Learning Objects and Instruction Components

Clark Quinn

KnowledgePlanet.com

Introduction

A new concept in Educational Technology is the 'learning object'. Learning objects, as defined by the IEEE's Learning Technology Standards Committee (http://ltsc.ieee.org), are "any entity, digital or non-digital, which can be used, re-used or referenced during technology supported learning."

In this paper I[1] introduce the concept, review current work in the area, and discuss ways in which our research is leading us to push the standard in a particular way. I conclude with some questions that arise from this work.

Learning Object Model

The learning object (LO) model is characterized by the belief that we can create independent chunks of educational content that provide an educational experience for some pedagogical purpose. Drawing on the object-oriented programming (OOP) model, this approach asserts that these chunks are self-contained, though they may contain references to other objects; and they may be combined or sequenced to form longer educational interactions. These chunks of educational content may be of any type—interactive, passive—and they may be of any format or media type. A learning object is not necessarily a digital object; however, the remainder of this paper will focus on learning objects that are exclusively digital.

An associated requirement for learning objects is that of *tagging* or *metadata*. For these objects to be used intelligently, they must be labeled as to what they contain, what they teach, and what requirements exist for using them, and thus exists the need for a reliable and valid scheme for tagging learning objects.

The LO model provides a framework for exchange of learning materials between systems. If LOs are represented in an independent way, conforming instructional systems can deliver and manage them. The learning object activities are a subset of efforts to creating learning technology standards for such interoperable instructional systems.

Benefits

The first major benefit provided by the LO model is the one imported from OOP-- reuse. A learning object designed by one person is made available to other instructors who can use them for different educational purposes. For example, a learning object that discusses how autos behave differently with and without anti-lock brakes might be used in several different educational domains: the physics of friction, automotive design, or insurance liability.

One of the benefits of the LO model is that it has the potential to reward the best educational content, by allowing objects to 'compete' in a market economy. In this scheme, there are costs to the consumer for the object, costs that are then delivered to the author as rewards. Rights to the objects are made clear, as is the financial responsibility. The objects can be customized, aggregated to produce courses, etc., as the Intellectual Property (IP) owner dictates. Then, as different authors produce different versions of the same content, the economy rewards those authors who produce the most effective objects. The Educational Object Economy (http://www.eoe.org) has parts of this, though they are limited to Java Applets.

Another benefit is to provide search operations for objects that meet a particular category. Instead of doing a web search on "+Railroad +US +Western +Expansion", for example, a teacher might do a search which specified a search for educational material aimed at fourth graders which described the western expansion of the railroads in the US, particularly which incorporated maps. This same capability could be used by

---

[1] While I wrote the first draft, important revisions have been made by Brendon Towle, Cindy Mazow, Edwin Bos, and Dan Christinaz. They substantially improved it; all remaining errors, of course, are mine. It is hoped that they will participate in the discussion as well.

learners to aid in their own educational processes. The richer the tag set, the higher the likelihood of being able to craft a query that generates a precisely targeted set of candidates.

Tagging

In any system that uses learning objects, the objects are manipulated by the system independent of their content, at least until delivery to the learner. Consequently, the objects must be tagged to indicate many things about the content. Tags have a syntax that indicates the name of the field or domain of the tag, and the value attached to that label. For example, the field might be author, and the value for this article would be "Clark Quinn".

Some tags are necessary, independent of educational use. Such tags would include technical issues of format, size, and delivery requirements. Other categories are authorship, ownership, and might include information about who did the tagging. Information might also track version number, status, and other issues associated with a lifecycle of the object. It might also indicate if there's been annotations or aggregations.

While tags like this are certainly useful, one can imagine a number of additional tags that might be useful for educational purposes. For example, it would be desirable to tag learning material as to the content. For objects at the level of courses or books, we might consider using any established library scheme, such as the Library of Congress subject headings. If our objects are smaller, how do we address this? Any librarian can tell you (and you should talk to them, they've been trying to solve this problem for years) that there is no overarching ontology that accounts for all knowledge. So unless we aggregate individual objects into larger buckets and label the buckets, we haven't solved the problem of semantic content tagging. If we do aggregate, we limit the flexible reuse of objects.

There are other tags to consider, as well. One, particularly for smaller objects, is the instructional role of the object, as well as instructional characteristics. Is it informational, or does it require activity on the part of the learner? Other questions might include how focused it is, whether it has navigation requirements, or whether and what the form of feedback is.

Others have supported the learning object approach, notably Merrill (1998), but there is lack of agreement on what needs to be indicated. While theoretically it might be valuable to err on the side of over-specification, pragmatically there are reasons to limit the amount of detail. The tradeoff, of course, is that for greater effort, you get greater power. The question is: where to draw the line?

Current coverage

There are several activities in progress to develop a tagging scheme for LOs, including the Dublin Core, the Instructional Management System (IMS) project, and the Learning Technology Standards Committee (LTSC).

The Dublin Core initiative was an early effort to standardize on what the core tags for any information object should be, and has been remarkably successful to the stage that most standard efforts start with the Core. The Dublin Core is now separately investigating the special case of educational objects (independently of the other ongoing work).

The Instructional Management Systems project of EduCause has made a tagging proposal that has achieved the level of a first specification (http://www.imsproject.org/metadata/index.html). Their work has passed on to the IEEE's LTSC, particularly working group 12, and is the basis for further work in this area. The LTSC have a draft that is close to voting standard (http://ltsc.ieee.org/doc/wg12/LOM-WD3.htm). Notably, the LTSC is having the work forwarded to ISO to work towards an internationally accepted standard.

The bottom line is that there is considerable work going into object metadata that the educational technology community needs to be aware of.

Currently, the LTSC proposal includes tag categories of: General, LifeCycle, MetaMetaData, Technical, Educational, Rights, Relation, Annotation, and Classification. Most of these are true of objects regardless of purpose, and would be true of knowledge objects as well as learning objects. It is only the 5th category, Education, that really concern us, though I will occasionally point to some other issues.

I will here note that the Classification category allows the introduction of other classifications for use in tagging. This allows people to propose and use new sets, and it is an explicit goal of the current tagging exercises to leave some difficult issues vague and allow actual use to drive further specification.

The educational category has several types of tags for objects. The first is interactivity type, covering flow of information between resource and user, with restricted values of active, expositive (passive), or mixed. Then comes learning resource type, describing the specific kind of resource (which can be a list, prioritized), and allows any terminology but recommended values are exercise, simulation, questionnaire, diagram, figure, graph, index, slide, table, narrative text, exam, or experiment. Next comes interactivity level, defining the degree of interactivity, and ranges from very low, through low, medium, high, to very high. Semantic density has the same values, and is meant to define a subjective measure of a resource's usefulness relative to size or duration. There are categories for intended end users (teacher, author, learner, manager), context of use (an open vocabulary, but examples include primary education, secondary, higher ed, different university levels, tech schools, etc.), typical age range, difficulty (again, a range from very low to very high), typical learning time. Also included are a space for a text description of the resource, and a language choice from the international standard codes.

Issues

Not surprisingly, a number of issues arise. These issues naturally divide into issues about the characteristics of the objects and characteristics of the tagging of the objects. Under object issues is the issue of level of granularity. Under tagging issues is the problem of vocabulary.

Granularity

Currently, people tend to develop instruction where a complete course is the smallest independent level of learning object. Certainly, that's the easy way. Can we find value in pursuing a finer level of granularity?

Several arguments can be made for a finer level of granularity. First, with smaller granularity, there's greater potential for reuse of objects. If the anti-lock brakes example discussed above had incorporated several problems specific to the insurance domain, for example, its reusability in the engineering domain would be limited. By keeping objects smaller, they are more likely to be able to be reused in different contexts.

Second, there's the opportunity to allow flexibility on the part of the learner, or even to support intelligent processing. If the objects are small enough, and instructional experiences are composed of these objects, then different learners can have different instructional experiences.

While developing an online course, I was trying to move beyond traditional instructional design to consider principles that might support people's choices in sequencing. Perusing different instructional design theories, I was struck that 'problem-based learning' (e.g. Barrows, 1986) provides problems first, before conceptual material, while Laurillard (1993) suggests conceptual material first. It seemed clear that one way I could support learners in determining their preferred learning path was to break material up along the lines of the role in the instructional process, and allow learners flexibility (while preserving a lifeline of a default path that followed a safe and standard approach). That led me to propose that instruction is composed the following components: Introduction, Concept, Example, Practice, and Reflection.

Introduction is material that motivates, activates relevant knowledge, and lists objectives. Concept is a presentation of the relevant abstraction. Examples are applications of the concept to problems. Practice is opportunity for the learner to practice the skill, including feedback. Reflection (as I use it here) is material that cements the learning and prepares the learner to transition beyond the learning experience. This includes reviewing concepts, pointing to further directions for exploration, suggesting ways to practice and keep the knowledge active, and a graceful segue from the learning experience. The smaller granularity provided greater opportunities for learner control.

Granularity is independent of object use, and the tagging standards have granularity (called Aggregation Level), under the General category. They talk about atomic units (raw media data or fragments), collections of atoms (molecules?), collections of collections, and full courses. Here, I am suggesting that granularity at the collection level is the one in which instructionally different individual choices would be made.

Vocabulary

With many tags proposed for learning objects, one stumbling block is whether to determine a fixed and controlled vocabulary for the tag, or to allow authors to extend labeling to meet their own needs (called "open vocabulary with best practice"). Although this is not an easy goal, I argue for a robust fixed vocabulary instead of the alternative, a lack of interoperability. We need categories designed so that authors or 'taggers' (a new job category that's part editor, part administrative) can easily discriminate how a potential object should be labeled and so that the objects are labeled consistently.

As an example that illustrates the issues related to vocabulary, consider the description of 'interactivity level'. We might have objects that are interactive, and we'd like to categorize this. However, I see several problems with using the interactivity level tag as it is now defined. First, it is difficult to imagine anyone using the 'low' category without guidance. If someone creates an interactive object, they are hardly likely to consider it only minimally interactive.

Second, it is not clear what distinguishes a 'high' interactivity object from a 'medium' one. Interactivity can come from several sources, whether navigation, or type of response, or quality and speed of feedback; and any of these sources can vary independently, and be more or less important than the others.

Ideally, we would have conceptual distinctions in a fixed vocabulary, but the definition of interactivity is currently unsolved. In the next best case, we would have categorical, demonstrated examples; and, here, I would argue, you can get traction (like pornography, you know interactivity when you see it). I'll argue that we can create rough examples for such categories. For interactivity level, this might be: no interactivity, page turning/linear progression, multi-dimensional navigation like web pages or multiple choice questions, or rich interaction such as SimCity or Doom/Quake with rich (or seemingly limitless) choice interaction possibilities and rapid feedback. While I am not committed to this particular set of distinctions, I believe this is an achievable and desirable intermediate stage on the path to a fixed vocabulary.

It's not easy to determine categories, nor to attempt to apply them to the myriad types of potential objects, but the guidelines for accomplishing the task can be by example as well as by theoretical principle. In places where the theory is still controversial, we'll need to do it by example.

I recognize that what I propose is not an easy task, but if we do not control the vocabulary, we ensure that systems cannot operate on the data. One important future use of learning object tagging is for intelligent systems, which will only be possible if the tagging is through a predictable vocabulary.

Just briefly, let me extend my interactivity level examples to two other categories—semantic density and difficulty level--to indicate that this is a generalizable approach. For semantic density, we could indicate something to the effect of: concept material implicit but not explicit, or buried in additional detail, as in a story; narrative and illustrated content; direct representations such as expositive text, charts, tables, or graphs. For difficulty, we could consider: introductory material; initial application or overview material; scaffolded practice or detailed example; and full application or for expert only.

Discussion

The sum total of what I'm proposing is a fixed vocabulary for a finer granularity and the discriminating feature (in addition to technical and IP properties) being the instructional role of the object. I'd like to stop here and suggest some questions for discussion.

What about a new instructional design? This suggests a different approach to instructional design, where the components of the instructional process are designed separately and designed to stand alone. Is that a good direction, and why or why not?

What about granularity? This level of granularity provides greater individualization of learning, but at an overhead for authoring. Is it worth it, and why or why not?

What about vocabulary? The powers of a controlled vocabulary are greater automatic processing. The costs are significant debate and perhaps premature limitations. Is the goal obtainable, and why or why not? Is it worthwhile, and why or why not?

What questions haven't we asked? What tradeoffs have I missed, and what are their pros and cons?

References

Barrows, H. S. (1986). A taxonomy of problem-based learning methods. *Medical Education*, 20, 6. 481-486.

Laurillard, D. (1993). *Rethinking university teaching: a framework for effective use of educational technology*. London: Routledge.

Merrill, M. D. (1998). Knowledge Analysis for Effective Instruction. *CBT Solutions*. Mar/Apr 1998, 1-11.

(Published in Educational Technology & Society 3(2), 2000)