Learning Objects and the New CAI:

So what do I do with a learning object?

Introduction

First things first: what is a learning object? The IEEE LTSC/LOM working group definition doesn’t provide much help in narrowing down the number of candidates:

> Learning Objects are defined here as any entity, digital or non-digital, which can be used, re-used or referenced during technology supported learning…Examples of Learning Objects include multimedia content, instructional content, learning objectives, instructional software and software tools, and persons, organizations, or events referenced during technology supported learning. (LTSC, 1999).

Think about it for a minute…what could not be *referenced* during technology supported learning? You could reference anything. So according to the LTSC, the set of learning objects is the universal set of all things. That's not very helpful. Perhaps the only useful definition that can be given (if you'll excuse my operationalism) is "something for which metadata has been created."

With the term learning object tentatively defined, I would like to register a complaint. The learning object movement suffers from the same delirium to which every other major instructional technology effort in recent memory has fallen victim: an almost exclusive focus on "technology" to the virtual exclusion of the "instructional." The LTSC, IMS, ARIADNE, and other groups are busily finalizing years of work on metadata specifications, language bindings, database structures, etc., but to my knowledge none of these groups has addressed instructional issues such as the pedagogical implications of the existence of learning objects online or in the classroom. Where's the "learning" in "learning object?"

We all (those in attendance at this technical workshop) understand the importance of metadata and its role in facilitating the location or discovery of learning objects. We understand that "a learning object undiscovered is a learning object unused," let alone reused. Taking basic assumptions of accessibility (technical issues such as file type, size, availability, and link integrity) and discoverability (the need not only for rich descriptive systems but also usable interfaces to such systems) as given, I would like to address the two most important instructional properties of learning objects: reusability and granularity.

Reusability

One of the stated purposes of the LTSC LOM standard is "to enable the development of learning objects in units that can be combined and decomposed in meaningful ways." Although they haven't stated it directly, one of the major interests in learning objects is their ability to be used more than once, and much has been said expounding the merits of "reusable learning objects."

A difference must be drawn between the *reusability* and *repurposability* of learning objects. By reusability I mean the ability to take a learning object as is and reuse it wholesale. By repurposability I mean the ability to extract portions of a learning object and adapt them to new learning contexts. This is a drastic difference, and the instructional opponents of learning objects have either intentionally or unintentionally missed this point, claiming that a significant degree of reuse will never be achievable because of the unique nature of each learning context. They are arguing against our claim to reusability, but ignore our claim to repurposability. In fact, it turns out that because of the relationship between context and meaning, there is an inversely proportional relationship between an object's degree of reusability and repurposability. This could be displayed graphically, as in Figure 1.

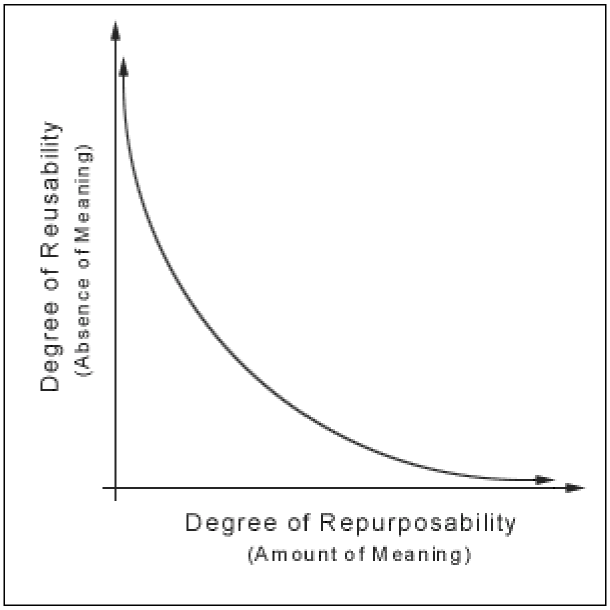

Figure 1. The relationship between reusability and repurposability in a learning object.

Granularity

This discussion leads directly into the second, more significant instructional property of learning objects: granularity. Granularity used to be defined by IMS as "the relative size of the resource," and is now called Aggregate Level by IEEE and is defined as "the functional size of the resource."

The previously stated IEEE purpose for having learning objects that can be meaningfully combined (reused) and decomposed (repurposed) is complimented by another statement from the IEEE list of purposes: "To enable computer agents to automatically and dynamically compose personalized lessons for an individual learner." I believe it is safe to condense these two goals into "to facilitate the automated combination of learning objects to compose meaningful personalized lessons for an individual learner." I would further suggest that regardless of whatever other meaning the lessons may have, if they are not instructionally meaningful they have failed to meet their intended purpose.

There are people who argue that context is everything. Philosophically they trace their roots to Kant and his Copernican insight. Outside of education there are fields such as semiotics that examine the relationship between context and meaning. Dewey picked up on the importance of context in education, and the concept of context finally reached pre-imminent status in the works of Vygotsky and the "constructivists" that followed him. A simple example of the relationship between context and meaning is displayed in Figures 2 and 3.

The picture above is Figure 2. I purposely omitted the caption in an attempt to show the role of context. What is the subject of Figure 2? It is difficult to say without surrounding information to contextualize the image. Someone once responded to this question "a helicopter shining a red light down into the ocean." Great answer, but no. Suppose I was to contextualize this graphic with a label and caption, and call it Figure 3.

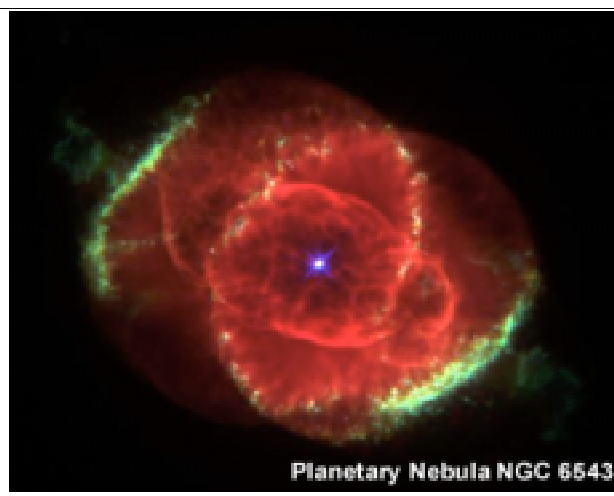

Figure 3. 6543 is 3,000 light years away
in the northern constellation Draco.

Figure 3 has more meaning than Figure 2. Could Figure 3 be used in a matching exercise on an assessment instrument? No. Would it be useful to an art student creating a collage? Probably not. Adding the label, box (to make the relationship between the caption and the image clear), and the caption increases the context of the photo, and therefore the meaning of the object, but it also decreases the number of ways the object could be recontextualized.

The implication of all this context hubbub for learning objects is that we may be defining granularity the wrong way. I stated above that IEEE defines Aggregate Level as "the functional size of the resource." Whatever function the LOM working group has in mind, it does not appear instructional, either from their definition or from their suggested values for the field, and so the granularity of a learning object is simply not expressed in instructional terms.

How are we, then, to meaningfully (i.e., instructionally) combine and recombine (i.e., contextualize and recontextualize), learning objects when the measure of their "size and shape" is not expressed in instructional terms? Another way of asking this question is, "What degree of granularity of learning objects best promotes the instructional use of the learning objects?" Three answers present themselves, which I will attempt to generalize to a single answer.

The course as a learning object

Some people, particularly commercial sellers of textbooks, etc., are supposedly jumping on the learning object bandwagon. What is it they are creating metadata for? Entire courses. Entire textbooks. Huge collections of vast resource. Is it possible to meaningfully combine these? To some degree, yes… a text with a course, or one course with others to form certificate or degree programs. But this is not what the instructional promoters of learning objects had in mind. This seems to violate the spirit of the learning object idea.

The Fundamental learning object

If combining and recombining learning objects is really a matter of meaningfully contextualizing and recontextualizing learning objects, then one way around the problem is to create learning objects with as little inherent or embedded contextualization as possible. This makes them extremely easy to contextualize and recontextualize. We could call the first level of the new CAI "Clip Art Instruction." A digital collection of images, sound bytes, video clips, and animations "pre-decontextualized" to the greatest extent possible could be searched for resources to incorporate into online or in-class instruction. This would be similar to the way clip art CD-ROMs are searched for material to use in Elementary School newsletters. I have elsewhere called the state of being decontextualized to the greatest degree possible fundamentality.

Combination and recombination of learning objects here occurs according to a brick-and-mortar metaphor. Small pieces are arranged in a useful way and held together by some kind of cement. In the case of online instruction, this cement must be an instructional design theory of some kind. This provides not only a blueprint of the manner in which the individual pieces should be laid, but also prescribes the type of cement which must hold the pieces together, e.g., advanced organizers, textual transition material from one idea to the next, etc. The benefit is obvious: compatible basic building blocks that can be economically mass-produced without specific knowledge of their end use.

The danger here is also obvious: although the promotion of resource sharing and reuse has the potential to cut costs, increase the speed of instructional development, and generally encourage the creation of higher quality instruction, it also opens the door to the same mindless display of gaudiness apparent in Elementary School newsletters. There are probably as many professors trained in instructional design as there are elementary school teachers trained in print design.

Granularity and Instructional Architecture

A third method for determining the proper level of granularity of a learning object is to agree on a contextualization scheme before hand. Adopting an instructional architecture, or instructional event model, can provide detailed specifications for the type and amount of context to build within a learning object. By instructional architecture, a "model that provides for all the events of the instructional process" is meant. One example is Gagne's Nine Events of Instruction. If developers were to adopt this model up front, learning objects could be built to fulfill the specific requirements of each step in the instructional process. Then, any learning object which meets the requirement "stimulate recall of prerequisite knowledge" for music theory instruction can be substituted in the place of any other, provided that certain assumptions are met. (Instructional designers would need to agree beforehand on the details of the architecture implementation, e.g., in the above example, the specific objectives to be achieved by the student, etc.)

This example is not meant to imply that this approach is only useful with older, behaviorist techniques. It could also be used with cognitive or constructivist approaches to learning, such as discovery or problem based-learning, etc. The only use restriction is that a robust model of the instructional process, i.e., a model that accounts for all the critical events of instruction, must be available for use as the instructional architecture. Instructional approaches that have not developed to this stage of maturity and specificity would not be applicable.

One construction metaphor for this type of learning object making and combining is the LDS temple in Salt Lake City, Utah. For this building a detailed blueprint was prepared beforehand which specified the exact shape of each granite block to be used in the structure. These specialty blocks were then laid directly atop each other without the need of cement or any binding agent.

This approach also has costs and benefits. The largest cost would first appear to be that each block must be specially cut to fill a specific purpose, which rules out the cheaper mass-production available with fundamental learning objects. However, a computer metaphor shows that this may not be as significant a problem as it first appears.

In the olden days of personal computers, when you bought a machine from a vendor you were stuck with that vendor if you needed replacement parts or support. However, over time the PC architecture matured to a degree of common understanding amongst vendors, to a point where everyone understands what a video card is, what it should do, the type of slot it plugs into (AGP, PCI, ISA), etc. A PC is a fairly sophisticated combination of mainboard, CPU, RAM, harddisk, video card, CD drive, keyboard, mouse, monitor, etc., that function together to perform a specific purpose. Note that in recent computers, different vendors supply many if not each of these main components, i.e., the individual components are mass-produced for a specific purpose. The PC peripheral market is huge.

Mac zealots will at this point be bitter that I have spoken about PC architecture and not Macs. Constructivists were previously bitter that I used Gagne as an example. However, the computer metaphor holds. If you don't like the behaviorist instructional architecture (the PC) then use another, like a constructivist instructional architecture (the Mac). Either way, once the purpose of the whole, the type and nature of components, and the way the components plug into each other is identified, a market can be created in which vendors can specialize and compete with each other, improving the number, quality, and cost of these instructional components available to the "consumers."

There is another, perhaps greater, benefit to learning objects whose level of granularity is a function of their instructional purpose. This type of learning object opens the door to the second level of the new CAI: Computer Assembled Instruction. If the components do actually plug into one another, then a computer could draw from an archive of objects of complimentary functionality and assemble them on the fly. This creates the opportunity to provide truly customized (and re-customizable on demand) instruction, because cosmetic changes (i.e., of attributes not affecting the fulfillment of the object's role within the instructional architecture) could be made to produce "instructional-purpose-compatible" blocks. Just the way that granite blocks could be painted different colors or styled to have different textures, if the basic functionality expressed by the learning architecture is held constant, other variables such as learning style (visual, hands-on, etc.) and level of conation (resistant, conforming, intentional, etc.) could be varied in other instructional-purpose-compatible learning objects. At any given point the learner could ask for more or less structure or a different presentation style, which the computer would be able to accurately retrieve and present to the learner. Navigation through the instructional material and cueing regarding location within the instructional space become completely automatable. Some sample metadata follows.

Educational.InstructionalArchitecture = Gagne9
Educational.InstructionalArchitecture.Level = 3
Educational.IndividualDifferences.Style = Visual
Educational.IndividualDifferences.Intention = Conforming

Although clever AI work could be performed to inspect and make inferences about individual learning objects and their metadata in order to facilitate the dynamic combination of objects into an instructionally meaningful unit, the "low-tech" common instructional architecture solution is not only less prone to machine assembly error, but is immediately (technologically) implementable.

The most effective degree of granularity

At the beginning of the discussion of granularity I asked the question, "What degree of granularity of learning objects best promotes the instructional use of learning objects?" From the examples above it would seem that there are many degrees of granularity that effectively promote instructional use. It would also appear that effective granularity schemes all share two common properties.

1. The degrees of granularity are measured in instructional terms.
2. The standard units of measure (of the degree of granularity) are agreed upon ahead of time.

Measuring degrees of granularity in instructional terms alone isn't enough. NASA recently lost a Mars orbiter because one of the engineering teams was using English measurement units while another team was using metric units. The confusion resulted in navigational instructions sent to the probe that led it to smash into the surface of the planet. Simply saying that we're all using instructional units to measure the granularity of our learning objects doesn't help if you're measuring instructional functionality in Concept Elaboration units and I'm measuring in Collaborative Problem Solving units. The instructional unit of measure must also be agreed upon ahead of time.

Conclusion

While current learning object metadata is capable of facilitating reuse and repurposability at the level of instructional clip art, its poverty of instructional design information suggests that it is incapable of achieving the more worthy goal of automating the construction and delivery of individualized, instructionally meaningful material from individual learning objects. That is to say, it currently seems to be incapable of supporting automated instructional development.

The herein proposed extension to learning object metadata, i.e., the introduction of fields conveying instructional design information (like instructional architecture and individual difference information), combined with a redefinition of granularity ("the instructional function of the resource") could facilitate an immediately technologically implementable method of delivering individualized, or mass-customized, instruction.