

Learning Object Repositories: Problems and Promise

By Gary W. Matkin

On September 22-24, 2002, under the auspices of the William and Flora Hewlett Foundation, a group of twenty – two educators and representatives of the elearning industry met on the campus of the University of California, Irvine to consider the state of the reuse and sharing of learning-related web-based material. These people were brought together with the notion that reuse and sharing of elearning material had not met the expectations of people in higher education and in the broader world of elearning, and, in fact, the lack of widespread reuse and sharing was becoming a crucial limitation on the growth and impact of education delivered over the web. This is a summary of that meeting and is based on the responses of the participants (see Appendix A) and on another summary of the meeting prepared by Reid Cushman, available at http://unex.uci.edu/distance/seminar/more_info/Hewlett_postseminar_report.pdf.

The Context

The elearning industry began to establish itself as a force in the early to mid-1990s as Internet technology captured the imagination of important institutions. Colleges and universities, some of the first institutions to make practical use of the Internet for research purposes, began experiments using web technology to enhance learning and teaching. Early on, there were two aspects of elearning that higher education institutions found very compelling. The first was the prospect that this new technology could, indeed, transform teaching and learning in very fundamental and profound ways. The second was money. Large amounts of money flowed into elearning projects and initiatives from a number of sources. In about 1994 the Sloan Foundation began its extensive investments in what it called asynchronous learning networks (ALN), working through a number of higher education institutions. Higher education meetings around the country began to be dominated by presentations and discussions about the impact of this new technology. At about the same time, private investors caught the same excitement, and began a series of highly publicized initiatives designed to exploit what was thought to be a very large demand for higher education which was not being met by the slow moving, traditional institutions of higher education. Some of the most prestigious universities in the country cooperated in these private ventures.

These early days, coincident with the dot-com phenomena, were characterized by the “first mover” mentality, the idea that those companies and institutions that could move into elearning quickly and with large amounts of capital could secure large market shares which would be hard to assail by late-comers. During this time also the notion that “content is king” was widely adopted. First movers sought to secure the best content experts to create courses using the latest and flashiest technology and the highest level of instructional design -- an expensive process, with the cost of a single course often going into six figures. The creation of elearning content became a major target for investment, particularly when it became clear that the delivery technology could be created or acquired relatively cheaply. The industry was flooded with elearning delivery companies offering turnkey services to companies and institutions wishing to deliver education over the web. This flood was followed by a flood of suppliers, colleges and universities (including very small institutions), private training companies, and corporate universities.

As it became clear that the demand for higher education delivered over the Internet was not as large as investors had hoped and that it was developing more slowly than expected, and as the number of suppliers grew, the focus of investment in the industry shifted and new players entered. First movers became first losers, and it became clear that, at least in the short run, the industry would retain its cottage character, and the development of content slipped a bit in the priorities of investors. The focus shifted to marketing and distribution and to the overcoming of what were seen as barriers to greater use of elearning. The Sloan Foundation dramatically reduced the amount it was willing to give for course development, and the Mellon and Hewlett Foundations, for instance, invested heavily in initiatives that would provide free courseware and learning content to a wide audience, an idea that struck a chord in many universities and has become known as the open courseware initiative. The value of pure content (courseware) dropped, threatening, among others, publishers. The federal government began to invest in elearning. The Department of Education, through its Learning Anyplace, Anytime Partnership (LAPP) program, sought to create partnerships and consortia among elearning providers on the theory that economies of scale and increased learning efficiency would result. The US Army, through its eArmy initiative, quickly followed by other branches of the US military and other governmental agencies, including the IRS, sought to exploit already existing elearning content and programs by setting up special groups of suppliers of elearning degree programs. They also added another dimension by supporting the development of elearning standards in part by requiring that all their vendors subscribe to those standards. The focus of private investment shifted to those aspects of elearning in which groups of learners could be more easily identified and motivated, to business and corporate training with an emphasis on compliance and product training. Universities with already established large distance or “convenience-based” audiences such as the University of Maryland, University College, and the University of Phoenix shifted their delivery to online methods and became very large suppliers.

The Purpose and Conclusions of the Seminar

It is with this background/context that the members of the seminar gathered to consider the state of learning content sharing. Everyone faced a version of the same two issues. How can already existing elearning content, including thousands of courses and hundreds of thousands of definable elearning content elements, be more effectively used? And, how can current investments in courseware and elearning content be structured to maximize future use? The Hewlett Foundation recognizes that the extent of usage is a crucial measure of the return on its investment in, for instance, the Open Courseware Initiative at MIT and the Connexions initiative at Rice University. As it continues its programs seeking to use technology to improve education in several areas, it has a stake in making sure that its investments will pay off as positive learning outcomes for large numbers of students. University administrators seek to improve instruction and reduce the cost of duplication of elearning content both within and among campuses (of multi-campus systems), and also to respond to those faculty wishing to share the content they have developed and to use the content developed by others. Publishers are seeking commercially viable and cost-effective ways of distributing the content they create and own. Professional course developers want easy and inexpensive access to learning content that they can use to include (plug into) the courses they are creating. Teachers and administrators of the K-12 establishment want free and easy (user friendly) access to the huge amount of material generated for their benefit.

While the questions posed at the beginning of the seminar were not fully answered by the end of the one and one-half days of the meeting, the members of the seminar reached some conclusions that we felt would be useful. First, our discussions defined and clarified the dimensions or domains of the problems. Second, as we considered the evolution of the elearning industry, we came to understand that its development was rather lop-sided in favor of the “supply side” and that many of its problems stemmed from a failure to understand the “demand side.” The industry concentrated its early efforts on the development of its products and delivery systems without understanding the market or the real needs and perceptions of its intended audiences. Third, this led us to the conclusion that the increased use of elearning content would depend significantly on the development of communities of users, that offering the content even under the best of circumstances would not be sufficient to drive greater usage without, at the same time, creating communities of users which would in some sense “own” the content. Finally we felt that an important element or criteria to be used in deciding on future investments in elearning should be the degree to which such investment promoted convergence in the elearning industry. That is, future investments should be made with the knowledge of other investments, past and present, in an attempt to build on that experience and bring the many different but interdependent elements of the industry together. While there is still much room for experimentation, duplication is costly, both in terms of the dollars spent and also in the foregone opportunities for faster advancement.

In reaching these conclusions we described some important aspects of the elearning landscape and developed some useful typologies. I will describe these under some general headings.

Defining the Problem

While the members of the seminar came to it with the idea of addressing problems and issues related to the sharing and reuse of learning objects, we spent some time defining the problem and discussing the concept of a learning object. Until now, I have not used the term “learning object” except in the title of this paper, because its definition is so fundamental to the logic of the discussion. There are arguments, some of which were referenced by participants of the seminar, that claim that learning is so contextually based that the breaking up of the learning experience into defined elements (objects), particularly in ways that are required by database technology, creates or promotes a destructive “autism.”¹ This philosophical problem has some practical symptoms. As one tries to define an element of the learning process, whether one calls it learning content or a learning object, the smaller it becomes (say a portion of text in a topic of a lesson of a course), or, in the words of the trade, the greater the granularity, the less contextualized it becomes, the more dependent it becomes on the creation of a context. In a sense, the more useful it becomes, as defined by the possibility of its use in many different contexts, the less interesting it is for any specific use. This dynamic has prompted discussions about whether there is or should be an agreed upon definition of a learning object as opposed to simply learning content. If learning objects are defined, for instance, as having greater granularity than learning content, say, possessed of some coherent combination of exposition, required action by the learner, and assessment, then any element of the object (a test, for instance) would not be separately assessable. On the other hand, if any element of learning content is defined as a learning object, say, a chapter in a textbook, then the problem addressed by the members of the seminar are no different than those facing librarians. One way around this

problem is to conceive of a telescoping or “Russian doll” approach where separate learning objects are created at all levels, course, lesson, topic, page, page element (media), with each level contained in the one above it. For the purposes of the seminar (so that we could get on with our discussion) we decided to accept this notion. Having decided that there could be such a thing as a learning object as distinguished from simple learning content, we could go on to discuss learning object repositories.

However, we also faced another related philosophical issue. Some members of the seminar raised issue with the constant use of economic and business terms to apply to the problems we were addressing—supply and demand, market and market share, elearning industry, products and product lines, production and consumption. At the core of this issue was the notion of an exchange of value that was presumed to take place, an exchange that seemed at odds with the underlying learning experience that was characterized as an internal transformation. We had to recognize, however, that once one accepts the notion that a learning object is something of value, that is in some sense, intellectual property, then one must use vocabulary that describes some form of exchange, or market.

Once we began considering the issues surrounding the development and use of learning object repositories (LORs), a number of more practical concerns surfaced about their ultimate usefulness. For instance, at whatever level of granularity, the use of an LOR may actually introduce inefficiencies in authoring. If one understands the course authoring process to be one in which learning objects are selected, placed in a sequence, and connected (or contextualized) by means of a narrative thread, how does a repository make a contribution? For some, the interruption of the creation of a narrative thread to search for just the right learning object is destructive, and may result in a disjunction. Often, it may be more cost effective to create a new learning object just to maintain the continuity that is so important for the learner. The problem is the reverse for learners—how would a learner find just the right learning object without understanding the context for the search? While we recognized these difficulties, we also know that the impulse for building and using a repository are strong; the overcoming of these difficulties may lie in the efficiency of the categorization (discovery) scheme.

Based on our collective understanding of the problems associated with LORs, we identified several domains of associated issues that needed to be considered. Any one of these domains may contain significant barriers to the creation and maintenance of effective LORs.

1. Standards. For LORs to develop standards must be agreed upon and issued. Standards may apply to the way learning objects are identified, described, and defined and may prescribe technical attributes that must be present. A good deal of progress has been made through the IMS initiative and through the government sponsored SCORM specifications, but more needs to be accomplished. Standards are needed both for specific LORs but also for all LORs so that they can be combined in the future.
2. Technology. Despite recent advances, there remain technological issues in the creation of effective LORs, particularly in the way they may be integrated with user systems.

3. Pedagogy. From the discussion above it is clear that pedagogical issues remain, particularly about how to extract greater learning richness and efficiency from the new technology.
4. Culture. There remain cultural barriers to the use of LORs in society at large, within specific institutions, and within individuals. These cultural issues are often hard to uncover.
5. Intellectual property. The laws in this area are still evolving as are the systems and practices related to the management of intellectual property in learning object form.
6. Organizational development. The creation of effective LORs may require the development of new organizational forms or the modification of existing forms.
7. Business models. New business models need to be developed for the creation and maintenance of effective LORs, models that include provision for honoring the intellectual property rights of contributors and that can develop the resources required to maintain the freshness and usefulness of the repository and its technical quality.
8. Policy. It is clear that current policies need to be changed and new policies adopted for LORs to operate effectively. Academic, intellectual property (policy as opposed to laws), and some financial policies are all candidates for revision.

Underlying these domains are two other aspects of LORs that need to be considered, scale and convergence. There are clearly scalability issues in the developing of LORs. Very large scale LORS will need to operate very differently from small scale LORs, particularly if the creation and maintenance of a community of users is important. Also, there is a strong impulse to build repositories that can ultimately be combined, either actually or virtually. While the standards movement theoretically serves this impulse, differences in attributes not covered by the standards are significant potential barriers to combination (business models, for instance).

Features of Useful LORs.

While describing these general domains of problems associated with LORs was useful, these descriptions fell short of providing the intellectual scaffold needed for an understanding of the real barriers to greater sharing and reuse of learning objects. For this framework we sought to list the attributes of the ideal LOR so that we could then compare this ideal to real examples to discover potential shortcomings. Our first typologies were simple—LORs needed to have systems of getting stuff in, maintaining and storing stuff, and getting stuff out. Another early model was the create, store, retrieve, and publish schema. As suggestions about attributes flowed from the participants, it became clear that these typologies were too simple. Instead we arrived at a list of features, rather than a typology.

Authoring/creation. The formation of a sound LOR should begin at the very beginning, the creation of the learning object. Without the discipline of some minimal structure that contemplates its eventual storage and retrieval, the creative process may have a very small impact, no matter how inspired. The very large body of creative courseware produced by early “hobbyists” in the elearning field has gone virtually unused because it exists in such disparate and often obsolete forms.

Review and input control. LORs should establish standards for accepting learning objects, standards that include peer review or some form of control over the quality of the learning

object. The notion that the users, through some form of evaluation or rating system, will provide the quality control has not proven effective. Opening the gates to all material with the expectation that users will sort out the good from the bad has resulted in LORs becoming so clogged with unrated material that users become frustrated. Among the standards logically set by LORs are standards with regard to compliance with industry-wide standards and specifications (IMS, SCORM, etc), regulatory requirements (Americans with Disability Act), and technical and data management specifications not covered by the industry standards.

Intellectual property and digital rights management. LORs must address the intellectual property issues inherent in publishing material. The way IP is addressed might range from a policy not to accept anything that has any property restrictions on it, to a single, standard, one size fits all licensing agreement. The Connexions project at Rice University provides contributors a choice of several licensing alternatives and restricts users according to the licensing arrangements they are willing to accept.

Modification, update, removal, collaboration. LORs are dynamic entities; their usefulness depends to a great extent on how up to date they are and how easily the material within them can be modified (updated or improved) and removed. For many user communities, the ability to work collaboratively on the improvement of learning objects is important. This collaborative work requires sophisticated version control techniques and records. The essential dynamic quality of LORs is what sets them apart from content repositories, particularly those created and maintained by library professionals who are usually motivated by strong archival instincts. To be useful, obsolete or substandard material needs to be periodically removed from the repository, which means that criteria for removal and removal processes need to be established.

Retrieval systems. LORs need elaborate retrieval systems designed to make it easy for users to find the elements of the repository they want and need. This usually means some form of authentication scheme to at least identify the user or the classification of the user, and to bar access to users not qualified to use the LOR. Good retrieval begins with a sound organization scheme for the data but also usually depends on sophisticated user interfaces designed to help the user understand the data organization and how his/her needs might be met. Creation of these interfaces requires knowledge of the motivation of the main body of users.

Publishing (getting stuff out). LORs should provide users with the ability to extract and use individual learning objects. Again, this means some form of digital rights management as well as the technical capability of sending material out. Of course, different types of users have different technical needs. A student may need to be able to retrieve a learning object over a low bandwidth computer and be willing to accept lower quality, whereas a professional instructional designer would require very high quality output. Interoperability is a key requirement in such publishing schemes.

Business model/ecommerce. LORs need a sound business model from which to operate, one that assures the long-term financial ability to maintain the repository and to serve its users. Achieving this sustainability usually will require some form of ecommerce system, designed to extract some value from users of the system and/or to manage intellectual property. The business model should provide funding for marketing the system or recruiting users.

Technical infrastructure. It goes without saying that LORs need an appropriate technical infrastructure. The hardware and software configuration, which includes elements establishing security for the system, telecommunication capabilities, continuous updates to the system, and appropriate operating and maintenance staff are essential, and they usually represent significant costs.

User information. The LOR should generate information about its users and how the material within the LOR is being used. Gaining information about the use of the system is invaluable in understanding how the system should be changed or updated and also how it should be marketed.

Community support functionality. If it true that LORs are really serving communities of users, then LORs should provide some form of support to those communities. The ability for users to communicate with representatives of the LOR and with each other is highly desirable. But directories of related resources, newsletters, and special web events might also serve user communities.

Scalability. Finally, since LORs are, by their nature, designed to grow, it is important that LORs be “scalable,” that is that they remain effective at higher and higher volumes, of material in them, of users (including simultaneous users), and of transactions, however defined. An important aspect of scalability is the effectiveness of the retrieval system to continue to deliver the right object to the user despite the proliferation of objects.

Conclusions of the Seminar

With these attributes in mind, the participants in the seminar chose to examine some specific examples of potential LOR applications. One such example was algebra. There are a number of LORs with wonderful learning objects in algebra, a subject that lends itself to the creation of learning objects. These LORs are relatively well known and yet the use of them is relatively small. From cases like these, considered against the background of our discussions as described above, the members of the seminar came to the following conclusions.

First, LORs, following a trend in the elearning industry, have been developed with a “supply-side” mentality. They have been structured in a way consistent with the underlying structure of the content as understood by the content producers rather than for the ultimate ease of use by the consumers of the content. Gross and untested assumptions have been made about the audiences for LORs and, in fact, often no studies or market research has been done to try to understand the needs and qualities of the potential audiences for LORs. No attempt is made to segment the audience for a particular LOR into logical user groups; algebra students and teachers are assumed to be parts of a single large audience that also might include sophisticated course designers and members of the general public. The failure of some repositories to attract more users may be because not only does one size not fit all, it really doesn’t fit anyone. In order to be more successful, LORs should be demand driven, designed to deliver value to the intended and well defined users in ways and with support appropriate for those users.

This demand side focus led us to a second conclusion, that users of LORs naturally form themselves into communities, communities that can be fostered and supported in ways that create

a sense of ownership of the material in the LOR among members of the community and a sense of loyalty to the LOR. The formation of these communities and their continued viability may be the higher goal for LORs, surpassing the goal of creating a cost effective exchange of learning objects in support of greater efficiency in the elearning industry.

Finally, as we completed the vast array of learning objects available from many LORs, we were struck with the potential for greater waste and needless duplication. It probably makes no sense to create yet another undergraduate course in statistics, just as it probably does not make sense to create another LOR for algebra. Further investments in elearning should promote convergence or at least should be made with knowledge of what other investments have been made. While early investments were understandably a bit scattered and highly experimental, and while there is still considerable room for experimentation, it is also possible to be more strategic about such investments, filling in the gaps that have appeared in the fabric of the elearning industry and building on the successes and failures that have occurred.

In general, the participants in the seminar reported that they had gained new insights and new perspectives on the issues that face us all as we watch and participate in the evolution of elearning in our own institutions. Despite the philosophical debates that surround them and despite the large number of obstacles to be overcome, there is little doubt that some standard forms of LORs will emerge and become central to the elearning industry and to organizations and individuals who engage in learning using the new technologies.

ⁱLambe, Patrick. "The Autism of Knowledge Management," www.greenchameleon.com.