

From Local Challenges to a Global
Community:
Learning Repositories and the Global
Learning Repositories Summit

Version 1.0
November 11, 2003



The Academic ADL Co-Lab
with support from
The William and Flora Hewlett Foundation

by Colin Holden
Academic ADL Co-Lab

Table of Contents

Preface.....	2
Introduction.....	4
Creating Quality Content and Metadata	6
Ensuring Access to Quality Educational Content.....	7
Quality and Consistency of Metadata.....	14
Awareness of User Communities.....	18
Creating Networks of Learning Materials	21
Communities of Learners and Communities of Learning Repository Leaders	24
Conclusion	27
References.....	29
Learning Repository Initiatives Explored.....	30

Preface

About this document

On the 7th and 8th of October 2003, the [Academic ADL Co-Lab](#) hosted its first [Global Learning Repositories Summit](#). The Summit, with assistance from The William and Flora Hewlett Foundation, brought together representatives of diverse learning repository projects to share information, identify and describe common challenges, and work towards approaching these challenges collectively.

A group of 45 persons, made up of thinkers and practitioners within learning repositories, digital libraries and e-learning, some representing learning repository projects from as far away as Australia, Scotland, and Belgium, was brought together to describe their projects, debate, and propose next steps for the future.

In preparation for the Summit, the Academic ADL Co-Lab conducted preliminary research into the state of learning repositories, producing before the meeting itself a description of the common challenges faced by many learning repositories, and an extensive listing of learning repository projects. This current paper expands upon the issues described in the initial description of challenges to learning repositories, and adds to them the insight and comments given over the course of the Summit.

The initial paper synthesized what the Academic ADL Co-Lab identified as the major important issues faced by most learning repository projects. All learning repositories, the paper argued, were concerned with the mechanical and institutional problems surrounding the collection or creation of quality content and the creation of quality metadata. As existing learning repositories were working to develop distributed architectures through which to share metadata from multiple collections, it was imperative that repository projects be aware of the diversity of institutional structures that informed the collections and metadata of various collections.

Combining materials both from the preliminary research paper and from the meeting itself, this paper is a guide to the various issues challenging learning repository projects today. Issues of quality, both of content and metadata, addressed in the first paper, are discussed and expanded upon. Issues of community definition and building, both with regard to user communities and the learning repository community itself, issues which were frequently discussed during the Learning Repositories Summit, are described and their importance for learning repository success discussed.

About the meeting topic

The Global Learning Repository Summit developed from a desire to share information about innovations and new developments within the realm of learning repositories among those who could benefit most from that information — those developing and maintaining learning repositories.

In the last few years, learning repositories (digital repositories containing resources used for training and education) have been developed to meet the needs of many various populations. They have met their user's needs with varying degrees of success and faced challenges common to all in differing and innovative ways.

Many parties, including the Academic ADL Co-Lab, believed a meeting would benefit individual repository projects for two reasons. First, there was concern that as multiple learning repository projects arose they would develop individual silos of learning materials, collections of materials that would each be accessible through a single point of access. Such a state of affairs would potentially slow movements towards bringing together multiple collections through the development of common search spaces.

The Academic ADL Co-Lab was concerned as well that learning repository innovators and leaders lacked mechanisms for sharing their own expertise and experience. Learning repository projects would benefit from an exchange of information as open as the exchange of learning materials they encouraged among users. Identifying possible directions for community building and collective action was seen as one of the preliminary steps towards collaboration in the creation of a network of interoperable repositories.

As we found, there are some preliminary steps still to be taken before the learning repository community has the community structures and institutional mechanisms that will better encourage collaboration and information exchange.

The original objectives of the Summit were quite ambitious: bringing the many existing learning repositories closer to interoperability, and permitting the sharing of their various collections. Once the various players were together, however, it became clear that this meeting was a first step towards community creation and building, and it was only when communities of interest and practice had been developed that learning repository leaders could move forward and implement plans that involved long term planning and working in concert.

Introduction

Information technology has entered the world of learning and it is here to stay. From the use of computing technologies in the classroom to the creation and delivery of entire online courses, new technologies are changing the ways in which we think about and practice education. With the development of the Internet has come the ability to easily deliver educational materials in electronic form to anyone, anywhere, at anytime. With developments in educational technology comes the promise that educational resources in electronic formats can change the ways in which we teach and learn.

But technologies do not transform societies or habits of practice all by themselves. The resources that develop from them must be easily located and retrieved, and these resources must be well selected to meet the needs of those persons to whom they are delivered. And those persons must know how to use the resources, and understand why it is in their best interests to use them at all.

This is why repositories, systems for the storage, location and retrieval of electronic content, are so essential to the further integration of information technologies and learning. They are the potential agents for breaking through many barriers to the use of new learning technologies.

While people are interested in repositories, and recognize their essential role in delivering educational content, few can agree on their precise definition. The search for an exact definition of the term “repository” has produced many potential candidates, some descriptive and some prescriptive. While theorists have debated definitions, developers and administrators have created networks that allow learning materials to be located and retrieved. In creating these networks, which they sometimes label “repositories,” these individuals often make do in a world shaped more by financial and institutional constraints than theoretical concerns.

Whatever the exact definition of a repository may be, all agree that **learning repositories can and should provide access to the increasing supply of digital educational content.**

The challenges to learning repositories are of two types. The first relates to the internal structure of a learning repository project, concerning those structures that a project will take on as well as the tools it will make available to encourage the development of a collection of quality content and metadata. The second set of challenges relate to the repository project’s relationships with outside communities, either its relationship and communication with user communities or the ways in which learning repository projects exchange information and expertise among themselves.

For the purpose of this paper, the definition of the term “learning repository” will be functional. **A resource is a learning repository if it is created in order to provide access to digital educational materials and if the nature of its content or metadata**

reflects an interest in those materials being used in an educational context. This definition reflects the current lack of agreement concerning the definition of learning repositories; it is not meant to discourage the creation of more precise definitions in the future.

The purpose of this paper is to describe issues unique to or especially relevant for learning repositories. This is because the intentions of learning repository initiatives are not solely the storage of quality digital materials. **Repositories are created in the hope that a community of practice will arise around them and that this community of practice will use the materials made available to it and, in many cases, improve the collection of those materials.** Any successful repository project will have to be successful at creating this community. The technical and administrative tasks of repository management within the context of education as a whole are what make “learning repositories,” into a subset of “content repositories,” with roles and challenges specific to them.

Unfortunately this work cannot be a comprehensive review of all the issues of concern to a learning repository project. The scope of this summary paper does not permit a full discussion of some subjects such as digital rights management and the long term sustainability of repository projects. There are many issues common to all digital repositories which are not discussed here.

Issues for consideration

- Do repository projects need to agree upon a single definition of success to benefit from collaboration and information exchange?
- Are there services and benefits other than making available their content that learning repository projects provide that would be strengthened by collaboration?
- What would be the benefit of a precise definition of “Learning Repository?”

Creating Quality Content and Metadata

The contents of any repository, as well as the metadata describing those contents, arise from a repository project's institutional structures and policies. This section will describe a limited number of repository attributes which have the potential to affect repository content and metadata, attributes most repositories have but which vary repository project to repository project. Some of these attributes arise from the techniques that repositories use to store electronic learning materials and to locate them. Examples are:

- The structure of the repository itself. Whether the resources or the metadata describing them are on a single server, or distributed over many servers.
- The application of metadata. Whether metadata is applied at all, what kind of electronic record it is contained in (i.e. XML, HTML, Database), and who creates the content of this record.
- Metadata standards adopted. Whether the repository's metadata schema uses or builds on an existing metadata standard such as [Dublin Core](#) or IEEE's [LOM](#) (Learning Object Metadata).
- Interoperability. Whether the electronic record containing metadata is readable by an off-site search engine, and the metadata contained follows interoperability standards.

Other attributes that vary repository to repository arise from the various ways in which repositories define and practice their relationship with a wider community of educators and/or learners. Examples of these are:

- The repository's intended users. These could be educators who are choosing course materials for their own classes, or learners who will access learning materials directly through the repository, etc.
- The effect of the repository's contents on pedagogical practice. Whether the contents of the repository reflect a specific pedagogical theory or encourage a specific pedagogical practice. This is related to the consistency of repository contents in terms of quality, granularity, etc.
- The capacity of the repository staff to create a community that is aware of its materials and is familiar with how to use them.
- The capacity of the repository staff to create or be part of a community that creates learning materials and is willing to deposit them within the repository.

A review of the materials and resources currently available online shows that many repository projects implemented successful policies regarding some of these domains, but few in all. Some have excellent systems of outreach and peer review but mixed collections of materials. Others have extremely good content but are not well known. Some hold fast to their metadata schema, applying it well and consistently, but lack a large collection of materials. To assess the impact this variety may have on repository content, we must describe what quality is and develop from there a discussion of various strategies.

Ensuring Access to Quality Educational Content

To encourage the use of repositories and thus of online educational content, repository projects should provide users with “quality” materials. However, the definition of “quality” is contextual, a function of the needs of specific learners and specific educators in specific and sometimes unpredictable contexts. An involved and interactive lesson on pendulums might be a “quality” resource for a learner studying physics, but for a teacher who simply needs a visual aid it will contain extraneous and distracting features.

To deliver quality materials a learning repository should:

- Predict the needs of its users,
- Provide users with the means to describe their own needs when searching the repository, and
- Deliver content that meets those needs.

To implement these tasks repository projects must adopt institutional structures capable of meeting user needs.

Almost all learning repositories either receive their materials through submissions from users or plan to receive such submissions in the future. This situation places responsibility for content development in the hands of persons over whom the repository project has no control. It is from within this population that the project must create its desired community of practice.

Repository projects have devised many means for coping with this problem, relying on repository users to perform much of the repository’s collection development. And in this paper many of these means will be described. But we must note that some repository projects do all of their collection development themselves; selecting, reviewing, and cataloging materials under a single roof. The [Eisenhower National Clearinghouse](#) for Mathematics and Science Education, as well as [LearningLanguages.net](#), for example, both have staffs of subject matter experts who locate and select materials for their collections. The size of both project’s collections suggests that the cost of the strategy is not prohibitive.

However, as the majority of projects do take and even depend upon user submissions for their collection development, we must ask whether this strategy can work in concert with a desire for collections made up of quality materials.

Predicting user needs

Learning repositories have tried to predict the needs of its users in a number of ways. They include:

- Allowing users to submit, rate, and comment on resources, thus allowing members of the user community to specify multiple definitions of “quality” themselves;

- Delivering content only after it has been reviewed by repository staff, or reviewing some amount of the repository content; and
- Delivering content to a pre-defined population (defined by subject matter or educational level) whose expectations can be more easily predicted and catered to.

Many users have taken advantage of sites that allow them to rate and comment on resources. This technique, used by commercial Web sites, allows sites to provide some level of direction to quality resources while reducing the labor demand on on-site staff. For repositories that provide access to many thousands of resources the cost of reviewing all materials might be prohibitive. [MERLOT](#) augments user reviews by providing “peer-reviews” performed by onsite staff who are specialists in the subject matter of the materials they review.

Other repository projects target specific professions or fields. Repositories have been established that deliver materials to educators and learners in such specific fields as the Earth Sciences, Computer Science, and Engineering.¹ By narrowing the range of a repository’s content, users can be assured that as long as they are members of that repository’s target audience, they are more likely to locate the kinds of resources they desire.

Describing and finding materials

The most powerful tool currently available for describing and locating digital materials is metadata and a thorough metadata schema. By creating records and making them available for searching, the creator of metadata allows a resource to be discovered by a user who might not know in advance the specific resource they require. By creating metadata records that describe multiple attributes of the resource referred to, the creator of metadata allows a resource to be discovered using multiple searching techniques.

A sufficiently nuanced metadata schema should permit users to locate materials using a variety of terms or phrases as search terms. The contents of each of a metadata record’s fields describe various attributes of a resource such as its title, its creator or creators, or its subject. Ideally, a metadata schema would take into account the needs of its presumed users. IEEE LOM, for example, includes fields of particular interest to educators or learners such as the intended age range of a resource’s users and the resource’s interactivity level.

To aid in the location of desired materials, metadata must meet the needs of searchers. However, repository metadata may meet many needs. Creators of metadata records may also be concerned whether:

- The metadata record is compliant with standards that will permit interoperability with other systems and thus allow records to be searched from multiple locations;
- The metadata record allows resources to be collated or classified with other similar resources through the use of consistent or controlled vocabularies; and

¹ Such projects include [DLESE](#), [SMETE](#), [CITIDEL](#), and [CSTC](#).

- The metadata record accurately describes the material it refers to once it has been located.

These priorities do not always lead to the creation of easily discoverable resources. The use of controlled vocabularies, for example, is very useful for consistently classifying a growing quantity of materials, and ensuring that when a controlled term is used the found materials will share common attributes. But when users are unaware of a controlled vocabulary's terms or meanings, the vocabulary will be unable to serve their needs. A look at Library of Congress subject headings, with such terms as "Cookery" for what would more commonly be called cooking, shows that the terms with which vocabularies are controlled and thus subjects classified are not always the terms with which users will search for records corresponding to those same subjects.

We can describe this issue as a conflict between the desire for precision in resource description, and the user's desire to search for resources successfully using a variety of search terms. Some basic tools exist to allow metadata records to be friendly to multiple searchers while still allowing for precision in resource location and description when they are required. The description field, and the keyword field, used in Dublin Core and LOM, allow for the use of natural language phrases and terms, as well as vocabularies specific to a field or community. Use of these fields grants the creator of the metadata record room to increase the chances that a record will be found while allowing the contents of other metadata fields to be precisely defined.

Controlled vocabularies, when used alongside thesauri, can also allow for both precision in classification and location of materials through a broad range of possible searches. However, while it is not difficult to imagine thesauri functioning in the digital context, the costs of creating and implementing them are large. Some resources have been developed, however. The [ERIC Processing and Reference Facility](#) has an online thesaurus of educational subjects. The Education Network Australia ([EdNA](#)) also uses thesauri terms to encourage resource discovery. EdNA has reported that this "has its challenges — each sector has chosen a preferred thesaurus."

Delivering desired content

A collection is as strong as the resources it can deliver, and a repository's collection development strategy will have an impact on what resources it will contain. The quality of a repository's collection, that is the likelihood that its contents will be of value to users, depends on the decisions that repository administrators make as far as:

- What qualifies a user to submit materials, and motivations for submission,
- What review process exists for submitted materials, and
- What criteria are used to determine inclusion of materials within the repository collection

Most repositories require that a user submitting resources be logged in to the repository interface. Logging in requires that users have created a user profile providing some basic information about themselves and, when required, specified that they are part of a member institution permitted to submit resources to the repository.

Resource submission

Permitting a body of volunteers, even of logged on members, to submit records for resources has its advantages, as is evidenced by the sizes of many repositories. However there are issues that must be faced by any repository project using this strategy.

- It is sometimes unclear or very doubtful whether the submitter of a record for a resource is the creator of that resource or instead simply its discoverer. This opens the possibilities that the same resource, or multiple versions of the same resource, could be submitted by multiple persons. Some techniques should be developed for both detecting duplicates and multiple versions.
- A body of volunteers may not be familiar with the priorities and philosophies that underlie a repository project's mission. Among other issues, it might be difficult to enforce any definition of "Learning Object," or maintain a consistent level of granularity throughout the collection, nor might they create or submit materials that meet repository project requirements for accessibility.
- Volunteers from the educational community, most of whom are not developers, do not always produce or submit materials that take full advantage of the digital environment. If it is in the interest of the repository project to encourage the use of interactive electronic materials that do not have real-world analogs, means must be developed to aid users in developing new kinds of materials, or appeal to a group of users that has the capacity to create interactive media.
- The legal and copyright status of electronic learning materials as intellectual property remains unclear. Creators of these learning materials might be discouraged from sharing them.

Duplication of materials within repositories is a difficult problem to solve. It is very easy for two people, especially volunteer submitters, to produce different records, even to assign different titles to resources, given the ambiguous structures of electronic documents. A small handful of repository projects have made available to submitters the AACR2 cataloging rules, which specify the techniques by which the title, the creator, and other attributes of a digital document will be specified, including the forms that names and titles will take. But AACR2 is complex, and it is unclear to what extent making the materials available online has encouraged widespread compliance with its rules.

Duplication prevention might also be assisted if electronic documents were assigned Uniform Resource Identifiers (URIs), Uniform Resource Names (URNs), or some other globally unique ID. It is unclear at present to what extent these identifiers would work on the Web as a whole, and how volunteer submitters of mixed technical backgrounds would be able to identify and use them when submitting sources.

Some techniques have been developed for encouraging consistency in granularity across collections of learning. One successful, though indirect, method to encourage consistent levels of granularity has been to collect specific kinds of software that deliver educational content or activities. Thus, collections of Java Applets or Flash files, for example, will contain materials with certain uniform attributes and relatively consistent levels of granularity established by the software.

One project, the [Wisconsin Online Resource Center](#) has a permanent staff that takes ideas for learning objects from educators and then creates a Flash learning object according to the project's own specifications. This maintains consistency within the collection, though it does slow down the production of materials.

A survey of the contents of the largest repositories shows that many of the online materials currently available for teachers, both in learning repositories and elsewhere, have real-world analogs. Such materials include lesson plans that are made up entirely of text, tests or review sheets that must be printed out to be used in class, or information for teachers about instructional technique or pedagogy itself.

The Internet is a delivery system and it should come as no surprise that when educators are given access to it they will use it to exchange materials that they are already familiar with and already exchange by other means. Many educators are trained in traditional teaching techniques and thus lack sufficient drive to exploit the potential that information technologies have to not only deliver content, but to run new kinds of interactive content that exploits the electronic environment.

The limited availability of digital technologies in many classrooms, especially K-12 classrooms, limits the potential demand for, and the capacities of many educators, to create materials that exploit digital technologies. If the technology is not available, educators will demand materials that can be delivered to students through traditional pedagogical means.

Some repository projects have responded directly to this demand for traditional learning materials. The [Gateway to Educational Materials](#), for example, collects lesson plans, projects, and other materials to be used in class.

Many repository projects have charted a course towards changing the ways in which learning technologies are used in the classroom, and have done so by collecting resources that meet specific criteria for either quality or educational style. The UK's [National Learning Network](#) provides users with a collection made up of interactive, browser based materials that can either be integrated into the classroom setting or accessed by learners directly. Another project, [Connexions](#), working out of Rice University, contains materials whose content and metadata fit a very specific XML format that permits resources to be easily edited, linked together in sequences, and allows content to exist independently of any one template, background, or lesson design (Henry 2003). The Canadian based Co-operative Learning Object Exchange ([CLOE](#)) has proposed a tentative system of incentives to create online learning materials of high quality. Users would participate in an online exchange of learning materials and continued participation in that exchange, and thus continued access to learning materials, would depend upon submission of materials that were accessed and used by others (Harrigan 2003).

Means for dealing with concerns about digital copyright remain under discussion. One solution, adopted by the Connexions project and others, is to encourage users to use [Creative Commons](#) licenses for all materials submitted to their site. Creative Commons

has developed machine readable metadata that can associate creative works with their public domain or license status.

Review and review criteria

There are repository projects that review some or all of the materials submitted to them. [The Learning Matrix](#) has a staff that reviews submitted materials. Of course, increased project staff can both drive up costs and slow down the posting of materials into the repository.

Many sites have developed a review process for submitted materials, thus allowing them greater control over the contents of their collections. This process allows for many benefits in the realm of control over a consistent and managed collection of materials.

We should not see user input as a potentially corrupting force for repository collections. Users respond to repository projects that accept their input, not only by allowing them to submit resources but also allowing them to comment on and rate resources, as well as create their own collections of recommended resources. MERLOT allows users to create their own personal collections of favorite resources that other users can then browse through. MERLOT users can also list assignments or units of instruction that they have created around the materials referred to in the repository collection. This creates a community actively involved in finding ways to bring digital content into the educational environment.

To encourage the use of digital content in teaching and learning, a repository project would do well to encourage the development of a community that develops, uses, and improves upon digital learning materials. Making the repository site into a place for information exchange and community building among educators increases the likelihood that users will visit the site, and thus use its materials and submit to it.

User profiles

Many repository projects are interested in tools that would allow the repository system to recommend materials to individual users based upon those users' past behavior within the repository site. The Science Math Engineering and Technology Education ([SMETE](#)) Digital Library is developing principles and tools for creating user profiles, a preliminary step in the creation of a recommendation system (Muramatsu and Wiley 2001).

Because of these developments, and because many learning repository projects track and analyze user behavior on their own sites, Learning Repositories Summit attendees explored issues arising from recording and storing information about users.

When tracking user information, it is important that a learning repository project keep in mind:

- Differing cultural attitudes towards the capture and storage of personal information, and
- Limitations of tracking technologies to generate the information that learning repository projects desire.

Attendees reported that willingness to have personal information recorded was dependent on a user's cultural or institutional attitudes regarding trust, and trust in authority. They also warned that tracking user behavior was not as reliable a source of information as one may think. They reported that use of materials differed from the downloading of materials or the visiting of a site. One attendee reported that he had stopped recording students frequency of downloading or accessing class materials in one of his classes because he found that with the copying and exchanging of materials there were more methods for students to acquire those materials than the one which he was able to track.

Repository content and pedagogy

During the Learning Repository Summit many attendees discussed the various educational and pedagogical philosophies that would affect the contents of learning repositories. Some attendees discussed whether digital learning materials should be constructed and collected in such a way that aggregations of those materials would easily knit together into seamless learning experiences.

There was no agreement among Summit attendees regarding the responsibility of learning repositories to pursue policies that facilitated the seamless aggregation of digital learning materials. One attendee pointed out that in the realm of non-digital materials, teachers and learners are accustomed to drawing upon a number of distinct types of media. In acquiring information, a student will turn to an encyclopedia, a history text, a specialized reference text, or digital reference resources. The capacity to understand this aggregation of materials, materials that do not easily develop into a seamless sequence of learning experiences, is one of the positive consequences of education itself.

Another attendee, who preferred to think of learning repositories as digital libraries, argued that if the demand for a specific type of learning material requires complex or tiresome labor on the part of users, this demand will go unheeded. He preferred to imagine that the principle of Ockham's razor applied to the realms of digital materials in learning and to learning repositories. The most successful methods for meeting challenges will be those that are simplest. Any solution that requires that users and administrators take on new and unfamiliar tasks for the sake of possible future benefits cannot be the solution that succeeds through wide acceptance.

Issues for consideration

It is clear there is a great quantity of digital education content. But as we develop ways to collect these materials we must be concerned with:

- What is the role of the general public in developing a repository's collection? Does public input further the mission of all repository projects? When does it hinder it?
- How can we create metadata records that best meet user needs to identify and retrieve those records electronically through a variety of search techniques while allowing these records to be classified according to consistent rules?
- What balance can repository projects strike between allowing users to exchange learning materials they deem useful, and encouraging the use of materials that users might not be familiar or comfortable with? Are all projects even required to be concerned with the later?

Quality and Consistency of Metadata

When discussing metadata we must distinguish between the various roles that metadata plays.

Some uses of metadata and metadata standards include:

- Describing materials in ways that aid users in identifying those that meet their needs;
- Providing access to those materials through multiple access points, and thus through multiple potential search strategies;
- Collating or classifying, similar materials; and
- Permitting compliance with metadata standards allowing interoperability with multiple collections.

It is also important to note that the implementation of metadata can be described as being made up of three tasks:

- The adoption of a metadata schema,
- The creation of a workflow through which metadata records will be created, and
- The actual creation of the metadata records.

These various and complex issues shall, for the purposes of this review, reduce down to the following concerns:

- Who should produce metadata records?
- What tools should be available for the creator of metadata?
- How can the creator of metadata records be encouraged to create records that meet standards and/or the needs of various users and administrators?

Who should produce metadata records?

The metadata that preceded digital technologies, library catalogs and the like, were produced by trained groups of professionals. Book catalogers do not discuss whether the library record of a given book should be created by that book's author or by those who donate the book to the library. The extent to which learning repositories currently depend on non-professionals for the creation of their metadata is unprecedented. We can view it to some extent as a giant social experiment, created by institutional and budgetary constraints, of the ability of the general public to create useful metadata.

Much like with content, the definition of "quality" within the realm of metadata records is difficult to specify. As the definition is contextual we can define a quality metadata record as a record that is useful in a number of different contexts, both with respect to the search strategies and terms that can be used to locate it. There are issues that any repository project must tackle if it chooses to grant general users the responsibilities of creating metadata.

Many learning repositories allow creators or submitters of materials to fill the metadata fields within the records for the materials that they submit. This allows the repository to receive ready-made records for its materials, thus avoiding the logistical issues surrounding the creation of those same records by on-site staff. However, it is difficult to both allow the public to create metadata records and maintain a certain level of quality and consistency within the metadata of the collection. Some collections that leave metadata submission open to the public suffer from blank fields and fields containing information of ambiguous quality.

An investigation of sites that depend upon volunteers for the content of their metadata suggests that while many are comfortable with the “description” fields that many schemas make available, providing long and discursive descriptions of their resources, they are less willing to fill out other metadata fields with more specific kinds of information. Thus records are created which are descriptive and perhaps findable by a searcher using the same vocabulary as the submitter of the resource, but are not responsive to search techniques that enhance precision or are aimed at finding pre-defined classes of resources.

Recent research, made available to Summit attendees by Dr. Terry Anderson (Athabasca University), shows that many of the metadata fields of various metadata records within repositories were left blank by users. This was especially, and ironically, true with respect to IEEE LOM metadata fields that described attributes particular to learning materials (Friesen and Nirhamo 2003).

While many repository projects provide documents guiding and advising users in the creation of metadata and the maintenance of any controlled vocabularies the specific project might maintain, it is unclear to what extent the availability of this information actually affects user behavior. Techniques are required to create a culture of practice that creates metadata that meet the needs of both repository administrators and users.

It may be that repository metadata does not need to hold to specific practices in the creation of metadata, precise vocabularies, or subject classifications in order to be useful. To assess to what extent repository searchers need or take advantage of chances to make precise searches, searching within multiple metadata fields simultaneously, and making use of controlled vocabularies, more work must be done to study how users search for learning materials.

Tools and workflow

Many repository projects have developed systems for the creation of metadata that involve multiple persons, each assigned distinct tasks. Some have allowed submitting users to assign values to some metadata fields, but assign the filling in of others (requiring what the administrators believe to be some level of expertise) to onsite staff. Others allow onsite staff to review, assess the accuracy of, and correct metadata provided by the resource submitter.

Many repository projects have developed interfaces for the creation of metadata. Thus, normally through a series of forms, users go through a step-by-step process in which they must fill out some metadata fields, and are constrained by drop down boxes to use controlled vocabularies, as long as those vocabularies are relatively small.

Such tools include GEMCat, developed by the Gateway to Educational Materials, cataloging interfaces for the Heal repository, and the resource submission page for MERLOT.

Vocabularies and metadata consistency

As the quantity of learning materials increases, it is becoming evident that attempts to categorize them should be aided by ontologies and controlled vocabularies. While some sites have their own limited schemas for defining and classifying the subjects for their materials, the growing collections of materials seems to reveal a need for systems that are as comprehensive as the subject classification schemas used in libraries and suited to the electronic environment.

When developing taxonomies, classifications, or ontologies for repositories it is important to discern what these tools are for. A taxonomy or ontology allows similar materials to be retrieved together through the use of a controlled search term, and it allows materials to be browsed by browsing a schema.

The creation of a classification schema for learning materials is a monumental task. The complexity of such systems as the Library of Congress Classifications and Subject Headings reveals how much work and institutional support can go into just one attempt to classify the subjects that define human knowledge.

There have been moves to create taxonomies and ontologies of universal scope. learndirect has developed the learndirect Classification System ([LDCS](#)) which it uses to classify its own training and educational materials. The system is free and available to all in text format. But such schemas have yet to be widely adopted. There are, however, controlled vocabularies and taxonomies that have been adopted by various professional or academic groups.

[MedBiquitous](#), in its own metadata specifications, permits the use of the Mesh, SNOMED, and UMLS controlled vocabularies to describe medical subjects to facilitate searching for and retrieval of materials by persons in the medical profession.

Much of what allows these tools to work can be located within the institutional structures that surround and enforce a given standard. The subject classification adopted by the Library of Congress is not the prominent and frequently used system it is because it is the best system possible. Indeed, any librarian can point out its flaws and inconsistencies. The Library of Congress classifications and subject headings are widespread because of institutional structures that educate persons in their use and create a network of libraries all using the same system. **We can learn from libraries to aid us in the cataloging and classification of materials. But the secret may not be to study the lists of subject**

headings and classifications they have developed, but to look at the institutional structures that transform those lists into real arrangements of real materials across various library collections.

Issues for consideration

- How can creators of metadata be encouraged to create metadata that meet the needs of other users and of administrators?
- How important is it that users be capable of performing precise searching? Does its importance justify the institutional and practical work that would need to be done to create a collection capable of responding to precision in searching?

Awareness of User Communities

In some ways all learning repositories create a community of practice that did not previously exist. A repository project that hopes to draw upon a body of users for its content and metadata must of course create a practice that leads to the submission of both to that repository. But reports from learning repository administrators at the Summit showed that some successful repositories develop from pre-existent communities of practice within which there was a desire for a repository to begin with.

Those at the Learning Repository Summit recommended that administrators of developing learning repository projects:

- Know the communities that they are serving, and
- Recognize and have means to overcome the barriers that might exist within that community to the sharing of materials.

Selecting a target audience and knowing that target audience's needs and practices can aid a repository project tremendously in facing a number of challenges. By identifying a community to serve, a repository project can:

- Identify what cultures of sharing already exist within a community and work to either facilitate sharing that already exists or extend that culture of sharing;
- Use existing vocabularies and practices to develop taxonomies, ontologies and metadata schemas that meet the needs of the given community and can be easily integrated into existing practices, techniques for information access, and workflows; and
- Draw on the desire of the community to exchange information and aid its members. The growing interest in blogs and RSS feeds within the learning repository community reveals a desire within many user communities to share and exchange information (Levine et al.). Repository projects can draw on this desire at the same time as they permit the storage and retrieval of materials.

Developing relationships with user communities is of benefit to learning repository projects for a number of reasons including:

- Sustainability. Meeting the demands of a population of users will open up the potential for new sources of support. This could be either in the form of expecting some payment once a resource has proven its worth, or drawing funding sources particularly aware of the population being served.
- Encouraging awareness of the repository. Learning repositories can depend upon the informal communication networks that communities already have in order to encourage awareness of their materials.
- Feedback and direction regarding policies and processes for the submission of materials and the creation of metadata to be contained within the repository. Feedback informing policies and procedures can come through communities with relatively well defined information needs.

Repositories and community practice

Many attendees at the Learning Repository Summit agreed that a learning repository project must meet the needs of its users and that many previous repository projects had suffered as a consequence of not thinking enough about the community they served. This sentiment was well summarized by one attendee who argued that “standards should not precede practice,” or that repository practices, ranging from collection development policies to metadata standards, should not be developed in the absence of an awareness of user practice and user communities.

Many repositories have been developed to address the needs of a very loosely defined group of “educators,” or “learners,” who would be interested in digital educational materials delivered over the Internet. But those in attendance at the Summit warned that it was possible to expend unnecessary energy producing or gathering resources that were not right for the community a given repository hoped to address. These projects faced the “Enabler’s Enigma,” the problem faced by any project developed in order to enable the creation of a community. The root of this problem is the assumption that the communities of practice that will use repositories in the course of education do not yet exist, and that it is the task of the repository project to create them.

MedBiquitous is a resource that owes much of its success to its identification of a community, persons in the medical and health professions, and establishing a useful controlled vocabulary of terms. The Eisenhower National Clearinghouse, addressing the needs of K-12 educators, also attributes its own success to its identification of a user group whose needs could be identified and addressed, and who could be convinced to use materials that were collected or developed to meet those needs.

Cultures of sharing do exist within specific user communities. But they can also prove infectious. Attendees reported that the resistance to sharing that they had often encountered within user communities was softening in response to the launching of MIT’s [OpenCourseWare](#) project, which makes course materials from many and eventually all MIT courses freely available online.

Listening to the community

Many Summit attendees argued that the key to any successful repository project was that it must “know its community,” but there was disagreement about what this would mean in practice.

While everyone knew that a learning repository project must address existent needs it was unclear whether it was the task of a learning repository to address user’s stated needs or to address needs in ways that the users were not aware of. Attendees pointed out that before their invention there had been no stated demand for a mouse or for Google. They were concerned that the technique of listening to users would limit learning repository projects by leading them to address only those community needs that members of those communities were conscious of.

From the conversation over how to listen to users and respond to their needs it became evident that:

- The learning repository community was working with a limited vocabulary for talking about user needs and responses to them.
- The learning repository community required familiarity with techniques for assessing and responding to user needs that were more nuanced and useful than “asking.”
- Those learning repository projects that had developed techniques for assessing user needs had not formalized them or described them in ways that were immediately transferable to other learning repository projects.

It would be useful in the future for learning repository projects to share how they use a variety of techniques for assessing user needs. Such techniques include use case analyses, and surveys of the repository’s potential user groups in order to gather sets of expectations from these groups. These techniques provide means for information gathering more sophisticated than “asking users what they want,” and more transferable than the tacit knowledge of and familiarity with their user communities that many successful learning repository projects currently have.

The importance that attendees gave to familiarity with a repository’s user community, and response to that community, revealed that the issues that each repository is dealing with are local. Even if every repository has a metadata or content management issue to deal with, the solutions that each repository presents to deal with these challenges will vary with the local conditions from which the challenges arise. Every city has a traffic problem, but the solution to those problems arises from studying the geographies and demographics of individual cities.

Issues for consideration

- What can be done to facilitate the development of learning repository projects within, instead of outside of, user communities?
- Can learning repositories integrate themselves into a community’s habits of technology use?
- What does the call among learning repository leaders to “think locally” mean for the creation of collaborative projects?

Creating Networks of Learning Materials

Distributed repositories and silos

The number of learning repositories is rising. Even as well-established hubs provide access to educational materials, repositories continue to be developed with their own content, and their own metadata. Educators and developers will continue to create their own repositories, and these hubs alone cannot provide access to all available learning materials.

In the face of the constant rate at which new repositories are developed, it is becoming more and more evident that if repositories are to provide access to the largest possible collection of learning materials, means must be developed by which the metadata of multiple repositories can be searched at once.

One of the most exciting developments in the realm of learning repositories, and repositories in general, has been the development of distributed repository architectures. A distributed repository consists of a network of physical repositories, which are capable of sharing metadata among themselves, and that together can function as a single repository. They allow users to search metadata stored within multiple physical repositories.

The distributed repository model stands in contrast to the more common centralized model, which has been described as the “silo.” A silo is a centralized collection of metadata and resources, centralized either because it is located within a specific device or available only through a specific institution or organization. While the database that supports each “silo” may be distributed, a repository can be described as a silo if the repository can only be accessed through a single point of access, usually a single URL.

At present, nearly all repositories can be described as silos, providing access to resources, or the metadata describing those resources, through a single gateway or URL. However, using protocols for metadata harvesting, some repository projects have made their metadata accessible to those who are not accessing the repositories Web page.

Protocols for metadata harvesting require that servers provide metadata that harvesters (server processes), can collect. The Open Archives Initiative ([OAI](#)) provides the OAI-PMH (Protocol for Metadata Harvesting), the protocol most commonly used in metadata harvesting among learning repositories.

At present, repositories that contain harvestable metadata and harvest other metadata records include MERLOT, EdNA (Hendrick 2003), ENC, and The Learning Matrix. [CAREO](#), in its final form, will be a network of collections, each able to access each other’s harvestable metadata (CAREO 2003). [NSDL](#) also sponsors many sites, each of which is required to be OAI-PMH compliant.

Distributed repositories can arise from the interconnection of pre-existent collections. But there are plans to develop distributed repositories that do not link previously centralized collections, but create instead decentralized stores of metadata and content. Portals for Online Objects for Learning (POOL), a Canadian consortium project funded in part by the Canarie Learning Program, has developed a Peer-to-Peer structure for distributed metadata searching. It has facilitated the creation of various collections by making available [SPLASH](#), a downloadable application for storing, searching and exchanging objects using the [CanCore](#) metadata schema, a schema developed from IEEE LOM for use by Canadian learning repositories. This tool permits metadata creation, storage of metadata and searching of the POOL network. SPLASH is a desktop client that communicates with other peers via the Peer-to-Peer POOL protocol (Richards and Hatala 2003).

It is possible to create distributed repositories, with multiple nodes where new metadata and new materials might be entered, the entire repository being searchable from any point within the system, or outside of it using the appropriate tools. At this point it would be wise to assess the benefits of these distributed systems and determine whether they are beneficial in any or all of the various domains in which repositories aid education.

Distributed Repositories, Distributed Institutions

The ability of an organization to act hinges on its institutional structure. Thus, the capacity of a repository to deliver content and meet the needs of its users depends first on the structure of the repository project itself. New technologies are allowing us to imagine new repository architectures, and it is important that we keep in mind that with these technological infrastructures will come new institutional structures, workflows, and distributions of power over repository content.

With the centralization of stored information comes a centralization of labor. If a repository project has a well-defined mission, and pursues it by delivering a specific kind of content, it is in the interest of the repository project staff to maintain firm control over the repository content and the metadata that describes it. It is easier for on-site staff to control these factors if the repository is centralized.

Distributed repositories move us closer to creating a large and comprehensive network of learning materials. But if we are not careful they may also remove many of the benefits that come from the centralization of energy and resources within a single hub. Distributed repositories make available and deliver more content, but repository projects have made institutional and technical innovations within metadata creation, collection management, and outreach of which we do not want to lose sight.

Creating distributed repositories raises the possibility that the contents of a distributed collection's metadata, both in terms of completeness and quality, as well as the completeness and quality of the content referred to, will vary with the policies of the individual collections from which the metadata are harvested.

It is important to ask, as we develop more and more distributed repositories, how collections drawn from different institutional structures are to be combined. Should users and administrators have information about the collections that make up a distributed repository? Will they be able to use this information to search one collection but not another? Or can the various collections within a repository work in concert in collection management or metadata creation?

Distributed repositories and metadata

One of the consequences of the drive to harvest metadata and encourage distributed repositories has been the development and promotion of communication protocols and metadata standards in the interests of promoting compliance across multiple repositories. This compliance of course, makes it easier for multiple repositories to share their collections with other repositories or standards compliant tools.

This has led to a change in the way some repository project leaders talk about metadata. Compliance with metadata standards is obligatory if distributed repositories are to function. Thus discussions of metadata can frequently turn into discussions concerning *compliance* with metadata standards. Yet as these standards are encouraged in the interests of interoperability we must not forget the other functions of metadata — to describe and locate resources, and to collate records.

The metadata standards with which we tie together the various collections of a distributed repository will either augment or limit the powers of users to describe and locate the resources they require. With the development of RSS feeds, which have no metadata requirements, the task of structuring metadata across collections becomes a more complex.

Some repositories have developed metadata schemas that facilitate both interoperability and the other uses of metadata. EdNA, which harvests metadata records from a number of Australian collections of learning materials, requires that harvested materials be compliant with Dublin Core or the EdNA Metadata Standard Version 1.1, which has been developed to meet the needs of educators and learners locating materials.

The Eisenhower National Clearinghouse, as was reported during the Summit, has developed crosswalks between metadata standards. These crosswalks allow ENC metadata to be expressed in Dublin Core or IEEE LOM compliant formats, as well as in a format compliant with USMARC, a standard used by library systems (Lightle 2003).

Issues for consideration

As we assess the powers of distributed repositories we are rightfully impressed by the quantity of materials they can potentially find within multiple collections. But it is worth asking:

- Does a distributed repository require similar institutional structures or habits of practice within each of its collections? If not, why not? And if so, how will uniformity among habits or practices be encouraged or enforced?

Communities of Learners and Communities of Learning Repository Leaders

Educators and learners demand a wide variety of digital content, and as digital technologies further establish themselves in learning environments, learning repository project leaders must decide how they will respond to a demand for content that is growing both in size and diversity.

The next few years will likely see the creation of distributed networks of learning materials. But the kinds of networks these will be and the needs they will serve will be functions of the technical and institutional structures that make them up. There is a difference between exchanging information and creating lasting catalogs of resources. The balance that is struck between free submission of resources and metadata on the one hand, and control over the quality of both on the other, will determine how the repositories of the future will situate themselves between the extremes of the library catalog and the searchable weblog.

To work towards the development of technological infrastructures that reflect the intentions of developers, the learning repository community must create social and professional infrastructures of its own to exchange information and to be a basis from which to organize collaboration.

The Global Learning Repository Summit revealed no technological barriers to collaboration. Instead the forces limiting moves towards collaborative projects such as single search spaces and distributed repository networks were to be found in the limited interaction and information exchange among projects and project leaders. The Summit proved to be a step towards creating connections allowing for future collaboration.

We can imagine three levels of interaction among learning repository projects:

- Awareness of each other's work,
- Mutually beneficial exchange of information, and
- Collaboration on projects and features.

It is useful to think about these levels in order to understand how the creation of networks for information exchange is essential for the easy recognition of opportunities for collaboration and action upon them.

Collaboration breeds more collaboration. By combining the tasks and projects of various parties, initial collaboration will create institutional structures through which projects must increasingly consult each other when making decisions as the consequences of those decisions will increasingly affect persons and priorities outside of any one project.

By creating networks and mechanisms to allow for what one attendee called "more collisions," more fruitful interactions between persons with similar interests or goals, we

can facilitate the development of various communities gathered around common interests, practices, or theory. Many at the Summit reported that the many forums within which information useful to the learning repository community was stored and exchanged were scattered discontinuously throughout the globe and the Web. There were multiple communities of inquiry familiar with and referencing the research conducted within each community, but unaware of other similar communities and the work they were doing.

The learning repository community lacks established mechanisms for information exchange and discovery because the community itself is made up of persons who come from their own distinct intellectual communities, be they computer science, library science, or education. Individuals often draw upon the literature relating to learning repositories arising from and being exchanged within their community. The Summit revealed that mechanisms need to be developed by which information held by the various attending parties could be disseminated. One attendee hypothesized that were everyone in attendance to write one page about what they had experienced and discovered over the course of their study or administration of learning repositories we would be provided with a wealth of valuable information. Another, describing research preliminary to establishing a learning repository project of her own, reported that she had found many similar projects online, but would have benefited from an information source that contained information and updates concerning learning repository projects and presented information in a way that was useful to those developing or maintaining learning repositories.

What the learning repository community needs is an open content community much like those its members encourage among users. Many attendees expressed an interest in potential online resources that would facilitate the free availability and exchange of information.

Repository projects are all at varying stages of development. And those with theoretical or tangential interests in repositories have them for varying reasons. These varying conditions will dictate the topics that repository project representatives and others will want to exchange information about and the information that those same persons require.

The information needs of attendees varied. Many, who were developing repositories of their own, had questions that could not always be answered by the technical, theoretical, and pedagogical debates that those with more experience with repositories were willing to and interested in conducting. Means other than the Summit were required to allow those with valuable experience to provide critical foundational information to those who needed it.

This complicating factor revealed during the Global Learning Repositories Summit has already been pointed out within discussions of open content communities (Cedergren 2003). In a situation in which there are providers and users of information, creators will provide the materials they would like to offer, not necessarily what the consumers of those materials or information want. And users do not have an adequate exchange mechanism to affect the behavior of creators. The creation of a free exchange of

information among learning repository leaders and innovators requires that these issues of motivations be addressed.

Issues for consideration

- Who is responsible for facilitating the creation of communities between learning repository leaders?
- Do learning repository leaders require traditional forms of community building and communication such as professional societies and journals?

Conclusion

Think globally, act locally

The Global Learning Repositories Summit was a chance for learning repository leaders to come together and discuss possible solutions to common challenges. Over the course of the meeting, attendees found that the young learning repository community was being pulled by seemingly contradictory forces.

Members of the learning repository community came to agreement that many of the problems facing learning repositories, relating to metadata, collection development and community outreach, had local solutions arising from the contexts in which specific challenges develop. Most of the solutions to most of the challenges faced by learning repository projects would not be best decided upon and imposed from outside of the community that the repository served.

This tendency to think locally when seeking solutions to problems posed a challenge to any attempts to build up communities of interest and practice within the learning repository community itself. The local nature of some challenges prevented any easy collection of wisdom into best practices or a set of principles applicable to all learning repositories.

The excitement with which many came away from the Summit, and the hopes described for the future, was evidence that the learning repository community would benefit from community building and the sharing of information. The next step is to describe and bring together the communities of interest and practice that will eventually exchange information and collaborate.

In the closing session of the Summit, attendees described many next steps that would build upon the events and discussions of the two days. Many proposals were raised, but threaded through all were these common themes:

- Building repositories that reflect/test new ideas,
- Continuing research into projects/innovations, and
- Facilitating dialog, knowledge dissemination.

There is a desire within the learning repository community to form communities that exchange information and expertise. But the issues that are researched and discussed must be issues that are of broad interests and that can lead to collaborative action. Those issues did not turn out to be the general challenges associated with collection and metadata management. But issues for future discussion and collaboration were specified as being specially suited to research and collaboration across repository projects and across user communities. These issues include:

- Interoperability between repositories,
- Construction of a single search space for multiple repositories,
- Manual and automatic metadata creation,

- Recommendation systems, and
- User profiles and their utilization.

The Academic ADL Co-Lab is working hard to facilitate the sharing of information and research on these subjects and to conduct and publish similar research itself. It is of great importance to the Co-Lab and to the learning repository community itself that our future steps recognize and focus upon those issues that can be addressed collectively. Visit <http://www.academiccolab.org/projects/repositories.html> for ongoing information and the link to a database of the repositories and papers mentioned in this document.

It is also important that we avoid discussing challenges to the solutions to which we can all come to individually without hindering the progress of learning repositories generally. There were many at the event who felt that to move forward, the learning repository community as a whole needed to adopt a new paradigm for thinking about any number of issues.

While paradigms are useful tools for formulating questions they are also constraints on the questions that will be asked, and there is no reason why we should push for the entire community to work under the same constraints.

As Paul Feyerabend has written, “The only principle that does not inhibit progress is: anything goes.” (Feyerabend 1993) It is simply too early in the game to advocate that learning repository projects pursue a single direction. The current wide distribution of interests and perspectives within the learning repository community is not to be lamented, indeed in the absence of institutional constraints or clear successes on the part of any one group, there are no pressures beyond weak suggestions for the learning repository community to achieve any level of agreement regarding these issues. We cannot demand uniformity of direction, but we can encourage members of the repository community to strike out in some direction, developing new strategies and techniques that better bring digital learning materials to users.

References

- CAREO 2003. *CAREO Repository Conceptual Architecture*.
<http://www.careo.org/documents/architecutre.html>
- Cedergren, Magnus 2003. "Open Content and Value Creation." *First Monday*.
http://firstmonday.org/issues/issue8_8/cedergren/
- Feyerabend, Paul 2003. *Against Method*. New York: Verso Books.
- Friesen, Norm and Lassi Nirhamo 2003. *Survey of LOM Implementations: Preliminary Report*. http://mdlet.jtc1sc36.org/doc/SC36_WG4_N0057.pdf
- Harrigan, Kevin 2003. *How Does the CLOE Repository Facilitate the Use of Learning Materials?*
http://www.academiccolab.org/events/oct78/CLOE_Presentation_Wisconsin.ppt
- Hendrick, Geoff 2003. *EdNA Online: A Services-Oriented Approach to Supporting E-Learning*.
http://www.academiccolab.org/events/oct78/CLOE_Presentation_Wisconsin.ppt
- Henry, Geneva 2003. *Connexions*. <http://www.academiccolab.org/events/oct78/Cnx-learning-repositories03.ppt>
- Levine, Alan, Brian Lamb, and D'Arcy Norman 2003. *Syndicating Learning Objects with RSS and Trackback*.
http://www/mcli.dist.maricopa.edu/show/merlot03/merlot_lo_rss_paper.pdf
- Lightle, Kimberly S. 2003. *Transforming Metadata: IEEE-DC-USMARC*.
http://www.academiccolab.org/events/oct78/LearningRepository_100703.ppt
- Muramatsu, Brandon and David Wiley 2001. *User Profile Information Draft v.3*.
http://www.smete.org/smete/public/about_smete/activities/technology/user_profiles/SMETE.ORG_User_Profile-v3b.doc
- Richards, Griff and Mark Hatala 2003. *POOL, POND and SPLASH: A Peer to Peer Architecture for Learning Object Repositories*.
<http://www.edusplash.net/upload/INET2-RichardsandHatalav2.doc>

Learning Repository Initiatives Explored

Apple Learning Interchange
ARIADNE -- European Knowledge Pool System
Blue Web'n
Campus Alberta Repository of Educational Materials (CAREO)
Canada's SchoolNet
Computer Science Teaching Center (CSTC)
Computing and Information Technology Interactive Digital Education Library (CITIDEL)
Connexions
Co-operative Learning Object Exchange (CLOE)
Digital Library for Earth System Education (DLESE)
Educanext
Education Network Australia (EdNA)
Eisenhower National Clearinghouse for Mathematics and Science Education
Enhanced and Evaluated Virtual Library (EEVL)
Exploratories
Fathom Knowledge Network Inc.
Gateway to Educational Materials (GEM)
The Harvey Project
Health Education Assets Library (HEAL)
Humbul Humanities Hub.
iLumina
The Learning Matrix
Learning Object Repository – University of Mauritius
Learning-Object.net
LearningLanguages.net
Maricopa Learning Exchange
The Math Forum
Multimedia Educational Resource for Learning and On-Line Teaching (MERLOT)
National Engineering Education Delivery System (NEEDS)
National Learning Network: Materials
OpenCourseWare
Problem-Based Learning Clearinghouse
Science Mathematics Engineering and Technology Education (SMETE) Digital Library
Scottish Staff Development Library
Telecampus
Wisconsin Online Resource Center