradigms

There is an urgent need for more research on novel access paradigms that enable an end user to zoom in on relevant LO's without requiring him or her to go through a lengthy process to formulate complex search criteria, to evaluate some of the results, refine the search criteria, etc. Two such paradigms are briefly mentioned below.

- *Information visualisation* enables end users to manipulate controls over the metadata, zoom in on potentially more relevant LO's and continuously keep an overview of how additional search criteria restrict the remaining number of LO's [13].

It is important to notice that this approach tries to empower end users and keep them in control of the selection process, rather than trying to transfer the locus of control to the Learning Management System.

- A somewhat "softer" approach to LO selection relies on *social recommending* techniques where similarities between end users are exploited to suggest potentially relevant LO's [30].

This approach has a number of interesting features:

- This kind of suggestion exploits data about previous interactions between users and LORs, and avoids the cognitive burden and time involved in browsing or searching the LOR.
- It may suggest LO's that one would never find through explicit searches, and thus open the view of the end user on the LOR.

In both of these approaches, LO's are identified in a more natural, less laborious way than by "filling in electronic forms".

5.3 Research Issue 10: Heterogeneous Searches

A specific and significant area of research is concerned with searches across heterogeneous repositories of content, metadata, learning paths and presentation or style information.

There is a really critical and so far apparently unnoticed need that this new component based content reference model requires: to be able to search for content and for metadata across any number of repositories which are "federated" in the sense that they can be treated as a virtual single repository. Note that there is a requirement for standardization efforts in this area to make this work.

6. INTEROPERABILITY

6.1 Introduction

Interoperability can be defined as "enabling information that originates in one context to be used in another in ways that are as highly automated as possible" [23]. More specifically, in our context, interoperability can be defined as the ability for objects from multiple and unknown or unplanned sources, to work or operate technically when put together with other objects. Examples include:

- content objects from different original creation/authoring tools working together when assembled together into a LO.
- LO's and content objects being able to work properly when moved from one infrastructure (operating system, LMS, etc.) to the next.

Typically, this requires standardization of common protocols, formats, etc. This kind of interoperability is a condition to realise the vision of an open, large-scale LO infrastructure with sufficient critical mass.

Interoperability is required at different scales:

- between LO's (section 6.2);
- between LO's and Learning Management Systems (section 6.3);

- between LOR's (section 6.4);
- between metadata schemas, understood in the remainder of this paper as the collection of interrelated data elements that structure a metadata instance (sometimes referred to as "vocabularies" of data elements (section 6.5);
- more general notion of semantic interoperability, that would enable crossing, among others, the boundaries of culture in general and language in specific.

6.2 Research Issue 11 - LO Interoperability

An important, and currently somewhat neglected, kind of interoperability focuses on interactions *between LO's*. The original EOE approach focused on a JavaBeans like approach, where interoperability between LO's was based on a common API [22].

Not only has that approach not been adopted so widely in the context of LO's, deeper interoperability also requires further considerations on interactions between for instance the models that underly the behaviour of LO's. An example that illustrates this requirement is that of LO's that simulate different parts of the human body exchange data between them [14].

This kind of behaviour can be thought of as a specific application of general software engineering component based approaches. In this specific context, there are many parallels between reuse of software artefacts in general, and LO reuse in particular. Even though this subject has a long research history in software engineering, large-scale applications in practice remain relatively few. Moreover, they are often constrained to particular domains or to particular technological approaches.

6.3 Research Issue 12 - LMS Interoperability

A kind of interoperability that has been worked on much more is that between a LO and a Learning Management System (LMS). In fact, among the earlier work on industry standards in the domain of learning was the Aviation Industry CBT Committee (AICC) Computer Managed Instruction (CMI) specification that deals with precisely this kind of interoperability. Additional work is needed though, as some of the original assumptions underlying that early work no longer hold in the current world of distributed Web applications.

6.4 Research Issue 13 - LOR Interoperability

To date, little work has been done on interoperability between LO repositories. The focus of this work should be on common *protocols*, so that

- clients can interact with a set of servers, rather than being developed for interaction with one specific server;
- servers can interact with one another, to exchange metadata and LO's (as well as potentially user profiles, and other kinds of "meta-metadata").

Again, there is a lot of more general work in this area that can be reused and applied to the specific context of LOR's: this includes for instance Z39.50, and, most importantly, the Open Archives Initiative (OAI): this initiative develops standards that aim to facilitate the efficient dissemination of content. The metadata it relies on are based on the Dublin Core Metadata Element Set. The core of OAI is a protocol for metadata harvesting, that allows crawlers to collect metadata from a variety of sources

An alternative, more novel approach relies on *peer-to-peer* technologies: among the main purposes is the desire to lower the barriers to participate in the LO infrastructure. Initiatives in this area include LOMster, Edutella, Pool, etc. [24].

A very interesting way to deal with these kinds of issues is to develop hybrid architectures where wrappers around servers make them act like super-peers, gateways peers interact with several kinds of peers from different networks, etc.

6.5 Research Issue 14 - Schema Interoperability

As metadata infrastructures develop, the need increases to enable automatic mapping between metadata schemas. The technology of metadata registries seems to offer a promising approach to tackle this issue: the basic idea is that descriptions of the interrelationships between metadata elements are made available for machine processing. As an example, one can imagine that a registry would include information about the definitions of data elements in the Dublin Core Metadata Element Set (DCMES) and the Learning Object Metadata (LOM) Base Schema. It could also include references from the DCMES data elements to the corresponding LOM data elements. In this way, software agents can "automatically" transform queries on DCMES data elements to queries on LOM data elements, or transform LOM instances to DCMES instances, etc.

Even though this approach is quite promising, it does seem to be in its very early stages. For instance, the design principles for metadata registries are not well understood, and it seems likely that a diverse set of registry services for metadata processing will evolve [10]. Also, relatively few operational metadata registries have been developed, and most of them are used by software agents that have been developed for that one registry only. Some early standardization work has taken place, most notably in the context of ISO 11179, which defines data elements that describe metadata schemas and their data elements. However, at the moment, there is for instance no standardized way to interact with registries, or to find out what their scope is, etc.

Metadata registries could enable interoperability between

- metadata standards, so that software agents can transform between LOM, MPEG-7, DCMES, and other metadata standards;
- application profiles, that can include data elements from several data element sets (the EUN application profile includes DCMES and LOM data elements, as well as some that have been defined by the EUN itself).

In fact, the very notion of registries can be extended further, to enable not only mappings and crosswalks between data elements, but also between the values that are appropriate for these data elements. This can be thought of as a vocabulary registry. Early work in this area is currently taking place in the CEN/CENELEC ISSS Workshop on Learning Technologies [3]. The CEN data registry enables one to discover for instance what kind of value spaces are being used for particular data elements, and by whom. A typical example query will reveal that communities use the Library of Congress Subject Heading, the Arts and Architecture Thesaurus, the Medical Subject Headings, etc. for the Subject element in DCMES or for the classification category in LOM. The main difference between this kind of registry and the more conventional one is that the number of values for a data element can be substantially larger than the number of data elements in a schema. Also, the interrelationships between values can be quite complex, in taxonomies, classification structures or ontologies.

7. TOOLS AND BUSINESS MODELS

7.1 Research Issue 15 - Tools

It may seem odd at first sight to include tools development as a "research issue", but it is important to emphasize that many of the standards and specifications under development are not meant as such for end users, but rather to enable tool developers to design and implement tools that provide advanced learning functionality to the end users. This means that we need to develop these tools in order to validate the standards against implementations, and, more importantly, against user requirements in realistic experiments.

The problem is that the usability of the few tools that have been developed is often rather poor... In many current tools, the LOM standard "shines through" the user interface: this is certainly a sign of the relatively immature status of these tools - not so much in terms of functionality and stability, but much more in terms of usability and effectiveness.

An obvious example is the classification category in LOM: it enables the inclusion of taxonomic stairways of arbitrary classification structures in a LOM instance, with a reference to the source of the classification, and an explicit indication of the purpose of this classification. This is an extremely versatile and flexible feature, but not one that should be shown to end users as such: rather, we believe that the tool developers should decide upon relevant classifications for their community, and should build into the tool facilities for the actual selection of the taxonomic stairways in the classification structures - a feature hard enough to implement in a usable way as such!

A number of related application domains with a longer R&D history may provide guidance on how to develop appropriate tools: the example of reuse in software engineering was already mentioned before. Some of the more professions Integrated Development Environments include facilities for reuse of object-oriented classes. Another application domain that can provide inspiration is that of technical documentation, where reduced product cycles and increased personalization of products have triggered the adoption of novel approaches based on structured data and Content Management Systems.

7.2 Research Issue 16 - Business Models

It is clear that LO production and (re-)use will only become widespread once business models of some kind will be developed. It seems like there will be room for a wide variety of such models, ranging from "share and reuse" for free approaches, as pioneered by ARIADNE and also advocated by the MIT Open CourseWare initiative, to more commercially oriented approaches.

In the academic context, a related question is that of appreciation of learning object authoring and (re-)use. One suggestion is to treat authoring of learning objects in a way that is similar to that of papers in journals and conferences. Quality could be taken into account by treating re-use of learning objects in a way that is similar to the evaluation of citation of scholarly work - similar warnings about inappropriate use of bibliometrical data would also apply.

Another related question is that of the relationship between the LO approach and Knowledge Management in general, the more so as the boundary between "learning" objects and "information objects" in general is quite fuzzy anyway in a world where "just-in-time" learning deals with small granularities that are anchored in the day-to-day context of the learner. Or, put another way, what kind of functionality does a Learning Content Management System (LCMS) need to provide that a CMS in general does not?

8. CONCLUSIONS

The intent of this paper was to formulate a number of research issues that the authors consider important to further the LO approach

to learning. We hope that this paper may be instrumental in raising the effectiveness and efficiency of research in this area, by making the different options to focus on more apparent, or, at least, by triggering a debate on what the most relevant research issues for the field are at this moment.

9. ACKNOWLEDGMENTS

We are deeply indebted to the many wonderful experiences and deep insights that we gained from working with many communities, including the ARIADNE Foundation and the IEEE LTSC LOM Working Group.

10. REFERENCES

- [1] Ariadne. <http://www.ariadne-eu.org>.
- [2] Cancore. <http://www.cancore.ca>.
- [3] Cen/cenelec iss learning technologies workshop. <http://www.cenorm.be/iss/Workshop/lt/Default.htm>.
- [4] Educational object economy. <http://www.eoe.org>.
- [5] IMS. <http://www.imsproject.org>.
- [6] Scorm. <http://www.adlnet.org>.
- [7] Singcore. <http://www.ecc.org.sg/eLearn/MetaData/SingCORE/index.jsp>.
- [8] Synchronized multimedia integration language (SMIL 2.0), Aug. 2001. W3C Recommendation, <http://www.w3.org/TR/smil20/>.
- [9] 1484.12.1 IEEE Standard for Learning Object Metadata. June 2002. <http://ltsc.ieee.org/wg12>.
- [10] Principles of metadata registries. 2002.
- [11] P. D. Bra, P. Brusilovsky, and R. Conejo, editors. *AH2002, Adaptive Hypermedia and Adaptive Web-Based Systems, Second International Conference*, May 2002.
- [12] S. K. Card, J. D. Mackinlay, and B. Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Jan 1999. <http://www.cs.umd.edu/hcil/pubs/books/readings-info-vis.shtml>.
- [13] A. V. Dam. Next-generation educational software. In *EdMedia2002: World Conference on Educational Multimedia, Hypermedia & Telecommunications*, June 2002. http://www.cs.brown.edu/research/graphics/EdMediaKeynote_files/frame.htm.
- [14] E. Duval, E. Forte, K. Cardinaels, B. Verhoeven, R. V. Durm, K. Hendriks, M. W. Forte, N. Ebel, M. Macowicz, K. Warkentyne, and F. Haenni. The ariadne knowledge pool system. *Communications of the ACM*, 44(5):72–78, 2001. <http://doi.acm.org/10.1145/374308.374346>.
- [15] E. Forte, M. Wentland-Forte, and E. Duval. The ariadne project (part 1): Knowledge pools for computer-based and telematics-supported classical, open, and distance education. *European Journal of Engineering Education*, 22(1):61–74, 1997.
- [16] F. G. Halasz. Reflections on Notecards: Seven issues for the next generation of hypermedia systems. *Communications of the ACM*, 31(7):836–852, July 1988.
- [17] R. E. Horn. Structured writing as a paradigm. In A. Romiszowski and C. Dills, editors, *Instructional Development: State of the Art*. Englewood Cliffs, N. J., 1998. <http://www.stanford.edu/~rhorn/HornStWrAsParadigm.html>.
- [18] R. Koper. Modeling units of study from a pedagogical perspective - the pedagogical meta-model behind EML, June 2001. <http://eml.ou.nl/introduction/docs/ped-metamodel.pdf>.
- [19] P. Maddocks. Case study: Cisco systems ventures into the land of reusability. *Learning Circuits*, March 2002. <http://www.learningcircuits.org/2002/mar2002/maddocks.html>.
- [20] F. Neven and E. Duval. Reusable learning objects: a survey of lom-based repositories. In *Proceedings of ACM Multimedia*. ACM, Dec 2002. Accepted <http://mm02.eurecom.fr/>.
- [21] J. Roschelle, D. DiGiano, M. Koutlis, A. Repenning, J. Phillips, N. Jackiw, and D. Suthers. Developing educational software components. *IEEE Computer*, 32(9):50–58, September 1999.
- [22] G. Rust and M. Bide. The `indexs` metadata framework: Principles, model and data dictionary, June 2000. <http://www.indexs.org/pdf/framework.pdf>.
- [23] S. Ternier, E. Duval, and P. Vandepitte. Lomster: Peer-to-peer learning object metadata. In P. Barker and S. Rebelsky, editors, *Proceedings of Ed-Media: World Conference on Educational Multimedia, Hypermedia & Telecommunications*, pages 1942–1943. AACE, AACE, June 2002.
- [24] J. van Ossenbruggen, L. Hardman, and L. Rutledge. Hypermedia and the semantic web: A research agenda. *Journal of Digital Information*, 3(1), 2002.
- [25] H. M. Walker, W. Ma, and D. Mboya. Variability of referees' ratings of conference papers. In *Proceedings of the 7th annual conference on Innovation and technology in computer science education*, pages 178–182. ACM Press, 2002.
- [26] N. Walsh and L. Mueller. *DocBook: The Definitive Guide*. Oct. 1999. <http://www.oreilly.com/catalog/docbook/>.
- [27] S. H. . M. Wentland. Ophelia, object-oriented pedagogical hypertext editor for learning, instruction and authoring. In *Proceedings of Hypermedia et Apprentissage*, oct 1998.
- [28] B. York, A. van Dam, J. Ullman, E. Soloway, J. Pollack, A. Kay, and T. Kalil. A teacher for every learner. June 2002. <http://www.cra.org/Activities/grand.challenges/slides/education.pdf>.
- [29] G. Zacharia, A. Moukas, and P. Maes. Collaborative reputation mechanisms in electronic marketplaces. In *HICSS*, 1999.