The Learning Marketplace

Meaning, Metadata and Content Syndication in the Learning Object Economy

Stephen Downes

Moncton, New Brunswick

January 21, 2004

Introduction 7
Nine Rules for Good Technology 10
Introduction 10
Good Technology: The List 10
Hungry Minds: A Commentary on Educational Portals 14
Content Syndication and Online Learning 21
Part One: Content Syndication 21
1. Channels and Channel Definitions 21
2. A Wee Bit of History 23
3. Syndication 24
4. Uses for Content Syndication 25
Part Two: Content Syndication and Online Learning 26
5. The MuniMall Project 26
6. Content Syndication 27
7. Link Syndication 29
8. Discussion Syndication 31
9. Why Things Didn't Work 31
10. A Success Story: The MuniMall Newsletter 33
11. All Together Now: Doing Educational Content Syndication Right 34
The Learning Marketplace – A Concept Paper 37
Overview 37
Why a Learning Marketplace? 38
Creating Course Offerings 39
Creating the Offer 41
Customization and Personalization 43
Presenting the Offer 46
Course Delivery and Completion 49
Third Parties 49
Cutting the Throat of the University 52
Introduction: Mensa and Academia 52
Eskow's Hypothesis 52
Why Defend Universities? 53
The Means of Production 54
A. Massification 55
B. Marketization 55

C. Managerialization 57
Language, Truth and Logic 58
A. Choice 59
B. Standards 60
C. Efficiency 60
Quality and Control 61
The Implications of Re-Form 64
Intellectual Wealth and Society 65
The Fragmentation of Learning 68
In Practice... 73
The Panel 73
The Issue 74
Epilogue 78
Education and Embodiment 80
1. The Experience of Cyberspace 80
2. Meaning and Experience 82
3. The Web and the World 85
4. Culture and Telepresence 86
5. Embodiment and Education 87
6. The Promise of Indirect Experience 89
Smart Learning Objects 91
Aggregators, Assimilators, Analysts and Advisors 95
Five Choices: Or, Why I Won’t Give Dave Pell Twelve Dollars 99
The Learning Object Economy 106
Abstract 107
A. The Need for and Nature of Learning Objects 107
i. The Idea of Learning Objects 108
ii. The Case for Online Learning 109
iii. The Cost of Online Learning 109
iv. The Argument for Learning Objects 111
v. Courses? No, Not Courses 111
vi. Sharing the Old Way 112
vii. Contemporary Online Sharing 114
viii. What We Need 115
B. Learning Objects from a Theoretical Perspective 116

ix. Course Construction and RAD 116
x. Object-Oriented Design 117
xi. Open Standards 118
xii. A Common Mark-up language 119
xiii. Common Program Interfaces 120
xiv. Standards and Standards Based Initiatives 121
xv. E-Learning Standards Based Initiatives 122
xvi. Classifications of Learning Objects 123
C. The Practical Application of Learning Objects 125
xvii. Creating Learning Object Content 125
xviii. Creating Learning Content 126
xix. Creating Metadata and Wrappers 127
xix. Creating Packages 128
xx. Delivering Courses 130
xxi. Course Format and Display 130
xxii. Pedagogical Issues 131
xxiii. Beyond Courses 132
D. The Learning Object Economy 132
xxiv. Course Portals 133
xxv. Course Packs 133
xxvi. Learning Object Repositories 134
xxvii. Certification and Review 135
xxviii. Publishers and Private Enterprise 135
xxix. The Learning Marketplace 137
The New Literacy 140
An Introduction to RSS for Educational Designers 143
Why RSS is Important for Educational Designers 143
The RSS Network Architecture 144
RSS Channels 145
Creating an RSS Channel 146
Weblogs 147
Aggregators 148
Headline Viewers 150
What RSS Does Not Provide 150

Componentization 152
Design and Reusability of Learning Objects in an Academic Context: A New Economy of Education? 155
Introduction 155
The State of the Art 155
Overview 155
Course Portals 155
Course Packs 156
Learning Object Repositories 157
Learning Content Management Systems 158
Problems and Issues 159
Overview 159
Proprietary Standards 159
Overly Strict Standards 160
Monolithic Solutions 161
Closed Marketplace 161
Disintermediation 162
Selective Semantics 162
Digital Rights Mismanagement 163
Design Principles 164
Overview 164
Open Standards 164
Royalty-free Standards 164
Enable, Don’t Require 165
Open-Source Infrastructure Layer 165
Open or Proprietary Service Layer 165
Component Based Architecture 165
Distributed Architecture 166
Open Access 166
Open Market 166
Standards Tolerance 166
Multiple Channels 166
Multi-Party Metadata 167

Integration with the Semantic Web 167
Multiple Data Types 167
Simple Digital Rights Management (DRM) 167
Brokered DRM 167
Permission Based 168
The Distributed Network 168
A Network Rather Than A (Single) System 168
Core Components of the Network 168
Contrast to Library Model 169
Component Registry Services 169
Functionality of the System versus Functionality of the Learning Resource 170
Secondary Components 170
Third Party Metadata 171
Digital Rights Management 171
Beyond Digital Rights Management: Employee/Consumer Rules 172
Learner / User Information System 172
The Lattecentric Ecosystem 175
1. 175
2. 177
3. 178
4. 180
5. 182
Copyright, Ethics and Theft 185
Preamble 185
1. 185
2. 186
3. 187
4. 188
5. 190
6. 191
7. 191
8. 192
9. 194
10. 194
Paying for Learning Objects in a Distributed Repository Model 196
1. How to ask for payment (and to specify use conditions) 196

2. How to actually make the payment 198
A. Presentation of a list of resources 199
C. Generation of a request for the learning resource 204
4. What Needs to Be Done 207
A. A Content Distribution Network 208
B. The DRM Substructure 208
The Aeffability of Knowledge Management 210
1. The Autism of Knowledge Management 210
2. Objects, Reusability and Universality 210
3. Applying Knowledge 212
4. Completeness and Liberation 213
5. Descriptions and Language 214
6. Deliverance 215
7. The Challenge Before Us 216
Public Policy, Research and Online Learning 220
Freedom in E-Learning 230
1. Discussion Topic - Theory for eLearning 230
2. Hypothesis 9 230
3. A couple of quick points... (in Response to Mark Nichols) 233
4. Some Notes 236
5. Freedom in E-Learning 238
The Regina Declaration 239
Closing Comments 239
Open Access Initiatives 240
The Regina Declaration 241
Design, Standards and Reusability 245
1. 245
2. 245
3. 246
4. 247
5. 249
6. 251
7. 252
Meaning, Use and Metadata 254
Resource Profiles 260
1. Introduction 260
1.1 Abstract 260

1.2 What is a Resource? 260
2. Describing Resources 261
2.1 Getting the Description Right 261
2.2 Multiple Descriptions 262
2.3 The Problem of Trust 263
3. Resource Profiles 264
3.1 Overview of the Concept 264
3.2 Vocabularies 265
3.3 Authorship 266
3.4 Distributed Metadata 268
3.5 Resource Identifiers 268
3.6 Models 269
3.7 The Concept in Retrospect 271
4. Types of Metadata 272
4.1 On Types of Metadata in General 272
4.2 Bibliographic Metadata 273
4.3 Technical Metadata 274
4.4 Classification Metadata 274
4.5 Evaluative Metadata 275
4.6 Educational Metadata 276
4.7 Sequencing and Relational Metadata 277
4.8 Interaction Metadata 278
4.9 Rights Metadata 279
4.9 Metadata Types in Retrospect 279
5. Using Resource Profiles 280
5.1 The Lifecycle of the Learning Resource 280
5.2 Generating Resource Profiles 281
5.3 The Metadata Distribution Network 281
5.4 Projected Metadata 283
5.5 Data Network Properties 284
5.6 Interoperability 286
6. Concluding Remarks 287
6.1 The Future of Metadata 287

6.2 The Intelligent Network 287

My Canada 297

1

2 Introduction

This is not a real book. If this were a real book, I would have to eliminate the redundancies, correct awkward expressions, employ a consistent chapter scheme (including the provision of names for all the chapters), and locate (or at least explain the disappearance of) a certain section B in one of the papers.

If this were a real book, I would have to make it interesting. And rest assured, unless you are deeply interested in such topics as learning objects, metadata and content syndication, this will be to you one of the least interesting books of the year. For myself, I find these issues fascinating, and the issues surrounding them pressing. But I do not expect my enthusiasm to carry.

So why does this book exist at all? This is actually two questions: why does the book exist now, as a book, when it did not exist yesterday? And why did I engage in what amounts to a discourse on these topics in the first place?

The second question is easy. This afternoon I received an email from someone announcing a number of 'discoveries' consisting of some of the works contained below. And while on the one hand I was pleased that he had found the (to me) important works, on the other hand I was disappointed that he had missed them for so long, and still missed some crucial papers. But in the end, this is my fault, as my work is scattered across the four corners of the web. This book is something like an attempt to create an authoritative account of the last three years.

Which brings me to the second question: why bury one's mind in such arcane matters at all?

In the first months of 2001 I accepted a research fellowship at the University of Melbourne where philosophy professor Tim van Gelder was attempting, like many professors, to put his course online. During the course of this work I explored the idea of making his courses available through learning management systems. On finding one roadblock after another, I concluded that the vision outlined in some of my earlier work was being hampered by a lack of vision of what could be realized by online learning.

In Australia, I wrote a paper that may be found below, The Learning Marketplace, in which I outlined the mechanics of a content syndication system. It drew on some of the concepts outlined in Learning Objects and Content Syndication and Online Learning. In bits and pieces over the next three years, I continued to fill in the details of the model. This book does not contain everything I wrote during that time – it is devoted solely to the topics of (as the title suggests) meaning, metadata and content syndication in online learning.

That's the beginning, then. Most of these papers are part of that work in progress, in a sense, drafts that build on each other (which is why you will find repetition (which in a real book I would have to excise, thus removing the essence of discovery). In addition to the straightforward mechanics of content distribution (which you would think is straightforward) I encounter issues as varied as ontology, legal policy, the open source credo, morality, semantics, and technical design.

So in a sense, I think that this book is much more than merely a technical text on how to create a content syndication network. It is at the same time an extended discourse on how we, as a society, ought to create and acquire knowledge, which in turn becomes a discourse on what constitutes knowledge, how we related to it, and how we talk about it. An essay like The New Literacy may seem as far removed from a technical treatise as you may get, but unless you create the analogy of

multimedia objects as words in a new language, many of the concepts seem muddled and disjointed. Allow us to communicate with learning objects, however, and the flow from Education and Embodiment to The Aeffability of Knowledge Management to The Lattecentric Ecosystem to Design and Reusability of Learning Objects becomes clear.

Well, mostly clear.

Anyhow, to preserve the sense of discovery and of wrestling with the issues, the papers are presented in mostly chronological order. I say 'mostly' because I have made some allowances for clarity of exposition and elucidation of concepts before their use. But not much. It is important to me that readers be able to follow the logic of discovery as well as that of the linear syllogism and the intermingling of threads.

But all of that said, I want to stree that there is a single, unified theory underlying all of my work, a theory that, because it can't be summarized in nine steps or a neat taxonomy, is perhaps a but hard to grasp, but is nonetheless as powerful, as expressive, and in my view, as correct as any other approach to the discipline. If I had to give it a name, I would call it 'network learning' (though that name has already been taken). In any case, it is very difficult to see the strands of the theory, much less the structural and methodological consistency between strands, without being able to view my work here as a single entity.

What I present here is as much an epistemological, moral and political theory as it is a theory of learning. It is based in philosophy of mind, the logic of connectionism, post-foundational epistemology, and empirical underdeterminism. It is as much a story about how we, in fact, construct our reality as it is a story about the nature of that reality (but though it shares affinities, it is not construtivism and should not be represented as such).

Think of this theory as the rejection of the unified field or the idea that there is one, knowable, fundamental law of everything. Think of it as the recognition that perception, reality, knowledge, learning, software and society come in discrete chunks – epistemological quanta, if you will – and that what we do when we create schools or civilizations is a function of how these chunks are organized. Think of there being layers of these chunks, such that each subsequent layer is composed out of entities that are emergent properties of the preceeding layer, and that the nature of these emergent properties depends in part on how the chunks are organized (like the pixels on a television screen) and in part on their perception at the higher layer (I say 'layers' but it would be more accurate to think of it as an n-dimensional ordering of reality, not a neat linear progression of layers).

On such a model, the idea of universal theory – in anything – is not only misguided, it is wrong, wrong empirically, since it results in false theories, and wrong ethically, since it seeks to impose order from one layer, where an ethical theory may apply, to another, where it makes no sense. One might say, for example, that the figure depicted in the telvision image is evil, but it would make no sense to say, therefore, that the pixels that comprise that picture are themselves evil. The goodness of a pixel has nothing to do with the goodness of an image, which in turn, has nothing to do with the goodness of a society.

If we cannot, then, employ the tradition of observation and generalization, definition and prescription, causal reasoning and effect analysis, then how should we approach these disparate phenomenal? That's what this book is about. Yes, true,

it is about the design of a knowledge and information network, but it is inherent in the design that the rules (so to speak) for such discourse emerge. The best I can do here is to point to instances and hope that you, the perceiver, 'see' the 'picture' in your head that corresponds with the system I am describing – understanding, always, that your picture will be different from mine.

Stephen Downes

Moncton, New Brunswick, Canada

January 21, 2003

This article began as a in DEOS as a response to some of the complaints about educational technology being posted. It captures some of my intuitions about technology and development and tries to express the idea that people need to choose their technology. Jim Morrison asked me to write the first paragraphs to adapt it for publication in Technology Source.

3 Nine Rules for Good Technology

Written March 7, 2000. Published in The Technology Source, July/August, 2000. Reprinted as "Mmm, God Technology", University Business, September, 2000. The CyberUnion Handbook: Transforming Labor through Computer Technology, by Arthur B. Shostak, et al., M.E. Sharpe, 2001.

3.1 Introduction

Today's educational technology is like a Rube Goldberg contraption. Enter any technology-enabled classroom or other facility, and you will see a mish-mash of computers with associated wires, video displays, modems, ITV, CD-ROM libraries, tapes, and more. To use this technology effectively and avoid being distracted by the usual malfunctions and dense manuals, teachers must spend a lot of time in the classroom themselves.

It doesn't have to be this way, however. As technologies mature, they tend to become easier to use. Consider the elevator and radio, for example. Once so finicky it needed operators to take riders from floor to floor, today's elevator functions flawlessly with little intervention on the part of users. Likewise, when the radio was first developed, it was the domain of specialists. Today's radio is a model of usability, requiring no special training for the listener who wants to find the nation's top ten hits.

It is true that not all technologies are so uncomplicated. For example, the person who operates a nuclear reactor must have some expertise and special training. But such systems are rare, overwhelmed by an array of far simpler innovations. If a technology is to become widespread, it is crucial that it be easy to use–so easy that it need not be packaged with an operating manual. Technology that teachers employ in the classroom must be of exactly that variety: widespread and easy-to-operate. A learning simulation, a conferencing tool, and a student record keeper should be as untroublesome to use as a television, a telephone, and a notebook.

I believe that we currently are in a transition phase; we are moving away from complicated technologies toward simpler innovations. For the most part, however, today's technology remains clumsy. We must question whether the time and money we are investing in that technology–in teaching teachers to use it–is well spent. Certainly training is necessary to get us to a higher level of technological advancement, but we must not take our eyes off the long term goal: good technology.

What distinguishes a good technology from a bad technology? The following nine characteristics define the former. Think of them as a checklist; a technology that has more of these features is, in general, better than a technology which has fewer of them.

3.2 Good Technology: The List

Good technology is always available. This distinction is what makes buses, in spite of all of their advantages, bad technology. People cannot count on catching the bus at absolutely any time of day; thus most people prefer cars. In the educational field, the technological equivalent of the bus is the equipment

trolley. It is necessary because only one projector (or workstation or overhead projector) is available to serve five classrooms. Imagine what life would be like if we had to schedule our use of the elevator. Or to make reservations to use the telephone. Good technology does not require scheduling, relocation, or set-up.

The availability requirement raises cost considerations. Equipment that costs less is more likely to be available. But cost is not the sole or even primary determinant. If a technology meets the other criteria described below, it will be made widely available despite the cost. Think of ATMs, electrical lights, and highways.

Good technology is always on (or can be turned on with a one-stroke command or, better yet, starts automatically when the need for it arises). One thing that makes the telephone useful is that we do not need to boot up the operating system before we make a call. Likewise, electrical lights are a significant improvement over systems that required individual ignition with a match or candle, and streetlights are practical because they come on when it gets dark outside. A weakness of motor vehicles is that they are not always on, a fact that causes endless frustration for users needing transportation on cold winter days.

Much of today's educational technology requires long and sometimes cumbersome initialization procedures. After wheeling in a projector from another room, for example, three teachers and a technician may spend time plugging it in, turning it on, spooling the film, and positioning the screen.

Admittedly, the "always on" requirement raises significant energy consumption considerations. A portable device that consumes a lot of energy, for example, cannot always be on because it must carry its own power supply. Energy itself-in inefficient forms like gas and oil-is too expensive to be consumed merely for convenience. Devices with low energy consumption, however, can always be on. Think of watches, telephones, and elevators.

Good technology is always connected. Good technology can send information when and where it is needed without human intervention. Fire alarms, especially institutional ones, are useful in this way. Indeed, if the detectors were not connected to warning systems, the alarms would be useless. Again, telephones are useful because no procedure is required to connect to the telephone system.

As recently as last month, I spent fifteen minutes in a room with a dozen or so highly paid professionals waiting for an ITV system to be connected to a remote location. I have spent much time listening to my modem dial up a local provider (and luxuriate today in the convenience of an always-on Digital Subscriber Line connection).

Good technology is standardized. One television functions much like another television (televisions became less useful with the introduction of brand-specific remotes). One telephone connects to any other telephone in the world. One brand of gasoline powers your car as well as any other-but cars that require different grades of fuel, such as diesel, are bad technology because of their reliance on non-standard fuel.

Standardization promotes interoperability. Interoperability means that you have choices, that you are not locked into one supplier or vendor. It means that you can adapt easily to improved versions of the same technology: you can upgrade to a bigger television or engine-cleaning gasoline without replacing your electrical wiring or car engine. A video that is designed to be played only on a specific computer platform and email that may be read only via a specific Internet Service

Provider are examples of bad technology. Video should be viewable on all platforms and email should be accessible through any Internet service provider.

Good technology is simple. Simplicity is a slippery concept, but the best technologies can be learned by looking at the input device, not by studying a manual.

Here's how I distinguish between good computer programs and bad computer programs: I try to install and run the program without the use of any manual. Installation is much easier today, thanks to a good computer program called "Setup." Running the program is a different matter. When I have to stop and think (and read very small print) about how to get rid of a paperclip icon so that I can type a letter, I know I am dealing with bad technology. Good technology, by contrast, is intuitive. To use an elevator, I press the floor number. Simple. To make a phone call, I dial the number. Easy.

Simplicity goes hand-in-hand with range of function. Features that you never use get in the way, and they make the product complicated and cumbersome. Look for technology that does exactly what you want: no more, no less.

Good technology does not require parts. Cars are bad technology: they require a never-ending array of parts, from gasoline to oil to air filters. It is easy to overlook parts because they seem integrated into the whole; consumables, like oil or ink cartridges, don't satisfy our intuitive definition of parts. But insofar as they must be replaced and are essential to the operation of technology, they count as parts, at least for the purposes of this article.

The bottom line is this: Do you have to purchase something on a regular basis in order to use your technology? Do you have to replace something that becomes worn out or depleted or that can be lost or stolen? The fewer times you have to purchase or replace, the better your technology; the best technology requires no ongoing purchases or replacements at all.

Sometimes it is not possible to do without parts, but this is a sign of a transitional technology. Perhaps even good technologies, such as portable stereos that require CD-ROMs, need parts. But a portable stereo that does not need CD-ROMs because it can download MP-3s from the Internet instead would be better. If parts are absolutely necessary, they should be widely available, standardized, and simple to install. DVD players, for example, will not qualify as good technologies until DVDs become as widely available as videotapes.

Good technology is personalized. Some of the simplest technologies succeed because they are personalized. One of the things that makes a telephone useful is that you have your own telephone number. In a similar manner, e-mail is useful because you have your own e-mail address. ATM cards would not be at all useful unless they opened your bank account and only your bank account. Credit cards, smart cards, pagers, cell phones, and eyeglasses are more examples of personalized technologies.

Bad technology forces you to fit its requirements. I purchased my copy of Microsoft Word in Canada, but the default dictionary was for American English. I could install a British dictionary, but Canadian English is distinct from both British and American English. Like many users, I am forced to add each distinctly Canadian word to a custom dictionary. This is bad technology. Why can't I simply tell Word that I am Canadian (or an architect, or a member of some other specialized group) and have it retrieve the appropriate spellings for me?

Good technology is modular. By "modular" I mean composed of distinct entities,

each of which works independently of the others and may be arranged or rearranged into a desired configuration with a minimum of fuss and effort. To a degree, this requirement is a combination of the requirements that good technology be standardized and personalized, but modularity takes technology a step beyond either of those features.

Bricks and wood are good technology because they interconnect neatly and can be assembled into custom configurations. Legos are even better because they do not require parts like nails or cement (which is why Lego, and not Mecanno, is the construction toy of choice).

The stereo systems we purchased in the 1970s are good examples of modular technology. Using the standardized RCA jack, we could assemble systems with or without pre-amps, tuners, equalizers, or even turntables. Today's Universal Serial Bus (USB) represents good technology because it allows computer systems to be assembled like the stereos of old. Books-and paper in general-are good because they are modular; a person may assemble a book, such as a binder, out of individual sheets of paper and a library out of a collection of books.

Good technology does what you want it to do. And it doesn't do something else. "Doing what you want it to do" means the same thing as "idiot proof." Good technology minimizes the potential for operator error and thus the possibility of unexpected consequences. Good technology is also robust-less prone to breakdowns and malfunctions-and reliable. Software that crashes instead of running is obviously bad technology. Telephone systems that connect you to India instead of Indiana are not useful.

"Doing what you want it to do" is a highly personal thing. If you want your daughter's clothes to protect her from the cold, then her selection of a light chiffon top and an ultra-mini skirt represents bad technology. But if she wants clothes to accentuate her physical features, then the same clothes represent good technology.

Conclusion

It is important to remember that no technology is perfect. No technology will satisfy all nine rules. However, some technologies will satisfy more rules than others, and some technologies will even break a rule or two and still be very good technologies (if only because no better alternative is available). That said, purchasers should insist on-and vendors should be pressed for-good technology as defined above. We spend too much time and money on new technology to be satisfied with anything less.

This article, like the last, was written before my trip to Australia and is included here in order to set the theoretical stage. While Nine Rules looks at technological design, this article looks at market design. These two concepts weave through this book, the one being of course dependent on the other. While on the one hand I am trying to suggest that portals will not become the destination of choice for internet users (a prediction that has now come to pass), I am also trying to suggest that branded institution-specific marketing of learning materials will likewise fail. People – especially those trying to promote a specific institution – still have a hard time wrapping their head around this idea. They want to create an environment where choices are limited to the offerings of single (or select sets of) providers.It won't work, not because it can't be done, but because people don't want it.

4 Hungry Minds: A Commentary on Educational Portals

Written November 15, 1999. Published in Online Journal of Distance Learning Administration, Volume III, Number I, Spring 2000.

The front end of Hungry Minds (http://www.hungryminds.com/) is an education portal modeled along the lines of another portal, About.Com (http://www.about.com; originally called The Mining Company). Hungry Minds' experts author a topical home page with commentary and links to resources and especially online courses.

In taking this approach, Hungry Minds is pursuing a path that intuitively seems correct – the idea that people searching for online learning opportunities will follow a topic-based pattern, and not an institutionally based pattern. For example, a person wishing to take a course in Roman History will search for 'Roman History'; they will not instead check out the University of Alberta or the University of California.

This is not entirely the case, of course. All universities have an established client base, represented by existing and former students with a stake in that particular university's credentials. And to some degree, the cachet of some universities, such as Harvard or Princeton, will draw additional students. But this is a client base on the wane; as continuing education especially rises to the fore, students will first search for the topic area, and only then consider the name and reputation of the institution offering the courses.

I belabor this point because it has been the practice of most traditional educational institutions to place most of their efforts into the creation of institutional, and not topic-based, portals. That is to say, most institutions list only their own course offerings, to the exclusion of other institutions. They rely on people using standard search engines to locate the course, though of course the portal is structured as though someone would first look up 'The University of Alberta' and then peruse the course offerings (indeed, the University of Alberta menu further requires that you proceed to individual faculty pages before you see any course listings).

I'll call this the 'restaurant' model of online course offerings. Like restaurants, traditional institutions are appealing to their name and reputation (and the occasional review). The only 'brand' present in a restaurant is the restaurants' own; the only choices offered are from the restaurant's own menu. Production and consumption is localized. Advertising promotes above all else the restaurant's name and distinctive quality.

It is hard to over-emphasize this point. For even where some various institutions have formed coalitions, the tendency has been to favor the individual institution

over the breadth of content and expertise. California Virtual University (CVU) is a classic example of this (Downes, 1999). Even though a consortium was formed, each institution clung rigidly to its own identity and methodology, even to the point of individualized course numbering systems.

Now online learning portals have existed for some time – I can off the top of my head list TeleEducation and the WWWDEV course database at the University of New Brunswick Even in such portals, courses are listed by topic. The name of the institution appears only as an attachment to individual courses.

What is new about Hungry Minds – and about UNext, a similar service recently featured in Wired News – is that these agencies are acting as online course brokers. Rather than merely listing online courses, they are acting in some capacity as a representative for both the course vendors and the potential students.

Thus, for example, Hungry Minds offers prospective learners a guarantee that offers a refund on any online course they take. And UNext, while it offers no financial guarantees, is focusing on the quality of its course offerings (it lists three Nobel laureates at the top of its academic advisory board).

These institutions are taking what I will call the 'grocery store' model of online education. They act as a distributor for brand-name products in such a way that the store – not the producers – manages purchases and refunds, and the store – more so than the producers – stands as the agent that ensures quality and price.

Now the easy approach at this point would be to argue that the grocery store model is on the ascendant, and the restaurant model is on the decline, because people prefer choice and selection, and because they prefer the guarantees and single interface that grocery store providers offer them. And to a significant degree, this will be what happens – while restaurant vendors may not see a decline in customers (they have a locked-in client base, after all), they will not participate in the enrollment boom online learning will engender.

More and more, restaurant vendors will focus on service and quality of offerings (some, such as ZD University (now renamed SmartPlanet as part of Ziff-Davis's new online learning initiative) will focus on price and availability). This is the main thrust behind, say, Michael Cenkner's remarks to ATLNet – "One idea that's been batted about in the Faculty of Extension is to provide a cluster of services for new students, in this case, foreign students (Cenker,1999)."

It has been remarked in the past that educational institutions will have to shift from being repositories of knowledge toward becoming service-oriented agencies (Downes, 2000) and this remains true. Such services, localized within an institution (even a large one, such as the University of Alberta), are costly, however. In order to keep costs down, even restaurant-type vendors will be drawn inevitably toward the grocery store model.

In the grocery store world, an inevitable battle is taking shape as different grocery stores try to establish themselves as the exclusive – or at least, the primary – broker for online courses. Hungry Minds and UNext represent two poles in this battle. Hungry Minds is focusing on expertise and service guarantees. UNext is focusing on expertise and exclusive offerings. Each is trying to segment the market, offering courses nobody else offers. It is as though Safeway – and only Safeway – sold Heinz products, while Save-On was the exclusive dealer of Kraft products.

But in the end, neither the restaurant nor the grocery store will be the primary

agent of online learning.

The machinations emerging in the online learning community mirror the machinations that occurred when previous monopoly services, such as long distance telephony or cable television, were opened to competition. A raft of competing vendors emerged, spending a pile of money on name and brand recognition (the recent press campaigns supporting Hungry Minds and UNext are instances of this). They began by focusing on quality (as in Sprint's 1-800-PIN-DROP campaign) and choice (as in the satellite-TV's 800 channels campaign).

But consumers were unable to find the suggested difference in quality - a long distance call is essentially the same no matter who provides it; FOX is FOX whether delivered by satellite or by cable. They next began to focus on price - but since the price of these services is essentially the same, various Byzantine pricing schemes emerged to obscure the difference. Expect a similar price war in the field of online learning, bolstered by commentaries complaining about the high cost of online learning.

But with quality and price being essentially non-factors, and with brand recognition able to carry only a small percentage of online learning institutions (who will charge a premium for this, thus maintaining their exclusivity), there will be nothing to choose from between Hungry Minds, UNext, and the many similar services which will operate across the world wide web.

So in planning for the future, education providers - both course-delivering institutions and aspiring portals - will have to look hard at what actually motivates the purchase of food, long distance telephony, cable television services, and any other commodity. And that factor (which also motivates love, marriage, crime and corruption) is proximity.

Think about it. Where do you buy your groceries? Do you drive across town in order to get the superior quality offered by the west-end Loblaws? Probably not; you probably buy your groceries within a kilometer of your home. Which restaurants and pubs do you patronize? I am a regular at the Inglewood, which happens to be two blocks from my front door. My long distance is provided by the local phone company, my cable from the local cable company, and if I were to commit a crime, it would probably be in my own neighborhood.

The same is true - has historically been true - of education. It is no coincidence that most students at the University of Alberta are from Alberta. While there is more mobility in education than in - say - restaurant selection, and while some name-brand institutions can attract scholars from around the world, in the main, people eat, sleep, learn and love where they live.

But what constitutes proximity on the world wide web? One truism is that the web breaks down physical distance. Once, people fell in love with and married people they met locally (and this is still how the vast majority of couples do it). But increasingly, couples are meeting and marrying online (http://www.webring.org/cgi-bin/webring?ring=netmate;list). In areas where physical proximity was paramount, the internet is breaking down that barrier and uniting people from around the globe.

Yet - even in the area of online romance - proximity is still vital. Browse through the hundreds of personal pages describing online romances and you will find (Alta Vista Search) in every one of them a reference to a particular MUD, chat line, IRC channel, discussion board, or other online forum.

Proximity on the internet falls under the loosely defined category of 'online

community'. Though only recently discovered by mainstream academics and corporate pundits, the proliferation of online communities is what has *always* defined the internet. In the early days, netizens populated particular MUDs, IRC channels or newsgroups. Today, people congregate around portals, mailing lists, discussion boards and chat rooms.

There is some research that reveals this pattern in web usage. Tauscher and Greenberg (1997) report, for example, that "People tend to revisit pages just visited, access only a few pages frequently, browse in very small clusters of related pages, and generate only short sequences of repeated URL paths." In other words, people find the sites they like and tend to stay with them.

While a variety of factors influence a person's choice of websites (for example, people will leave sites which are too slow), the primary determinant is interest. The site discusses some topic that is important to that particular person. Indeed, a person's interests may be deduced from the sites they frequent – one person may visits news, gardening, astrology and self-help sites, for example, while I frequent news, technology and education sites.

It stands to reason – though I have no statistics to support this because the practice is not yet widespread – that people who take online courses will take those courses listed on the sites they most frequently visit. If, for example, I wanted to take a course on XML, I would be far more likely to take such a course offered from one of my regular haunts than I would to search for XML courses in general. And the idea of searching a particular institution – say, the University of Alberta – for XML courses would not even show up on my horizon.

There are some strong caveats to this, of course. I would have to be sure that the course was offered by a reputable institution and taught by people knowledgeable in the field. I would have to be convinced that they would not merely take my cheque and disappear. It would have to be offered at a reasonable price, and at a place and time convenient to me. But these are all factors that emerge after the initial course selection has been made – factors which influence whether or not I select a particular course, and not how I begin my search (if I search at all) for a course to take.

So now – Terry Anderson observes (1997) that "There is huge 'land rush' now in progress between third party portals, seeking to combine and generate courses and student services from many institutions vs. schools who are working in-house trying to build delivery via systems such as WebCT and Blackboard and to integrate these services with registration, student support etc., providing students with a customized view of 'their institution'. In both instances the goal is an integrated 'one stop shopping' approach to life long learning."

Quite right – but if he asks, "What should the University of Alberta portal look like," he is traveling down the wrong road. If he is asking, even, whether the University of Alberta should team up with one of these grocery store portals, offering exclusive access in exchange for brokering services, he is still traveling down the wrong road. While both an institutional portal and a commercial portal will offer some short-term success, neither is likely to be the dominant model for online course delivery in the long run.

The conceptual leap that must be taken – which will be taken first by potential students, and only later by established institutions – is that the traditional gap that exists between learning and practice must be transcended.

Today, education exists in one sphere – in schools, colleges and universities –

while work and play exist in another sphere – in the workplace, job site, or the home. When we decide to learn, we stop our other activities, remove ourselves to some distinct place (and often at a preset time), and for a certain period of time, dedicate ourselves solely to learning. Our knowledge of learning opportunities – courses, programs, and resources – is distinct from our knowledge of work-related or play activities.

But work and learning (especially) and play and learning (to some degree) will converge online. The same site we use to chat with people who share our interests will be the site where we find our research materials, our examples of best practices, and our online courses and programs. We are likely to drift toward a site devoted to – say – gardening, there to chat with our online friends about roses, to look up fertilizer mixes for tulips, to buy seeds, and to take that course in hydroponics.

Such online communities – today misleadingly called "vertical portals" – are on the rise. They will focus on particular topic-based niches. Some will cater almost exclusively to a corporate environment, while others will cater to a person's general interests. In many cases, the two will combine – some people study the history of the Roman Empire professionally, while others merely find it an engaging hobby (and yes, I am a Roman Empire buff).

Traditional educational institutions need to do two things. First, they need to devise mechanisms that will enable their courses to be embedded in the offerings of a vertical portal. And second, they need to study the mechanics of vertical portals to best understand how learning could even fit into such a context. It is not clear that they should actually build such portals (there will be endless complaints that they are reaching beyond their mandate if they take up such activities as selling seeds), but they should place themselves in a position where they may partner with established government and non-government partners.

A more complete metric of exactly how the traditional institution should position itself is probably beyond the scope of this diatribe. But a few observations are in order.

With respect to the development of online learning materials and support systems, the institution must:

learn how to deliver materials in a distributed environment, where the primary point of interaction is *not* the university site

learn how to develop and deliver learning materials 'on-demand'

learn how to produce customized or tailored learning programs for particular corporate or individual clients

learn how to provide a completely online learning experience (this includes the such things as registry services, books and other resources, testing and grading)

learn how to promote the authority and trustworthiness of its online course offerings

With respect to the development of sector specific online communities, the institution must:

develop a framework for partnering with non-institutional partners

learn how to develop sector-specific resource sites in general

and how to pay for them without offending their partners

learn how to structure resource and learning databases so that materials and

courses are available on-demand

learn how to partner with other educational institutions offering courses and programs in the same field (including credit transfers, common registration, etc)

These are just a sampling of the mechanisms required to support sector-specific online learning. No doubt a wide variety of technical, administrative and political issues will emerge in practice. In my own experience - trying to develop a sector-specific learning environment in municipal affairs http://www.munimall.net - all of these issues and more have arisen. And as we build this community in Alberta, many more issues - unanticipated issues - will arise.

References

Distance Education Courses at the U. of A. Web Page. Available http://www.atl.ualberta.ca/disted/distlist.cfm. Viewed 11 January, 2000.

Downes, Stephen. What Happened at California Virtual University. 14 April, 1999. Available http://www.atl.ualberta.ca/downes/threads/column041499.htm. Viewed 11 January, 2000.

TeleEducation New Brunswick. Website. Available http://tenb.mta.ca/. Viewed 11 January, 2000.

WWWDEV Members' Courses. Website. Available http://www.unb.ca/wwwdev/c3.html. Viewed 11 January, 2000.

Unext.Com. Website. Available http://www.unext.com/. Viewed 11 January, 2000.

Dean, Katie and Kendar Mayfield. A Top-Drawer Education Online. Wired News, 12 November, 1999. Available http://www.wired.com/news/culture/0,1284,32474,00.html. Viewed 11 January, 2000.

Course Guarantee. Hungry Minds. Available http://www.hungryminds.com/aboutus/fl/help_courseguarantee.html. Viewed 11 January, 2000.

SmartPlanet.Com. Website. Available http://www.smartplanet.com/sp_home.asp. Viewed 11 January, 2000.

Cenkner, Michael. Distance Education Portal. Email to ATL-Net, 12 November, 1999.

Downes, Stephen. The Future of Online Learning: Accreditation. Available http://www.atl.ualberta.ca/downes/future/accreditation.htm. Viewed 11 January, 2000.

I Met My Mate on the Web. Website. Available http://www.webring.org/cgi-bin/webring?ring=netmate;list. Viewed 11 January, 2000.

Alta Vista Search Results for "We met on the internet". Available http://www.altavista.com/cgi-bin/query?q=%22We+met+on+the+internet%22. Viewed 11 January, 2000.

Tauscher, Linda and Saul Greenberg. How people revisit web pages: empirical findings and implications for the design of history systems. Academic Press, 1997. Available http://ijhcs.open.ac.uk/tauscher/. Viewed 11 January, 2000.

Anderson, Terry. Online Learning Portal. Email to ATLNet, 12 November, 1999.

I include finally this third paper to conclude the setting of the stage for this volume. It is my first real use of RSS in a paper and expresses what I was trying to do with munimall.ca (which continues to run to this day, albeit without the content syndication). These thoughts were fresh, very fresh, in my mind as I headed Down Under. In one sense, this paper is a report of the failure of MuniMall to meet my expectations, a failure that is analyzed in the work below. But in another sense, it is an explicit recognition that some things worked, and worked very well, giving me sufficient empirical validation to continue. What's important about this paper – and about RSS in general – is that it gets at the relations between individual entities. It is a systemdescribed here as a means of syndicating them, and in so doing, is a mechanism for organizing them.

5 Content Syndication and Online Learning

Written (probably the night before) and presented at NAWeb, October, 2000. Published in USDLA Journal, November, 2000.

This paper divides into two parts. In the first part it defines and describes the RSS (Rich Site Summary) format and its emerging use as a format for content syndication by news and media organizations on the World Wide Web. Through the use of working models and demonstrations, the development, display and distribution of content modules via RSS will be discussed. In the second part, the theories and practice employed by news and media organizations are applied to online learning. Using MuniMall, an online learning community developed by the author, as an example, the method of integrating syndicated content with online courses and learning materials will be described and illustrated.

5.1 Part One: Content Syndication

5.1.1 1. Channels and Channel Definitions

If you surf the web using a Netscape browser and followed the 'My Netscape' button to its logical conclusion, you will have encountered a description of something called RSS, or "Rich Site Summary." An RSS file allows a website publisher to produce a on Netscape's site; Netscape users, in turn, may select your channel as one of several channels on their 'My Netscape' page

A channel, typically, looks like this :

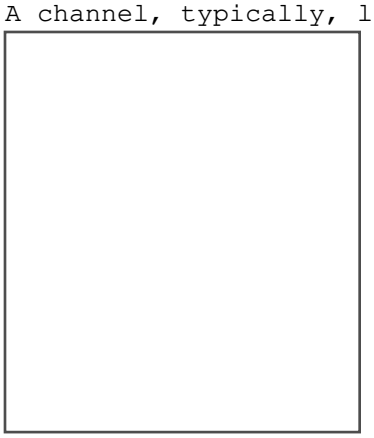

Figure 1: Netscape RSS Channel

The idea of a channel is that it is a brief summary of a website or online publication. It is composed of a channel name, a logo, and a set of headlines listing items on the site. Each headline points to a different article or column and may be supplemented with a brief description of its contents.

So far so good, and when Netscape launched its service early last year I was quickly on board with an RSS file of my own. It was a frustrating experience: Netscape's validation didn't work properly and I found myself re-registering over and over with the site's somewhat slow interface. Eventually the wrinkles were smoothed and my Rich Site Summary was accepted into Netscape's interface. Here's an abbreviated version of what it looks like:

```
<?xml version="1.0"?>
<rdf:RDF
   xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
   xmlns="http://my.netscape.com/rdf/simple/0.9/">
   <channel>
       <title>Stephen's Web Threads</title>
       <link>http://www.munimall.net/scripts/downes/clist/clist.cgi</link>
       <description>Stephen's Web Threads</description>
   </channel>
   <item>
      <title>Distance Education vs. Traditional</title>
      <link>http://www.munimall.net/ /topiclist.cgi?topicid=969550119</link>
      <description>
Does assigning distance students more work make up for the lack of classroom
contact? Well, no.
      </description>
   </item>
   <item>
      <title>Interview with Presidential Candidate Jackie Strike</title>
      <link>http://www.munimall.net/ /topiclist.cgi?topicid=969464710</link>
      <description>
The one-on-one chat with the talking 3-D candidate sets not only a
political precedent, but is a technological first.
      </description>
   </item>
</rdf:RDF>
```

Figure 2: RSS File

As you can see from the diagram, there are two main elements to an RSS file: the channel definition, and the item definition.

A channel is a set of related items. Items are descriptions of individual articles. A channel may describe items from a single website or items which discuss a particular topic. Items in turn may be anything at all, though typically they are a particular essay, news item, column, or similar chunk of content.

Channels and items each have properties. In the example above, a channel will have a title, a link or URL, and a description. Channels frequently have images

associated with them, may be provided by a publisher or website, and may have keyword descriptors. In a similar manner, items also have properties: a title, a link, a description, and perhaps some keywords, author and publisher information.

The idea here is that an RSS file is a structural description of a website or a group of related websites. Because the information is structured, when it is retrieved by a remote service – such as Netscape's NetCenter – it can be manipulated, displayed in various templates, and made the subject of intelligent searches. But more importantly, for the author of the RSS file, it allows content to be created and published once and distributed and viewed on many different websites. This is the heart of the concept behind RSS and of content syndication generally.

5.1.2 2. A Wee Bit of History

Where there is Netscape there is always Microsoft, and it should be no surprise to the reader that the Redmond software company developed an alternative channel format. The Microsoft format is called 'Channel Definition Format' and was introduced in 1997 for its Internet Explorer 4.0 web browser. The specifications were described in the November, 1997, issue of Microsoft Interactive Developer and a software development kit was released.

The idea behind Microsoft's 'Active Channels' was that website summaries could be displayed in the browser itself via a 'channel bar.' For some reason, Microsoft abandoned this feature in its release of Internet Explorer 5.0 thinking, perhaps, that it might incorporate it later as part of the Windows desktop. Ironically, a Netscape version of the channel bar was one of the major features added to the Netscape 6.0 release in April of this 2000.

Both the Microsoft and Netscape initiatives centered around a set of protocols described by the World Wide Web Consortium as RDF, or Resource Description Framework. The purpose of RDF was to provide a generalized format for online resources; major implementations thus far have included the Dublin Core for publications and the IMS Protocols for instructional materials.

But RSS channels need not be defined in an RDF format. Dave Winer's Scripting News, for example, adopted a non-RDF version of RSS. Started in December of 1997, the Scripting News Format, as it was called then, was launched to introduce the use of XML to news pages. By June of 2000, the Scripting News format had evolved into something called RSS 0.91 – which should not be confused with Netscape's RSS for while Netscape's 'RSS' stands for 'Rich Site Summary', Winer asserts that that there is "no consensus on what RSS stands for, so it's not an acronym, it's a name" .

Finally, in August, 2000 (which, by the way, explains why my paper is late), a group of developers adapted the best of RSS 0.91 and re-adopted the RDF format, producing the widely accepted RSS 1.0 specification. This design allows content developers to design and employ "RSS modules" in their RSS files, thus greatly increasing the potential vocabulary and use of RSS files. Content designers can now include, for example, threading, referencing, categorization, and more to the core RSS data set.

5.1.3 3. Syndication

The purpose of creating RSS files is to allow for the syndication of news content. Syndication on the world wide web works in much the same way syndication works in the world of print and electronic journalist: somebody writes a story, it is posted on 'the wire', and somebody else picks up the story for inclusion in their

own publication.

On the web, the earliest syndicators of online content were the portal sites such as Yahoo and Excite. The basic idea behind these portals was that a reader could locate information from many sources from a single web site. Syndication on Yahoo has become extensive. The site no longer merely lists headlines; it also prints complete sets of news stories from suppliers such as Associated Press, Motley Fool and Forbes. While attracting remarkably little attention, Yahoo has become the most comprehensive news service on the web.

Syndication can be time consuming and expensive. Content syndicators want mechanisms that allow headlines and articles to be collected automatically. Programs that search through the web – called crawlers – have been around since the early days; the first well-known crawler was WebCrawler . Today, the most popular crawlers are AltaVista and Google.

But these are very generic crawlers and they do not organize their information in any systematic way. That's why they are better known as 'search engines' than as syndicators. Nonetheless, the technology for automatic syndication is essentially the same as for web crawling, and it was only a matter of time before automatic syndication came to the fore.

Perhaps the largest such syndicator is Moreover.Com. . This site collects headlines from 1500 newspapers and content providers around the globe and organizes the results into 280 separate categories. What Moreover has to do is retrieve the headline page from each of these content providers, parse the HTML in order to find headlines and links, and then store these in appropriate categories. It then outputs a series of RSS files, one for each category. News and information sites around the world use RSS feeds. Providing a similar service is iSyndicate.Com, which enters into content distribution agreements with publishers and provides RSS feeds and complete articles for syndication.

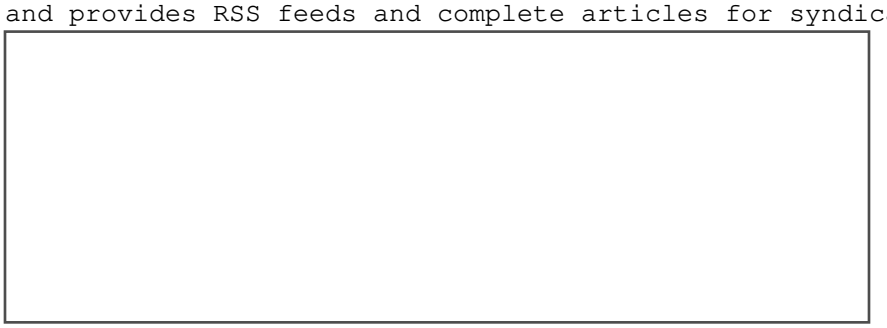

Figure 3: RSS Data Feeds

Pictured above is a flow chart diagramming the syndication process. Original content sites (Site 1 and Site 2) produce headings or content on different topics. The aggregator retrieves this content, which sorts the retrieved content, producing topic based news feeds in RSS or JS format. These news feeds are in turn retrieved by other content sites (Site 3 and Site 4) and are displayed as HTML pages.

Earlier content syndication sites collected content from content providers in the form of HTML pages. This is not nearly so simple as it looks. HTML is not designed to organize content; it is designed to display content. It turns out that it is a lot easier to retrieve and parse XML files – and in particular, RSS files. Sites that do this are called 'aggregators', and today's new breed of aggregators is focusing almost exclusively on RSS files.

RSS was used to good advantage by Netscape, but a major problem with the My Netscape directory was that users could not view the actual RSS files – Netscape would only let readers access the site summaries through its portal. The same was also true of another repository, My Userland, the portal application for the Scripting News Format discussed above. But RSS files may be located through yet another repository, XMLTree.Com, which indexes a wide variety of XML and RSS files. Launched early in 1999, the site has grown over the last year to include thousands of sites sorted by category.

5.1.4 4. Uses for Content Syndication

Although the easiest and most obvious use for content syndication is in the production of relatively current lists of news links on a given topic, RSS developers are beginning to perceive that a wide range of uses will be possible. In a document released in September, 2000, Ian Graham and Benet Devereux suggest the following :

New bulletins or news summaries, currently largely distributed using a simple XML dialect called RSS. An examples of this is My Netscape.

Web site content replication or distribution (often done using tools such as rdist, which is Rdist is a program to maintain identical copies of files over multiple hosts.

Database-related content distribution, such as gathering event calendar data for use in a local calendar.

Gnutella-like file/resource sharing services. This is a serve where multiple copies of the same file (for example, a music video) are located on different servers, with syndication information being used to facilitate retrieval.

Dmoz.org-like catalogues. The Mozilla Open Directory project is a human-created directory of Web-accessible resources. This directory is available as an open -source archive (in RDF), and is integrated into many other Web cataloguing systems (for example, Google or Lycos).

The HEML (Historical Event Markup and Linking) Project . This is a project aimed at creating a world wide collection of history-research related XML resources, with each academic research group being able to create their own resources, which can then be syndicated and distributed amongst the different institutions.

To aggregate proprietary scientific data, as described by David Detlefsen.

As Graham and Devereux point out, in each of these cases, "one organization publishes 'origin' data and makes it available in some form, and another organization downloads the data and processes the data to integrate it in some way into their own database or application."

5.2 Part Two: Content Syndication and Online Learning

5.2.1 5. The MuniMall Project

MuniMall, a project funded by Alberta Municipal Affairs, was intended to provide a common services and information platform for people working in Alberta's municipal sector. It would provide resources, learning and points of contact to elected officials, municipal administrators, and students of municipal government.

Figure 4: MuniMall Home Page

As such, it was intended to be what has since come to be called an "online community of interest" or "vertical community." The original design was modeled on the concept of online community as described in Hegel and Armstrong's Net Gain. At that time, the concept of content syndication had yet to reach the mainstream; it was envisioned as a portal for all things municipal in Alberta.

Because MuniMall was perceived to be a threat to existing services (and especially websites hosted by the Alberta Urban Municipalities Association and the Alberta Association of Municipal Districts and Counties), the commercial aspects of MuniMall were quickly removed and the website was re-purposed to provide a strictly educational function. To enhance its value as an educational site, MuniMall would include, in addition to resources and links to resources, an online simulation of a municipal website, MuniVille, to act as a training tool.

The removal of the commercial component probably doomed the project to failure: one the one hand, when government funding runs out (as it will in the spring of 2001), the project will be unsustainable. But more importantly, to act as the locus of a community of interest, the site would have had to be able to link to and contain information about all aspects of the community; to draw an artificial boundary around the content MuniMall is 'allowed' to have and that which is not is to limit its effectiveness as a community of interest.

As a research project, however, MuniMall remains invaluable. Unlike most work in the fields of online learning and online community development, MuniMall had an explicit mandate to merge educational content with information and resources used by the community of practice. In other words, MuniMall would be a tool used by municipal administrators in the course of their day to day activities, and at the same time, function as a teaching tool for students in the Government Studies certificate program.

The next three sections will describe three approaches taken to accomplish this.

5.2.2 6. Content Syndication

The first area of integration looked at by the MuniMall team focused on the resources used by both students and administrators. In particular, Alberta Municipal Affairs has over the years developed a Handbook for Municipal Administrators. This handbook contains detailed instructions on how to conduct a municipal election, draft and pass a by-law, approve building permits, and more. The Handbook, in turn, refers extensively to legislation and regulations governing the conduct of municipal affairs in Alberta.

Although an important – indeed, essential – resource, the Handbook was paper-based and not available online anywhere at all. It was maintained, as many similar Handbooks are, as a set of loose-leaf inserts into a massive binder. Periodically, updates would be issued from Municipal Affairs; these updates would be delivered to individual municipalities and also to the Government Studies program, where they would (sometimes) be placed into the binder.

An examination of the Handbook also revealed that it was out of date and in many ways redundant or internally contradictory. The maintenance of the Handbook was a major task for staff at Alberta Municipal Affairs, and the output was of minimal usefulness to practitioners in the field.

MuniMall proposed that the content of the Handbook be placed online and syndicated. Placing the content online would mean that it could be updated online, through a forms processing system, and thus, much of the time and expense in maintenance would be eliminated. Syndication, moreover, would allow the same (always up-to-date) document to be used in a wide variety of locations: and in particular, in online courses, in the MuniMall portal listing, and as help for any online forms or documents employed by municipalities.

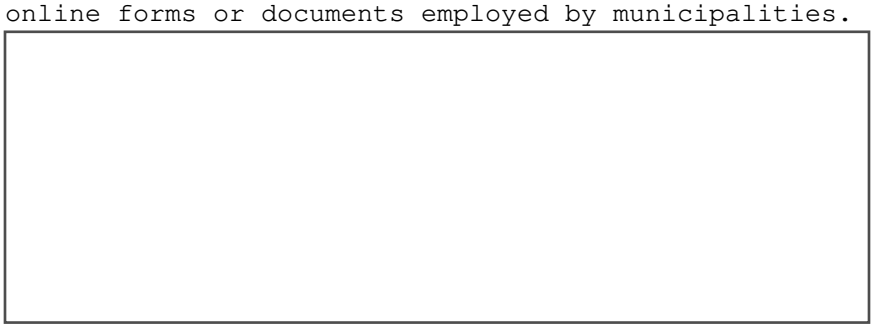

Figure 5: Content Data Flow

In the end, this model of content syndication was never put into place. Several major obstacles emerged:

First, the Handbook was (as mentioned above) in a considerable state if disrepair and would have required extensive revision, a task in which Alberta Municipal Affairs was unwilling to engage (as events transpired, they instead launched an extensive 'Best Practices' initiative which may have as a final outcome a content syndication model as described here). Moreover, Alberta Municipal Affairs had no mechanism for assigning authority or responsibility for the upkeep of the Handbook.

Second, it was not clear that MuniMall, or even Municipal Affairs, could get permission to distribute the content of relevant legislation as described. Copyright over the legislation is held by the Queen's Printer, which currently returns revenues to the provincial government through its printing service.

And third, even were the content available, there was no place to put it. The online course design for the Government Studies program adopted a mixed mode of delivery, with the course outline and discussion occurring online, but with course materials distributed as part of a paper-base package.

A modified version of content syndication was instead employed in the MuniVille to serve as a demonstration of how a similar technique could be employed by local governments for the wide distribution of documents and information. The MuniVille website consists of a set of topical pages, such as 'Industry', 'Recreation' and 'Restaurants'. Content in each of these pages is updated via an online form, and

the content is available for insertion into multiple pages. Thus, for example, a real estate agency could draw upon the community website to provide up to date demographic information; the municipal website could in turn draw from a real estate agent's site an index of new listings.

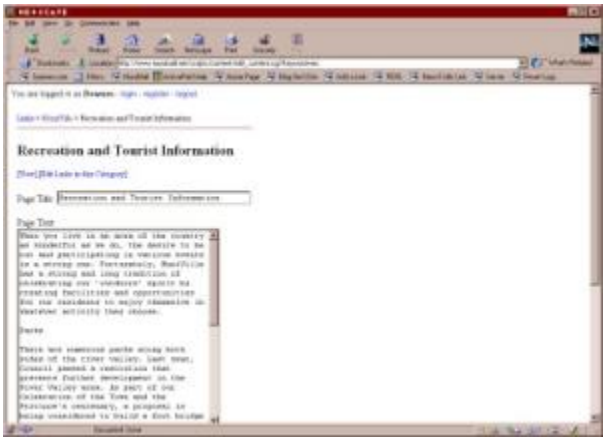

Figure 6: Content Input Window

Figure 7: Content Display Window

5.2.3 7. Link Syndication

As part of its mandate to provide resources and information, MuniMall developed a portal of links relevant to Municipal administrators and elected officials. To date, more than 1200 resources have been added to the portal, with more being added each day. Links are entered into a common database and then displayed in a set of topic-based pages, much like traditional portals such as Yahoo.

The idea behind the link syndication system was to act as a means of accessing resources that could not be stored as web pages on MuniMall itself. The most common type of these resources is the external link; MuniMall staff added a large number of links and MuniMall users were encouraged, through an online submission form, to submit their own links. Three major categories of links emerged: links that dealt with specific municipalities, link which addressed aspects of municipal governance (especially as it related to the provision of online services), and

links that related to some aspect of a community (in other words, links that correspond to one of the topic-based pages in the simulation).

Figure 8: Links Display in a Portal

In order to facilitate this system of link syndication, four systems were developed over-and-above the link submission forms and syndicated output. First, an automatic categorization tool was developed to sort the links as they were submitted. Second, an automatic link-retrieval engine (similar to a web crawler) called Grasshopper was built. Third, a link editing tool was created. And finally, a search tool or "drill" was added to the system.

Although stored in a common database, these link lists are available to multiple web pages. As new links are processed, output files in both RSS and JS are produced (the JS file is a server side Javascript file which can be used by any HTML page without special processing). Thus, the same list of links can be used in the MuniMall portal and also (for applicable categories) in the MuniVille simulation.

Figure 9: Links Display in MuniVille Simulation

The system was originally designed to allow for up-to-date resource lists to be used in online courses as well. Ideally, both students and people working in the field of municipal affairs would submit links. These links would then be embedded

in a WebCT course page (using the single-line Javascript command to embed the content).

To date, however, the link system has functioned mostly as a portal. Part of this is due to the fact that the tools are not as reliable as would be liked (the editor, for example, still has some major bugs in it). Part of it is due to the fact that there has not been a consistent and useful flow of content into the system – such a system needs multiple contributors, and more importantly, contributors expert in the field of enquiry. And part of it has been due to the fact that, other than the "today's links" page and MuniVille, there has been no place to display the syndicated content.

5.2.4 8. Discussion Syndication

As MuniMall was intended to foster an online community, a forum for discussion and communication was essential. To this end, a discussion list program (Allaire Forums) was added to the site, where it sat – empty.

It became apparent that the discussion forum had to be seated much more closely to the main content; indeed, the discussion forum had to be a part of the main content. Once again, the idea was that posts, lists of posts, and list of discussion topics should be syndicated, so that they were available to a large number of web pages.

Because no discussion list program currently offers this feature, a specialized discussion list program was developed and used in place of Allaire Forums. The program – CList – provides output in RSS and JS as well as HTML. In addition, CList, like many other discussion list programs, allows email notification as well (in other words, if the user selects the option, the program will send an email message when somebody adds another post to the discussion).

Discussion on the MuniMall site still languishes; the two threads today have a combined 17 posts. Indeed, the most effective use of CList has not been on MuniMall at all, but rather, on my personal home page, where I used the discussion list program to format and display articles – like this one – on one website, while using the JS feed to list and link to the articles on another one, my main home page. And even in this system, discussion is minimal.

Part of the reason for the ineffectiveness of the discussion tool is the low traffic. Although, starting September 2000, the tool was employed in one of the Online Courses; it is accessed only as an external link, and not embedded in the course content as designed. In addition, on both MuniMall and on my home page, traffic is low, proving once again that a certain level of traffic is necessary in order to sustain a discussion board. Third, there has been no concentrated attempt to foster discussion: no events have been scheduled, no course requirements for discussion, no moderation or introductory articles. And finally, the sort of people who use MuniMall are just the sort of people who do not have time to engage in unfocussed online discussions.

5.2.5 9. Why Things Didn't Work

I am standing before you and saying that, in three major areas of content syndication, the MuniMall project failed. As I suggested above, perhaps it was doomed to failure in any case because of the segregation of its potential audience. But it also failed as a result of a number of structural flaws. These flaws are worth investigating, especially when placed against an area of substantial success, yet to be discussed.

First and foremost, I think, an entity like MuniMall cannot exist in isolation. Like any form of syndication, it needs content at the input end, and it needs recipients at the output end. MuniMall suffered from shortfalls on both ends.

Input:

commercial and provider content was banned from the site almost immediately

government content, such as the Handbook, the manual, and even web site contents, was not forthcoming

there is a dearth of subject matter experts (or even knowledgeable participants) providing links, articles, discussion list posts and other materials

Some of this could have been addressed through better management. For example, a coordinated campaign to generate user contributions might have helped. Course professors should have been recruited to provide expert commentary. Students should have been recruited to provide discussion.

But in the absence of the more substantial content – especially content the target audience really wanted, such as business contacts and government documents – MuniMall was bound to suffer.

Output:

no external sites used MuniMall as a content source

A syndication site that cannot market its materials anywhere is a site which is in deep difficulty. Obvious locations for syndicated content would have included the online courses, community and government sites, and the AUMA and AAMD&C sites.

These problems are indicative of a second and deeper cause for the difficulties faced by MuniMall. The project, from its inception, ran counter to two major features of information networks:

First, the market was just too small. And as Metcalfe's law states that the value of a network increases exponentially with an increase in the number of participants, its corollary, which I'll call Downes's law, states that the value of a network decreases exponentially as the number of members decreases. A variety of factors, structural, organizational, personal and political, led to successive reductions in the numbers of people using MuniMall, and this led to its exponential decrease as a network.

Second, prospective participants in the network didn't participate (in other words, the size of the network decreases), an instance of Downes's second law of networks, which is, that the value of a network decreases exponentially as the size of the network decreases. As the associations, the commercial entities, online courses, and the governments were removed from the network, the value of the network collapsed.

5.2.6 10. A Success Story: The MuniMall Newsletter

The MuniMall Newsletter was launched in September, 1999, and circulation has grown steadily since that launch date (it now stands at 359, about a quarter of the total market population). It is widely read, often printed and distributed in municipal offices, commented upon favorably at conventions and in research studies.

Figure 10: MuniMall Newsletter

The MuniMall newsletter is an example of syndication in action. Published once a week, it contains links to websites and articles of relevance to municipal administrators and elected officials. It draws from oft-ignored sources, such as local newspapers and government press releases, and presents this list of links, each with a short description, as a weekly email message. The newsletter is also published on the MuniMall site, and as Items are added to the site, the "What's New" page is automatically updated.

Figure 11: What's New Display

The MuniMall newsletter address the two major weaknesses identified in the previous section.

First, it has content. The typical newsletter is a collection of links from external sources and articles produced by MuniMall staff. Moreover, this content is highly filtered, designed to reflect the specific interests of the community it targets. Such highly filtered content is possible only if some form of syndication is employed, whether the process is implemented automatically or by hand.

The Newsletter, in other word, incorporates the first two of the three types of syndicated content described above: it contains textual content, in the form of articles, and it contains resources, in the form of links. In only the third form

of content – online discussion – is the Newsletter lacking, though there is every reason to believe that with better content filtering and integration, a discussion component would be a useful addition (as it is in so many list services around the world).

Second, it has recipients. The MuniMall newsletter circumvents the usual channels for syndication, bypassing websites almost altogether, by being placed directly into readers' email in-boxes. Because it is an email newsletter, it is easy to read (people tend to use email a lot more than they tend to use a particular website), and because it provides a list of filtered resources, it is easy to use.

The MuniMall Newsletter thus offers two of the best features of content syndication: content and convenience.

5.2.7 11. All Together Now: Doing Educational Content Syndication Right

What can be learned about content syndication in the educational domain from the MuniMall example?

First, and not trivially: it is technically feasible. Using the tools described in this paper (or tools which are becoming widely available on the internet) any course (or program of courses) or any online learning application can tap into up-to-date resources from remote sources, and tap into them in such a way that content is tailored specifically for the course in question.

But second, and also not trivially: because content syndication requires the development of a network, the practices and politics of building networks must be observed. Especially where the syndication network is breaking new ground (which today, is everywhere), the ground rules and principles of participation must be laid out in advance of any development.

Because, third, a content syndication network needs content, and in an educational setting, it needs authoritative content, which means that the providers of that content – whether they be government agencies, university professors, or professional associations – must be on board and willing to provide that content.

Of course, this is a two-way street: fourth, no content provider can go it alone. The reason for this is clear: in our examination of the municipal sector, we found dozens of agencies which provide authoritative content of one sort or another, agencies such as newspapers, community websites, research institutions, multiple government departments, a dozen professional associations, and more.

Fifth, there must be an audience, which means that at least as much care must be taken to present content in contextually useful situations as is taken in gathering the content to begin with. Even less comprehensive content – such as found in the MuniMall Newsletter – can be widely used if it is presented in an attractive format; conversely, even the best content will not be used if it is not accessible. The mechanisms employed by the Newsletter, including content filtering and a gentle push, tell us what an attractive format is likely to look like.

And sixth, although the temptation is often to start small – a pilot course, a pilot class – in endeavors which depend on a network phenomenon, it is best to start with as large a set of participants as possible. A large network may be scaled back or subdivided if it becomes unwieldy, but a small network may never get off the ground because the interactions upon which it depends are not there.

[pic]

Netscape Website image. http://my.netscape.com/publish/images/mozillazine.gif

[2] Microsoft. Channel Definition Format. http://www.microsoft.com/mind/1197/cutting1197.htm

[3] Microsoft. What's New in IE 5: Significant Changes. http://www.activewin.com/articles/article_14.shtml

[4] Downes, Stephen. My Netscape 6.0. NewsTrolls, April 5, 2000. http://newstrolls.com/news/dev/downes/column000405.htm

[5] World Wide Web Consortium. Resource Description Framework. http://www.w3.org/RDF/. See also Downes, Resource Descriptions, unpublished (1999) http://www.atl.ualberta.ca/downes/threads/column050599.htm

[6] Dublin Core. http://purl.org/dc/

[7] IMS Protocols. http://www.imsproject.org

[8] Scripting News. http://www.scripting.com/

[9] Dave Winer. Scripting News in XML. Scripting News, Dec 15, 1997. http://davenet.userland.com/1997/12/15/scriptingNewsInXML

[10] Dave Winder. RSS 0.91. June 4, 2000. http://backend.userland.com/rss091

RSS 1.0. August 24, 2000. http://purl.org/rss/1.0/

[12] Yahoo. http://www.yahoo.com

[13] Yahoo Daily News. http://dailynews.yahoo.com/

[14] Webcrawler. http://www.webcrawler.com

[15] Alta Vista. http://www.altavista.com

[16] Google. http://www.google.com

[17] Moreover.Com. http://www.moreover.com

[18] iSyndicate.Com. http://www.isyndicate.com

[19] My Userland. http://my.userland.com

[20] XMLTree.Com http://www.xmltree.com

[21] Ian Graham, Benet Devereux. The Syndication Project. http://www.java.utoronto.ca/news/index.html

[22] MagniComp. Rdist Home Page. http://www.magnicomp.com/rdist/rdist.shtml

[23] Historical Event Markup and Linking Project. http://www.heml.org/

[24] David Detlefsen. How I want to use Manila, MyUserland & RSS. August 14, 2000. Backend.Userland.Com discussion list. http://backend.userland.com/discuss/msgReader$84

[25] John Hagel III and Arthur G. Armstrong. Net Gain. Harvard Business School Press, 1997.

[26] Stephen Downes. MuniMall: A Comprehensive Proposal. September, 1999. http://www.munimall.net/about_munimall/proposal.cfm

[27] Independent research report, as yet unavailable (but we saw preliminary results).

This is the paper I wrote in Australia and finished on my return to Edmontonin the spring of 2001. This is an architectural paper – that is, it is intended to describe the major components of a learning resource distribution system. It was adapted as a principle architecture document for one of my projects, peggasus.ca (which launched November, 2003) and outlines the basic mechanisms described in much greater detail in the essays that follow.

6 The Learning Marketplace – A Concept Paper

Written in April, 2001. Unpublished.

6.1 Overview

Online learning involves the delivery of courses and course components via the Internet.

These courses and course components are delivered from an educational provider to a student. Usually, learning objects are created and managed from within a learning management system (LMS) hosted by the educational provider. The LMS acts as the student's primary point of contact with the educational provider, and in addition to delivering course materials, also provides a testing environment and manages the student's course grades. The LMS also may provide interactive facilities such as a discussion board or online chat.

The educational environment depicted by the typical LMS may be described as a 'one to many' relationship between institution and students. That is to say, for the typical LMS, there is only one educational provider and many students who interact with this educational provider. Administrative tasks such as course selection, registration, tuition payment and accreditation are handled externally to the LMS, sometimes by student management systems, sometimes by more traditional means.

This system works well in an environment where one student will take many courses from a single provider, as is the norm, say, when a student is completing a program of studies such as a university degree or college certificate. But as the educational environment becomes more fluid and as post graduates seek to extend their education, we envision a need for a system which may be depicted as 'many to many', that is, where one institution offers courses to many students, and where one student may take courses from many institutions.

To date, support for a 'many to many' model of online learning has consisted of online portals offering lists of courses sorted by course category. And while such course selection portals offer a useful indication of the range of courses available from different institutions, these course portals are passive: while they list online courses, they do not provide access to online courses. In order to participate in online learning, the student must revert to the 'one to many' model for each educational institution.

The purpose of an online learning marketplace is to provide a 'many to many' environment for online learners. Using a single interface, a student may select, register for, and be delivered courses and course components from a number of institutions. The learning marketplace acts as a broker between institution and student, passing information from providers to students, and passing information from students to providers.

6.2 Why a Learning Marketplace?

The advantages of developing a learning marketplace may best be described by analogy. Suppose that the learning institutions are similar to food producers (such as, say, Kraft, General Mills or Heinz). And suppose that the student is similar to a food consumer, a person who wishes to purchase food for consumption. Then under the 'one to many' system, each food producer sells food directly to the consumer. But on the 'many to many' system, the learning marketplace is like a grocery store, offering a single point of sale to consumers for products offered by many consumers.

This advantage becomes even more apparent when we allow the consumer to have an account with the seller. In order to purchase food by account from each food producer individually, the consumer would need to establish a separate account with each of them. Each producer would have to perform a separate credit check; the consumer, meanwhile, would have to make and keep track of payments to each of the many food producers. It would take a significant effort for a consumer to purchase food from a new supplier, and it would take a significant effort for a food producer to acquire a new customer.

Under such a system, a food consumer would tend to make purchases from only one or a few food producers. This would force the food producer to provide a wide selection of foods, in order to ensure that the food consumer is able to purchase the desired item. But the food consumer would find the selection overly restricted, because no food producer can produce enough variety to satisfy the individual needs of large numbers of food consumers. And while the consumer may be aware of the existence of other foods, perhaps by browsing through a catalogue, the effort required to make a purchase from a different producer would limit such purchases to all but the most important.

Colleges, universities, and other educational institutions are in the position of food producers selling directly to consumers. They are forced to offer a wide selection of courses and programs to meet their students' varied needs. But even the most adept of them can offer at most a few hundred courses because of the logistics of offering increased numbers of courses. Students, taking courses from a single institution, often find the selection restricted and inadequate to their needs. But the effort required to create an account with a new course provider limits the number of institutions they can work with.

From the point of view of a new institution, the situation is even more difficult. It cannot easily specialize, because it must reach a large number of students in order to become sustainable. It has no place, other than online catalogues, to advertise its courses, and students seeking to take a course from a new institution must overcome a built-in disincentive. This is why, today, large institutions dominate the educational market; there is no easy means for a small or medium sized enterprise (SME) to market and sell its wares.

For the student, then, the advantages of a learning marketplace are clear:

A much wider selection of learning is available

Because of the wide range of offerings, a much more personalized listing of offerings is possible

There is no extra effort involved in purchasing learning from a new institution

And purchases are made from a single point of contact

Or, in summary: the student has a better choice of offerings, and needs to take

less time in accessing those offerings.

For the educational institution, the advantages of a marketplace are also clear:

A much wider audience is available for its course offerings

The institution can focus and specialize on particular types of offerings

It is much easier to acquire new students

Or, in summary, the institution is able to reach a large market and therefore to focus on specialized delivery or on extending market reach.

These advantages become even more apparent when we bring third parties into the equation. Suppose, for example, the student works for a corporation and is funded by that corporation for professional development. There then exists a three-way relationship between the student, the corporation, and the educational institution. In the 'one to many' model, for each institution offering courses to the student, a separate set of transactions must occur between the institution and the corporation in order to establish billing and reporting procedures.

Under such a scenario, the corporation is likely to select one or a few educational providers as 'preferred institutions' and students would be limited to selecting course offerings from those providers. So even if a student wishes to take a course from a different institution, an additional barrier is put into place preventing this transaction, because a relationship between the corporation and the institution must first be established.

And in fact, in today's educational environment, there are multiple parties involved in course development and delivery. There are students and institutions, corporations and governments, professional associations, standards bodies and certification agencies, testing centers, regional learning centers, educational consortia, and more. When each course offering to a student requires a multitude of bilateral agreements, a significant disincentive exists against any change or development. Course selection is severely limited. And course enrollment is cumbersome and slow.

6.3 Creating Course Offerings

We now move into a description of the mechanics of a learning marketplace in order to describe how such a system would work. While this description is necessarily speculative, it draws on developments already taking place in the field of online learning. The purpose is to show that a learning marketplace is feasible and to suggest how it would be structured.

In order for there to be a marketplace, there must be a product. In the case of a learning marketplace, the produce consists of learning opportunities. We have assumed in the discussions above that these opportunities would consist of online courses, but there is no need to be so restrictive. Learning opportunities may come in a wide variety of sizes and delivery modes.

To see this, we need to distinguish between the learning opportunities themselves and the descriptions of the learning opportunities. Think of the former as being like the contents of a can and the latter as the label on the can. In a grocery store, a purchaser looks at the label of the can in order to decide whether to purchase the contents. From the point of view of the grocery store, the can may contain anything at all; what is important is that the label provide certain information to the customer. Indeed, the contents of the can are completely untouched by the grocery store. The transactions between the store and the producer, and the store and the consumer, are based entirely on the label.

A similar situation exists in a learning marketplace. For each educational offering there is a corresponding label. The label describes the contents of the educational offering. The marketplace displays this label to the student, and on the basis of the label, the student decides whether or not to purchase the contents. Upon paying the grocery store, the store delivers the label to the consumer. Of course, what is important here is that the label is attached to the contents, and that the label accurately describes the contents.

In the terminology of online learning, the combination of a can and its label is called a learning object. The label itself is called the metadata describing the contents of the learning object. The learning object may be anything at all, so long as the label accurately described the contents, and so long as the label is attached in some way to the contents.

In grocery stores, certain standards have evolved regarding the labeling of cans. Consumers expect, for example, to find the name of the producer, a title or brand name, a description of the can's contents, the price offered for the can, usage or shipping instructions, expiration date, seals or certificates of quality, and recommended use. Some of this information is contained in a Universal Product Code (UPC), the bar code scanned by the grocery store on checkout.

The metadata for learning objects must also be machine-readable; this is what allows transactions to be automated. The machine-readable code for learning objects is written in a language called XML. The vocabulary for that language is expressed by learning management systems such as IMS or SCORM. The details of IMS and SCORM are unimportant at this point (moreover, these vocabularies are in transition and are to limited for a learning marketplace; it is impossible in IMS for a provider to specify a discount for government employees, for example) but the mechanism is important.

For each learning object that an institution wishes to offer for sale, a corresponding set of metadata must be created. From the point of view of the institution, this is like filling out a form with the required values.

The form consists of two sets of elements: the name of a field and the value of the field. The name of the field is a word in the IMS vocabulary, such as "author". The value is the actual value for that learning object, such as "Fred Smith".

These names and values are stored in a database located at the learning institute. Thus, for example, the University of Alberta would have on campus a database consisting of all the courses it wishes to offer for sale through a learning marketplace. This database is connected to a web server. When a remote user accesses the web server, the server retrieves the information from the database and presents it to the user as IMS compliant XML files. The content of the XML file corresponds to the content of the database.

In order to list a learning object, a learning marketplace system connects to the institution's web server and requests that it send the metadata. This transaction happens completely automatically; the learning marketplace computer program automatically connects to the institution's web server at predetermined intervals of time (this is known as polling).

Some filtering may occur at this stage, since there may be no need for the entire database to be transmitted to the learning marketplace. The learning marketplace may access the database using a search procedure, thus retrieving only a certain set of records from the database. A learning marketplace that focused on geology,

for example, would retrieve metadata only for geology courses. Moreover, the educational institution may also filter its results. For example, it may quote different prices to different learning marketplaces. Or it may make only a subset of its courses available to certain learning marketplaces. These are policies determined individually by each learning marketplace and each educational institution; these policies are determined autonomously and implemented locally.

The resulting XML file is thus transported automatically to the learning marketplace where it is parsed and stored. More filtering may occur at this stage; the learning marketplace may wish to store only a subset of the information provided by the educational provider. In addition, the learning marketplace may add some information at this point. For example, it will add information about when the file was downloaded from the educational provider. It may also add information about the provider which may not be in the provider's database, such as the name of the provider, its geographical location, or its user satisfaction ratings.

The learning marketplace then stores the information in its own database, ready for display and distribution as needed to prospective students.

Some systems exist which almost perform this function. IMS compliant repositories of learning object metadata already exist. For example, MERLOT collects information from member providers and displays the results in HTML or XML format. However, MERLOT does not poll member institutions for metadata; staff at each institution must submit the information using an online form. This means that for each learning marketplace, staff would have to complete an additiona for for each learning object, an obviously cumbersome procedure. In addition, MERLOT collects metadata only from member institutions, thus limiting the number of course offerings available to prospective students.

No learning marketplace is obliged to list offerings from all educational providers, just as no grocery store is obliged to stock cans of food from all food producers. Learning marketplaces can and will determine which educational providers they will poll for information. This list is contained in a secondary database in the learning marketplace and includes not only the address of the institution to be polled but also the exact request that will be made of the institution's database.

Some learning marketplaces will charge fees for inclusion, as MERLOT does. Other learning marketplaces may charge for exclusive listings in certain product areas, just as a grocery store may charge Kraft if Kraft wants to be the only producer seloing peanut butter at the store. Other learning marketplaces may charge producers for premium placement of their product, just as Safeway extracts a discount from Heinz for an end-or-row display of Heinz Ketchup. These policy decisions, must be balanced against the marketplace's usefulness to the student. The value of the marketplace to the student is created by offering choice and convenience, and this value is diminished by any exclusive agreement restricting choice and convenience. If a store only offers Heinz ketchup, fewer people will shop at the store, reducing the profit earned by the store (and not incidentally, its value to Heinz as a distribution agent).

6.4 Creating the Offer

Once a store is stocked with product offerings, its next major task it to create a mechanism for customers to view and purchase those offerings. In a store, the products are clustered by category and arranged on shelves, labels facing outward.

Aisles are created to allow the customer to view the offerings and shopping baskets are provided to allow customers to collect the desired products. Near the exit of the store, cash counters are provided where the objects are scanned, a bill presented, and payment made, thus completing the exchange. Some stores also provide delivery or other post purchase assistance.

The dominant model employed by grocery stores is the browse model; this is the model employed by online course portals today. Products are clustered into categories, and prospective purchasers are invited to walk through the aisles and view a selection of individual products for each category. Thus, the aisles of a store are labled in much the same way as the sections of a course portal: there is a pet food aisle, a pasta and sauces aisle, a meat aisle; and there is a biology aisle, a geography aisle, and a mathematics aisle.

The problem with the browse model is that the customer must know how the store categorizes its products in order to locate a certain type of offering. Is Heinz ketchup located in the sauces aisle, for example, or is it located in the condiments aisle? Often, stores will provide multiple product placements (putting ketchup in both sauces and condiments, for example) and will sometimes provide associative placement (putting ketchup where the meat is sold, for example). But even with these innovations, customers nonetheless must spend a fair amount of time browsing, especially if they have specific needs, and there is a certain risk that they will never find the desired product at all.

Stores are aware of this and attempt to interpret their customer demographics in an effort to minimize browsing. During certain seasons of the year, for example, a store may offer 'barbeque kits'; knowing that a certain percentage of shoppers will arrive at the store in early July looking for charcoal briquettes, steaks, steak sauce and lighter fluid, these products will be grouped and placed at the entrance to the store during the summer months. In winter, these displays are replaced by collections of wrapping paper, take, ribbons, and popular gift selections. Department stores generally place men's wear items on the main floor exit based on the perception that men are more likely to value convenience than selection; the women's wear is further inside the store but offers much more choice. Stores also record consumer purchases – this is the purpose of affinity programs such as Air Miles or Club Z – in order to determine shopping patterns and product clusters, again in an effort to minimize browsing.

Online course portals provide none of these conveniences. This is because the user of an online course portal does not fit into a certain demographic (except, perhaps, the desire to purchase a course offering). Because the actual shopping occurs outside the course portal, it is not possible to create affinity programs or detect purchasing clusters. Indeed, the online course portal knows almost nothing about the purchaser, and so is forced to rely almost exclusively on the browse method, inconvenient though that may be. And as the number of online course offerings increases, the browse mechanism becomes increasingly unwieldy. This is alleviated to some degree by a search mechanism, but as web users have learned, in an arena where there may be tens of thousands of offerings, the search must be very specific.

In the retail arena, this is mitigated to some degree by the segmentation of specializations. No store (except maybe Wal Mart) sells everything; stores will focus on a certain market niche in order to reduce the browsing required by potential customers. A similar trend is beginning to happen in online course portals, where a given portal may list courses only law courses, say, or only

computer courses. But this trend is still a stop-gap. In a global education market, there may be tens of thousands of courses in a certain field, thus reproducing at a lower level the same sort of difficulties in searching. Moreover, a new difficulty is created: that of locating the appropriate specialty course portal. A portal of portals is required, but without being able to scan the contents of each individual portal, browsing becomes even more hit and miss.

The learning marketplace circumvents many of these difficulties by creating individual profiles for each potential customer. In addition, by handling the transaction as well as the offering, the learning marketplace is able to learn the sort of offerings that may be appealing to the potential students. It is likely that niche learning marketplaces will be created; for example, a professional association may create its own learning marketplace. But these niches will be based on some similarity among the customers rather than on some similarity between he products offered.

This latter distinction is important. No person buys products from only one store; at various times, a person may need shoes, food, computers, art supplies or talcum powder. Thus if products are groups according to product type, consumers will still have to shop at many stores. However, if products are grouped according to consumer type, then stores can offer a wide range of product types, but product types which are typical of consumers in that demographic.

Such stores exist, although because consumer types tend to be geographically distributed, they are less common. Nonetheless, at a university we will find stores specializing in the purchasing needs of students: such stores will stock pens and paper, binders, school jerseys, caffeine tablets, snack food and newspapers. Such a clustering of products makes no sense downtown, but makes eminent sense in the Student Union Building. A general store at a beach will display a certain sort of clustering, offering towels, sunscreen, paperbacks, souvenirs and beach balls.

Thus a learning marketplace can refine its offering in two major ways: it can create demographic profiles of each of its users, and it can cluster a group of users by user type. Thus, we envision a multitude of learning marketplaces, each developed or sponsored by a professional association or some similar affinity group, in which members share a set of common interests with other members, and obtain personal characteristics my means of a personal profile.

6.5 Customization and Personalization

In order to arrive at a description of the offers that will be placed in front of individual prospective purchasers, it is necessary to discuss the means by which these offers are tailored for each individual. There are two major mechanisms, known respectively as customization and personalization. Here we adopt a terminology that has become common in the field of enterprise portals.

Customization is individualization performed by the portal or online service. For example, if a person elects to purchase a Beatles album, a music purchasing portal may then recommend albums by similar artists (such as the Kinks or Klaatu). If a person reports that he lives in Montreal, then he may be offered Montreal Canadiens merchandise instead of the more popular Toronto Maple Leafs merchandise. Customization occurs when the online service detects some feature of the individual and proposes product offerings associated with that feature.

Personalization is individualization performed by the individual user. For example, if a person enters his name as ’John Smith’ and the service greets him by

saying ’Hello John,’ this is an example of personalization. Or if a user requests to be shown new books by Stephen King, then a placement of The Stand on a list of offerings reflects personalization. If a person has already purchased The Stand, then this book is not displayed among the personal offerings. Personalization reflects a user’s tailoring of his or her own environment and is necessary because customization can at best reflect group demographics, and never individual demographics. Personalization can be explicit, that is, the result of some direct action by the user (such as selecting a background colour) or implicit, that is, the result of some other action of a user (such as the purchase of a particular product).

Customization and personalization are affected not only by user demographics but also by events external to the user. For example, if an educational institution releases a brand new course offering, then this course is more likely to be displayed than an older course offering. If the individual is sponsored by a certain company, then courses favoured by that company are more likely to be displayed than courses not favoured by that company. The exact listing of courses for any individual on any given display will be the result of a dynamic interplay between user actions and events and external actions and events.

The exact details will vary according to a wide variety of parameters but the mechanics will be similar in every case. For any user at any time, there are two sets of variables:

The set of offers which may be displayed (in a learning marketplace, the set of learning offers which may be displayed)

The number of offers which can be displayed

This latter feature is the limiting variable. For all practical purposes, only a limited number of offers may be displayed on a screen at any given time. A listing of ten thousand courses is not useful to the prospective purchaser (nor, for that matter, is a listing of only one course offering).

The number of courses that may be displayed is subject to several variables:

The mode of access – it is possible to list more courses on a broadband web browser than on a wireless PDA

Limits set by the user – the user may elect to see more or fewer options. The user may subselect, effectively filtering the offers displayed (today the user is interested only in customer services courses, while tomorrow he may select from a more general list of courses)

The number of courses available – obviously, the marketplace can only display available courses; if the number of providers is limited, or if the topic area is very specialized, then only a few courses will be available for display

Out of ten thousand course offerings, then, only a small number – say, ten – may be displayed on the screen at any given time. The actual ten to be displayed are determined by two major factors: filters and preferences.

A filter is a mechanism that eliminates from consideration a certain number of courses based on course information and user data (as expressed by customization and personalization). The most common filter is the search function; only those courses that satisfy the search parameters are candidates for display; the search process filters other courses. But in addition to search, a number of other filters may come into play, for example:

Mode selection – a user may elect to view only online courses, for example

Platform selection – courses which require a MacIntosh platform are not displayed to a user using a Windows system

Accreditation – a user may elect to view only certified courses

Funding Limitations – the user may elect to view only courses funded by the employer

Pre-requisites – A user may not be qualified to take certain courses, or may require formal admission to an institution before taking certain courses

Filtering may dramatically limit the number of candidate courses, but depending on the display limit, more refining may be necessary. This refining is accomplished by means of preferences; the purpose of a preference is to sort the candidate courses from the most preferred to the least preferred. Each candidate is assigned a preference weighting; customization and personalization features determine the preference weighting. Preferences may be defined by the learning marketplace, the user, or by external agencies, such as the user's employer.

For example, higher preference may be given to courses that:

Are newer

Are from a provider that provides discounts to the employer

Are associated with recently completed courses (for example, have a recently completed course as a pre-requisite)

Are Canadian

Are three hour courses (as opposed to three week courses)

Have higher student evaluation ratings

Each candidate course remaining after filtering is thus assigned a preference value; those courses with the highest preference value are displayed first, up to the total number of courses which may be displayed as determined by the number of offers that can be displayed.

The preference selection is enabled by a sequence of operations performed on the contents of a set of databases. In a relatively simple example, the databases involved are as follows:

The Course Offerings Database – this is the database created by the learning marketplace as a result of polling the various educational providers. This database contains information about each course on offer, thus providing values according to which preferences may be assigned.

The User Database – this is the database created by (and about) the user in question. This database contains normal demographic data (such as the user's name, age, profession, and location). The user database also contains links to other databases. For example, the user database will list the user's employer; this is a signal to the system to take into account the employer's preferences. The user database will also contain historical information about the user, for example, those courses already completed.

The Session Database – this is a dynamic database updated each time a user logs on for a new session. This database includes information about the user's mode of access (eg., desktop browser or wireless PDA), current location, and session-specific variables (such as search parameters or category selection)

The Employer Database – this includes global information about the employer, such as the employer's name and billing address; it may also contain marketplace-

specific information, such as a list of those courses it is willing to pay for, or those certification agencies it is willing to recognize. The employer database may include subsections for individual employers, thus allowing, say, one type of employee to enroll in management courses while directing other types of employees to technical courses.

The Preferences Database – this is a list of all candidate courses and preference values for each course, based on values obtained from the other databases

The actual list of course offering, therefore, consists of a set of operations over the contents of these databases. For example, it might look like this:

For each candidate course:

assign it a preference of 10

if it was released in the last week, then if the user has assigned a preference for recent courses, then add 10 to the preference

if it is certified, then if the certification agency is preferred by the employer, add 10 to the preference; otherwise, subtract 10 from the preference

if it leads to a degree selected as an objective by the user, add 10 to the preference

if it is available online, and if the user has indicated a preference for online courses, add 10 to the preference

and so on, until:

write the resulting preference value into the preferences database

6.6 Presenting the Offer

The presentation of course offerings consists of two types of display: the list display, and the object display. In the list display, a course is displayed as one of a list of courses; in the object display, the course is displayed by itself. Think of the list display as being like what you see when you see a row of cans on a shelf, and an object display as what you see when you take a can from the shelf and look at the label.

In the list display, a certain, restricted, set of information is displayed to the user. Three major factors determine what information is displayed:

Information deemed essential by the learning marketplace (for example, the learning marketplace may require that course titles always be displayed)

Information deemed essential by the course producer (for example, the University of Alberta may require that the course instructor always be displayed)

Information selected by the user (for example, the user may require that the price always be displayed)

It is important to keep in mind that the list display may be provided to the prospective student in a variety of formats. Of course, the student may obtain the display any time by logging on to his or her learning marketplace home page. Or this information may be coded in RSS or JS format and embedded into some other web page, such as the student's association web page or corporate desktop. Additionally, the list may be compiled and sent as a weekly email reminder. Or it may be formatted into WAP or some similar wireless protocol and made available as a wireless web page or instant message.

The user moves from list view to object view by clicking on the link for a given course object from the list provided (the user may alternatively elect to see more

of the list by selecting the 'more' option, or may elect to recreate the list by entering search or other parameters).

In the object display, all relevant information about the course is displayed. The contents of the object display may be determined by the same factors that inform the list display. Additionally, the learning marketplace may at this point introduce new information. So in addition to seeing relevant course information, the user may also be able to:

Read reviews of the course from previous users

Enter a discussion forum held by students in the course

Search all discussions for references to the course

Contact the educational institution for more information

Having reviewed the course and having decided to take the course, the user may then elect to purchase the course. A link is provided on the object display for this purpose, at which point the student enters the transaction phase.

The Transaction

The transaction consists of three essential stages:

Entrance into the course

Work during the course

Completion of the course

In this section we consider only the first of these three stages. The entrance into the course consists of two essential components:

Course registration

Course delivery

In the first component, the user applies for and is granted admission to the course in question. In the second component, a transfer of educational materials (the course contents) occurs between the educational institution and the student.

In order to complete course registration, a series of transactions must be completed involving the exchange of information between the student, the educational institution, and any third parties which may be involved, such as the funder. Generally, the registration consists of the following steps:

Admission into the institution – many institutions restrict admission into certain programs of study based on the qualifications of applicant. The mechanics for admission into each institution vary and there is often a delay as applications for admission are evaluated manually. Two possibilities occur here:

The student has been previously admitted. In such a case, the learning marketplace will provide the institution with the student's institution-specific identification, such as a student number

The student has not been previously admitted. In such a case, the student is transferred to an admission subroutine. With the student's permission, demographic information and required documentation is provided by the learning marketplace to the institution. The learning marketplace automatically generates transcript requests. Upon acceptance of admission, the institution transfers the student's institution-specific identification to the student directly or to the learning marketplace, depending on expressed preferences and institutional policies

As mentioned before, the admissions process is often not automatic and there may

be a delay at this stage. It is important for a learning marketplace to enable automatic application for admission, but whether that is possible will depend on the educational provider. Institutions will be encouraged to enable automatic admission as those that don't will see prospective students go elsewhere.

Ideally, courses from institutions where a student is not admitted, or could not be admitted, will not be displayed (or will have a very low preference).

Admission into the Course – assuming that the student is admitted into the institution, the institution must determine whether the student will be admitted into the course. Several factors may affect course admission, including course pre-requesites and the number of seats available. Provided that the student has an institution-specific identification, the institution should be able to determine whether the student may be admitted to the course; again, however, not all institutions will make this available automatically. However, the learning marketplace should make automatic admission possible, though whether automatic admission is available will depend on the institution. As above, institutions which do not permit automatic admission to courses will find students seeking courses elsewhere.

Presentation of Admission Information – assuming that institutional and course information are automated, the student will then be presented with admission information. Such information could include, say, a start-date for synchronous course offerings, a room and building location for in-person courses, tutorials available, required texts or online resources, total costs, and other course specific information. The student at this point has the option either to accept the admission as presented or to decline the admission.

Acceptance and payment – upon electing the accept option, the transaction is concluded with the payment of fees. Depending on the user, the institution, and third parties, such as the user's employer, one of several outcomes may occur:

The student is funded by the employer, in which case, upon acceptance by the employer (which may be automatic), the employer will be billed by the institution

The student must pay directly, at which point a credit card payment subroutine will be invoked

In either case, this transaction occurs via the learning marketplace. That is to say, the student or employer pays the learning marketplace, and the learning marketplace pays the institution. This is necessary for the following reasons:

The marketplace may extract a percentage of the course fees

The employer is billed by a single entitiy, and receives a single invoice, thus simplifying its accounting procedures

The institution receives money from a single source, thus simplifying its accounting fees

Placement into the course – once the financial transaction has occurred to the satisfaction of both the student and the institution, the student is placed into the course. Again, it is preferable if this happens automatically, because the institution may then provide course-specific information to the learning marketplace. Minimally, the institution should provide (a) the online access point to the course, if it exists, and (b) logon or other course-specific identification.

This allows the learning marketplace to place current course data into the student's database. If a student is currently registered in a course, then access

to the course may be provided through the marketplace. Thus, the student is able to access all his or her courses, no matter who the course provider, from a single point of contact.

The transaction is the most critical component of the development of the learning marketplace. As insisted in several points above, it is desirable that the transaction be as automated as possible. Just as Heinz is comfortable in allowing a consumer to select and pay for a jar of ketchup at the grocery store, so also educational institutions must be comfortable in allowing a student to sign up for and pay for a course through a learning marketplace.

Thus, in the development of a learning marketplace, it is essential that protocols be developed which enable such automatic course registration. This involves a certain amount of political work, as institutions tend to keep a tight reign on admissions and registrations. And it requires a certain amount of standards and protocol development, as the learning marketplace must speak and listen in a language understood by university computer systems.

6.7 Course Delivery and Completion

In most cases, the only role played by the learning marketplace in course delivery will be to provide a point of access to the course from the user's learning marketplace home page. On selecting the course, the learning marketplace may send login information to grant the student quick access to the course, though this may vary according to user and institutional priorities.

Some interaction between the marketplace and course providers may occur. For example, the learning marketplace may poll the course provider on behalf of the student. It may, for example, request any course announcements of immanent deadlines. These announcements may then be relayed to the student either via their learning marketplace home page, via email, or via wireless instant message.

Otherwise, the learning materials are delivered directly from the institution to the student. This is important because different institutions may use different learning management systems – one may use Blackboard, while another may use WebCT. Or the learning materials may be customized application, such as Java Applets or downloadable programs.

Only when the student has completed the course (or when the completion date has expired) is the learning marketplace again involved.

Upon course completion, the relevant information is then delivered from the educational institution to the learning marketplace. The learning marketplace in turn relays this information to the student, and with the student's permission, to any relevant funders or certification bodies.

6.8 Third Parties

A number of third parties – institutions other than the student and the educational course provider – are involved in a learning marketplace.

Foremost of these is the entity that operates the learning marketplace. As mentioned above, learning marketplaces are intended for certain types of students. A learning marketplace, for example, any be constructed for members of a professional association, or for people who work in a certain market sector. Indeed, learning marketplaces will often become essential components of a sector specific or interest based online community, being one of the range of services offered to members by that community.

The learning marketplace operator has three major tasks:

The selection of educational institutions that will be allowed to provide course offerings through the learning marketplace

The recruitment of additional third parties who may play a role in the offering and purchase of educational activities, and

The recruitment of a body of students who are potential customers for those course offerings

As mentioned above, the learning marketplace operator may elect to be as open or as restrictive in the selection of educational institutions and course offerings to be included in the marketplace. Generally, it will be preferable to include offerings from as many institutions as possible, but in a limited a range of topics as possible. This provides potential students with maximal selection with minimal browsing.

As the operator of an online community, the learning marketplace will need to establish a relationship of trust and interaction with its members. As mentioned above, the learning marketplace should be one of an array of services offered to members. The online community should provide a forum for member interaction, should be a locus for industry related news and information, and should serve as a marketplace for other sector specific goods and services. While specialization in the definition of membership is desirable, specialization in the range and type of services offered is not.

Other third parties will be those entities affiliated in some way with the members of the community and especially as regarding the members' learning interests. Some key third parties may be identified:

Employers – employers will often play a role as interested third parties in a learning marketplace. Employers serve two major functions: first, they may guide some members (either current or prospective employees) in the selection of appropriate learning for employment, and second, they may provide funding for some employee learning.

Employers are likely to interact with the wider community in a number of other ways. An employer will find a sector or association specific community an excellent location for recruitment or subcontracting. Employers may also elect to advertise goods and services to members of the community employed with other companies.

An employer interface to the learning community specifically will consist of a set of preferences and filter protocols. As described above, employer preferences may determine which courses are offered to prospective students at all (assuming the prospective student has the employer filter turned 'on') and employer preferences may affect the placement of course offerings on a student's learning marketplace home page.

The employer may also exchange information about the student (with the student's permission). For example, an employer may wish to be notified of completion of a specified course which it has funded; it may be interested in funded or unfounded courses completed as part of the employees overall profile (which may in turn be used as part of the expert-detection function in the employer's own knowledge management system).

Employers may also play the role of educational providers. An employer may elect to provide course offerings to prospective students both inside and outside the company. It may be worth while, for example, for an employer to offer free or low

cost company specific training to prospective employees; this reduces the cost of training on the job, and also gives the employer better data when making hiring decisions.

Certification Agencies – there are two major types of certification agencies that may be involved in a learning marketplace. On the one hand, there may be agencies that certify course offerings. On the other hand, there may be agencies that certify individual competencies.

Agencies that certify courses – provincial government departments of education, for example – may be invited by educational institutions to review and accredit course offerings. Such certification agencies would then report course certification information in much the same way educational institutions provide course metadata. The learning marketplace, after loading course metadata from educational providers, would then poll certification agencies for information about the courses listed; this certification information would then become part of the course record on the learning marketplace.

For any given course, many certification agencies may come into play. Some courses, recognized by the provincial government, may also be recognized by, say, Microsoft as being a standards compliant course. The same course may also be recognized by a professional association as being eligible for completion of a certification. And the same course may also carry a stamp of approval by an unaffiliated user group.

The learning marketplace managers, since they control the polling of certification agencies, have the responsibility over deciding which agencies to poll and which to omit. Of course, such a decision is best made in consultation with members of the community.

Agencies that certify individuals (as, say, being eligible for a certification) may wish to interact with the learning marketplace for information on course completions. An individual who desires a professional certification may register this desire as a user preference, at which point the certification agency preferences come into play in determining the list of offers available to the student. Moreover, with the student's permission, that agency may be notified of course completions, thus enabling it to recognize an individual student's satisfaction of the requirements for certification.

While I was in Australia I was able to watch students and radicals protest the University of Melbourne's signing with Universitas 21, a commercial consortium that almost seems to exemplify the sorts of complaints posed by people like David Noble and Steve Eskow. But while I have a lot of sympathy with such complaints the appropriate response is not to combat them with ineffective methodologies and technologies. In a world of autonomous agents, options and choices, the rigid authoritarian structure of the university system and academia is no longer appropriate, and if we are to embrace our goals of liberty and diversity, then the new technology – which is designed for that purpose – should be employed to those ends, not rejected out of hand. This is a theme I return to a lot – the use of new technology to accomplish objectives that make sense only in a world of universals and hierarchies.Anyhow, if this article seems a bit strident, it is because I am expressing frustration about dealing with the old issues well after having engaged the new issues considered here.

7 Cutting the Throat of the University

Written August 7, 2001. Published as Unrest in the Ivory Tower: Privatization of the University, USDLA Journal, October, 2001.

Academics must resist the trend toward the commodification of education, claims Steve Eskow, or universities will become privatized. On the contrary: the more professors resist, the greater the liklihood that privatization will happen, and that would be a tragedy.

7.1 Introduction: Mensa and Academia

I once had a desire to join Mensa. I'm bright enough; their IQ tests are pleasant diversions but no real challenge. And I enjoy hanging around with bright people, as I have for the two decades I've spent in an academic environment, quaffing a few fine ales at Dewey's or Dinnies Den, debating matters far and wide of varying degrees of importance.

As I learned more about Mensa, however, disillusionment set in. While one would have thought that society's brightest minds were focussed on the pressing issues of the day, these minds were focussed most of all on puns, word games and clever tricks. By comparison, even my regular trivia games have more merit. And my sudsy sermons Pulitizer prose.

Over time I have become less enamoured of the university environment for similar reasons. Not that universities even approach the banality of Mensa; the people across the road continue to amaze (islet transplants, fun with phage cancer treatments, human livers in mice...) and the university is recognized – here, at least – as the city's key economic engine. But university professors can and do obsess over the minute. They can put their own momentary comfort over the needs of academia and society. And they can be as self-absorbed as the most narcissistic Mensa meeting.

I say this lovingly, of course. Nobody spends two decades associated with institutions and people they despise (or even dislike). Twenty-first century academia is a treasure, one of humanity's shining pillars of achievement. It is worth saving, or at least, spending a few hours on a Tuesday morning talking about how it may be saved.

7.2 Eskow's Hypothesis

I now turn to Steve Eskow's interesting words:

> Hypothesis: There is a growing movement afoot in the US and elsewhere to use

> distance education as one of the knives to achieve the dismembering, and the death, of the university. It is often unconscious, as in John Hibbs's quote, but one does not have to be a Freudian or a literary critic to detect one of the organizing patterns of this death wish.
>
> It is often disguised--hidden from the speaker or writer--as a desire to "improve" the university, to make education "more affordable" and "more efficient." I'll call it THE PRIVATIZATION SCENARIO.
>
> It is a pattern, easily detectable, in that it plays variations on the theme of privatizing education on the grounds that the "market" is a more efficient guarantor of quality than the "elite" guild of academics who are more interested in protecting their own turf, etc., etc.

How typical that Eskow's own words indict him. I most certainly agree with his hypothesis: that there is a movement toward privatization. But if Eskow would poke his head beyond the ivy-lined campus windows, he would see that the privatization movement encompasses all of learning, not only universities. The move toward charter schools, home schooling and various alternative education projects highlights this trend in the elementary sector. Trade schools and colleges face increasing competition from private institutions.

Moreover, it is not only the institutions of learning that are being privatized. Their product: the books, journals, ideas and opinions produced by professors and their ilk are being increasingly placed under corporate lock and key, whether they be through funded research or collected in fee-based archives such as XanEdu. Patents and copyrights are moving the learning that used to be freely circulated in the public domain into a closed marketplace of privatized knowledge.

Universities and especially university professors are easy targets precisely because, like Mensa members, they become self-absorbed. Part of that comes with the territory – you cannot be expert at anything unless you become a little fanatical – but part of it comes from a blindness, an inability or unwillingness to look at some wider trends sweeping society, trends that have the potential to sweep the university system with them.

So let's subsume the 'privatization hypothesis' under this larger picture, the one in which human knowledge itself is being privatized.

7.3 Why Defend Universities?

As I mentioned above, I am a defender of the university. Perhaps you may not believe that, given my staunch defense of distance and online learning, and given my occasional carping about universities and university professors. But I am a defender of the university because I am a defender of knowledge, and in particular, that view of knowledge where it is a public trust, intended and to be used for the benefit of all of humanity, freely shared and freely used.

If we were talking about money, not knowledge, I would be classed as a socialist, perhaps even a Marxist or communist. I am not sure whether there is a corresponding term for the public ownership and free distribution of intellectual capital (I may as well take yet another stab at historical significance and call it Downesism). Whatever it is, it is that that I support; my support for the universities is as a means to this end.

This is important: universities are not worth defending in and of themselves. They are worth defending only insofar as they foster the free distribution of knowledge, whether it be by means of allowing people an affordable education,

whether it be by means of discovering and announcing fundamental truths, or whether it be by means of advancing our science, technology and human sciences for the good of society as a whole.

Knowledge is different from capital, and from material goods, in that there is no inherent scarcity to knowledge. A piece of knowledge, once produced, may be replicated almost for free, distributed around the world in the blink of an electron, fed almost as easily to one person as to one billion people. Oh sure, there are some pragmatic issues: knowledge can be expensive to create, and as those of use involved in distance and online learning will attest, distribution is not free. But for the greater good people in a society – and across societies, in a global society – pool our resources, funding public universities for the production of knowledge, and a public education system for the distribution of knowledge.

We allow and accept a market system for the distribution of knowledge where it is appropriate. We recognize that a person owns his or her own ideas, and that the inventors of new technologies have the right to profit from their work. We allow that money may be exchanged for knowledge. So long as the objective – the widespread creation and distribution of knowledge – is met, we can allow a multiplicity of methodologies. And just so society today has created great public universities, great private universities, public-private collaboration, government sponsored research, and corporate research. When we look at the intellectual achievements of the twentieth century, we regard not only M.I.T, Harvard and Stanford, but also Xerox PARC and Texas Instruments, NASA and National Geographic.

Now Steve Eskow is concerned about the privatization of universities. He argues that the privatization of universities is being accomplished via a set of processes and paradigms that I will look at below. And so it is true: these processes and paradigms are being used as the thin edge of the wedge by those who would privatize universities, and indeed, privatize knowledge generally.

But: these processes and paradigms only accomplish the goal of privatization if they are effective. Were they not effective, they would not be a danger to universities at all. Nobody is trying to privatize universities by means of beer sales or fox hunting competitions, because there is no great demand from the public for university beer sales or fox hunting competitions. The people who are advocating privatization are hitting the universities where it hurts: and they are appealing to society's larger objectives in an effort to transform the university system.

They are aided and abetted by those who resist many of these changes, for while many of these changes would result is an improved educational system for all, the reluctance of public universities to adopt them is by itself the single greatest cause of the privatization of universities. University professors, by taking the narrow, self-serving view, hasten their own demise.

7.4 The Means of Production

Let's look in some detail at this:

Eskow continues,

> Here are some detectable pieces of the pattern. (Not all who are impelled to destroy the university subscribe to all of these, and there are those who subscribe to some of them who do not wish the university harm.)
>
> 1. The Three M's--"Massification; "Marketization;"Managerialization" These

may be self-evident; I'd like to write more about these later

When an academic writes that something is "self-evident," it almost certainly is not. But I digress.

7.4.1 A. Massification

Massification, in my understanding, is the employment of the instruments of mass production for the development and distribution of knowledge and learning.

People entering a contemporary research lab would be astonished at the degree of massification already in progress. Modern medical labs, for example, resemble production centres much more than they do Thomas Edison's garage. Teams of scientists, following strict protocols, work in assembly to synthesize, test and produce thousands of compounds. The sequencing of the human genome was possible only through mass techniques. Such researchers also use the means of the mass to disseminate their knowledge: journals are mass produced and shipped to every corner of the globe where identical scientists in identical labs reporduce their discoveries. Scientific progress is no possible without massification.

Only in the field of teaching does academia seem to have successfully resisted massification. Only in the field of teaching is the product the result of the individual craftsman, toiling alone, each bit of lecture a custom fit for the small group of students assembled before him. It is a source of continual frustration to society as a whole – why can't we devise a means of reaching everyone, and not only those favoured few with the time and money to spend attending university lectures? And when we look at the challenge or providing a university-level education to a global population of 7 billion and more, it becomes obvious that teaching must evolve. Were cars hand-made, only a fortunate and wealthy few would have them. The same is true for education.

A profession that insists that all its products must be hand-crafted dooms itself to oblivion. As long as university professors assert that the only form of teaching must be the in-person lecture they are hastening the development of non-university alternatives that prove otherwise.

7.4.2 B. Marketization

Marketization is, in my understanding, the treatment of education and learning as a commodity, to be displayed and selected for consumption by a paying public. Marketization (and not online learning per se) is the major objection David Noble offers in his critiques of distance and online learning.

Defenders of the university may then be surprised to hear me defend marketization. I have even written (half-written) a paper called "The Learning Marketplace." Why would I do so, if marketization is so contrary to the university culture?

The fact that it is contrary to the university culture is why the paper had to be written, but I have no intention, subconscious or otherwise, of thereby dismantling the university system. Quite the contrary, in my view, marketization may be the salvation of much of the university as we know it today.

Private enterprise theorists, as Eskow comments elsewhere, often argue that the market is the most effecient way to distribute a resource. Universities have steadfastly resisted that doctrine, maintaining instead an monopoly and control over the distribution of knowledge, reserving it either for their peers or for the select few who attend university classes. But clearly there is some evidence, is there not, that markets do provide an effective means of distribution? Otherwise we would not have grocery stores, we would have government food outlets. Otherwise

we would not have restaurants, we would have government eating stations.

Markets work on the principle that the exercise of choice is more efficient that the exercise of control. The reason for this should be obvious: people are much more willing to decide for themselves what they want that to have it decided for them. Moreover, when somebody must decide for them, there is an increased liklihood that they will make incorrect decisions. As John Stuart Mill famously observed, the best indication in a society that something is valued is that people value it. The best indication that something is good for people is that people desire it as a good.

Where market theorists err is in their slavish adherence to the principles of the marketplace in all times and in all contexts. But marketplaces are known to fail, as anybody buying lumber in Flrida in the wake of Hurricane Andrew can attest. Markets work only if there is a sufficient supply of a commodity. Choice is only efficient where choice may be effectively practiced. Where choices are forced, where commodities are in short supply, the markplace collapses in on itself, spiraling out of control, rewarding the rich and powerful and leaving the mass without.

When something is in short supply, a call for the marketplace to distribute that good can (and should) be seen as folly: for the advantage rests entirely with the distributors, and none with the consumers. Thus it may seem that having the market distribute education may be called a folly, because education is, as anyone can see, in short supply. People today spend the equivalent of a price of a small house for a university education. I saw recently five day courses offered by Queens at a price that would buy a small car. Putting education into the marketplace in such conditions would be folly: it would be licensing the owners of knowledge to print money, and condeming the vast majority of humanity to doing without.

But there is no reason why learning must be a scarce commodity. Indeed, it is arguable that it is a scarce commodity only because universities and university professors have created a false scarcity. It is as though the news media of the world decided that the only way people could really understand the Balkan conflict would be to hear about it in person from a professional journalist. The result of such folly would be evident: people would pay thousands of dollars to listen to average journalists (not everyone can afford a Cronkite) while the vast majority would have no access to this information at all.

There is no reason why education must be scarce, and every reason why it can be produced in mass quantities for mass consumption. And in such an environment, there is no reason why learning cannot be distributed via a marketplace, and every reason why it should. For the best indication that something needs to be learned, as Mill would say, is that people want to learn it.

7.4.3 C. Managerialization

Manegerialization is, to my understanding, the process whereby an academic relinquishes some control over the production and distribution of knowledge to a team and where that team is run, not by the academic, but by a manager. The manager, of course, knowing nothing about the subject in question, can be relied on to make poor decisions.

As a sometime software designer, I am certainly sympathetic to this line of reasoning. Indeed, an entire culture – the Dilbert Culture – has developed in the software community to make fun of the pointy-haired bosses who think they have

some understanding of software design. I have no doubt that the same is true in other areas of endeavour, and were I promoted to coordinate the design of, say, a learning project in the field of microbiology, my academic interference would be as welcome as a focus-group expert at a hacking convention.

The problem, of course, is not the practice of employing teams to develop learning material: the problem is pointy-haired bosses. In the software industry, almost nothing is created outside a team. Even some of the most heralded individual achievements – unix, say, or Linux – have over time become the project of dozens, even hundreds, of dedicated individuals, each person working on their own area of expertise, suffering the indignities of more or less coordination by a manager. Indeed, looking at the wider world, only professors, it seems, have the wherewithall to resist working as part of a team, so much so that the term 'Lone Wolf' has been coined to characterize much of academic endeavour.

And, of course, no professor (or very few, at any rate, since I obviously count myself as one of the exceptions) has the expertise to professionally provide all aspects of educational delivery. It is no wonder professors say that the best and only means of teaching is in-class and in-person: no professor has the skills or the time to do anything else! But by their own dognatic adherence to individualistic 'lone wolf' production methods, they make their own prediction a self-fulfilling prophecy.

A prophecy, moreover, which is demonstrably false. Teams of people working in unison in other fields have managed educational attainment far beyond that of any individual professor. Hundreds of millions of people could reliably create a Big Mac (two all-beef patties, special sauce, lettuce, cheese and a pickle on a sesame seed bun). Millions more could state with conviction why a 2-5-5 defense is ineffective in a football game, analyse the comparative merits of Randy Johnson and Nolan Ryan, comment knowledgeably on the weather, sing a Beatles song and play a round of golf (correctly, within the rules, though perhaps not well).

And: so long as professors refuse to work as part of top-flight educational teams, more or less competently managed, their achievements will be eclipsed, over time, by teams of skilled professionals producing top-flight educational materials. And when professors, teaching alone in a classroom, are widely recognized as an inferior (not to mention expensive) form of education, the call for privatized education will take full flight.

Massification, Marketization, Managerialization – to the degree professors resist these, rather than embrace them, they are hastening their own demise. It seems to me that the best minds in society could find ways to make the 3Ms work for all of society – but instead they sit in their little offices, careless of society, wondering how their pleas of 'quality' can be possibly relevant ot the many millions of people who never shadow their hallowed walls.

7.5 Language, Truth and Logic

Eskow continues,

> 2. Changing the Rhetoric of Education
>
> Students as "customers"; the college as a trading company, importing and exporting; "standards" that can be "measured"; "brokers"; and, of course, "productivity," "efficiency", "accountability." And: "quality control."
>
> Profound clue: John Chambers of Cisco popularized the notion now a cliché in the forprofit community that "education is the next 'killer app'." One

> doesn’t have to be Freud or Jung to see the implications of both "killer" and "app." Or to sense the possibility that one of the things that has to be killed by the app is the university.

As any linguist knows, the words we use are used to reflect reality, either as it is or as we would like it to be. Thus, words such as ’phlogiston’ fell into disuse as our concept of reality came to encompass oxygen, and our use of the word ’girl’ declined, with much encouragement, as a reflection of our desire to minimize the diminuation of women. A vocabulary is like a mirror into a person’s world view: words express meaning, meaning expresses reality, either shared or solopsistic, either faithful or fancied.

The words Eskow lists fall mostly into the category of representing the world as we wish it to be, a fact he seizes upon to assert that their use reflects a hidden agenda. For any person without effort can find instances which prove that the university system is not, as he suggests, customer (or learner) centered, efficient, effective, or productive. The typical university lecture does not adhere to any standards (at least none that I can detect), learning is measured only in the crudest of fashions, and professors – the bearers of ultimate job security – are certainly not accountable.

Your words, not mine.

Eskow quotes John Chambers as describing learning as the next ’killer app,’ implying that university education is what will be killed by some new technology. Perhaps so. It is worth noting that the term ’killer app’ was devised, not merely because it was fatal to some preceeding category of products, but because it was widely used, wildly popular, and became a paradigm for the applications that followed.

Mosaic – later Netscape – became a killer app, popularizing the World Wide Web and the internet in general because it bucked conventional (and I might add, professorial) wisdom, by allowing people to view graphics. People familiar with the history of the internet are familiar with its academic origin: and such people sometimes cynically say that only university professors thought that pictures and graphics would not be needed for online communication.

Email flourished as a killer app because it replaced an outmoded and inefficient organization: the post office. Today the flow of messages by email far exceeds the capacity of the post office. The writing of messages on paper, the placing of paper in envelops, the procurement of tariff stickers (called stamps), the trek to the post office box, the wait while the physical package is collected, sorted and distributed (by foot, no less) – all this was a technology waiting toi be superseded by a more efficient, productive (and dare I say, standards-based) replacement.

I have heard the lament more often than I care that the web has produced a wealth of poor graphical design and that email has produced an endless supply of drivel: and perhaps it has, from people who never engaged in graphic design before the advent of the web, and from people who never set pen to paper when mail was a combersome task. And the same critics overlook the awkward design of most publications in print (not everything is National Geographic or the National Post) and the steady deluge of junk mail that flows, even today, into our mailboxes. Much less the time and cost of producing pens, paper, evelops, stamps and a worldwide pedestrian delivery system.

The fact is, killer apps become killer apps because they’re better, and so when

John Chambers suggests that online learning will become the next killer app, it is because he thinks it will be better – much better – than the contemporary pedestrian product.

And how might it be better? The new vocabulary – used not only by potential privateers but also by people genuinely interested in education – tells the story.

7.5.1 A. Choice

Students as 'customers' – or in the more common parlance of educators, 'student-centered learning' or even 'learner-centered learning' – a reflection of the desire to create a system where universities exist to serve students' needs, and not professors' needs. This does not (necessarily) reflect a 'customer-is-always-right' attitude – as any patron of McDonalds will tell you, the customer is often perceived as wrong (you get a pickle whether you like it or not). But it does reflect an understanding and even an ethos that the purpose of the institution is to provide students – the customer – with what they want (not to mention, paid for).

A lot flows from that assumption, but I will key in on one thing which encapsules the difficulty contemporary universities have with the student-centered approach: choice.

Aside from some very broad choices (will I study engineering or philosophy?) students have very few choices in a university. Having selected a program, they are routed to a faculty, given a small selection of options and a bevy of required courses, and are assigned professors (if they are lucky, they will learn about and manage to avoid the particularly bad professors). Inside the classroom, they have very little choice about the course content, nature and number of assignments, criteria for passing, time and place of course offerings, labs, workshops or seminars. They have no choice at all regarding their classmates, limited choice in assigned texts and readings, and are unified in their quest for a single (obligatory) goal, the university degree.

No doubt all of these decisions are made for the benefit of students. Sometimes – often, actually – these decisions will in fact be correct decisions. It is a nice healthy line-up of educational nutrition. But imagine a grocery store where, once you have selected your food type (Italian, Chinese, Indian), you are routed through a certain set of aisles. You are given one or two of each product item to choose from, and you have a set of required products you must purchase. You are required to show that you are able to prepare the food correctly before you leave, but you must prepare it in a certain way 9depending on the whims of the cashier). You will buy – and only buy – a full year's worth of food. No doubt many of these decisions regarding food and nutrition are correct decisions, but the experience is entirely unsatisfying, and to a diabetic, fatal.

It's a simple thing, choice. Yet if John Chambers can develop an application that provides educational choice, the killing fields will be littered with ivy-covered rubble.

7.5.2 B. Standards

Standards – of course university professors are notorious for resisting standards, at least so far as the practice of their profession is concerned. This has the result of creating frantic student consultations in the halls and campus pubs in a determined effort to avoid the notoriously bad professors (my own experience cannot be that unique, can it?). It is difficult even within a single institution to determine what constitutes a first year logic class, let alone to determine

this across a nation (much less world-wide).

In no other field is such a crass disregard for the nature and quality of the component parts of a product or service so brazenly displayed. Those very academics who rail against standards would be appalled were they to learn that the airplane they are flying was assembled, ad hoc (no doubt by a team of skille craftsmen) without regard to wiring, fuel or avaiation standards. They would not dare drive were they to learn that the reliability of their tires was not proven. They would not eat food that may or may not contain arsenic (much less peanuts), would not drink water which could not meet certain criteria of safety. They expect that the wiring in their home will not only be up to standard, but also that it will be inspected by a third party to make sure. Yet in this, one of the most important investments of time and money a person can make, they expect to fly without standards.

I personally see no reason why there cannot be a 'standard' logic 101 in use world wide. The principles are fairly well understood and have been accepted without a significant change for the last two throusand years. A common base of examples exists and makes the rounds in any case. Tried and true techniques for teaching reasoning – from Venn diagrams to truth tables – exist. Yet there is no such thing, and no concept of what would constitute quality teaching of logic, and successful learning of logic. Except, I should add, for the innovation of a private standards-based test in logic, which is used only grudgingly (if at all) by academics (but most enthusiastically by people who teach logic online – what a surprise).

7.5.3 C. Efficiency

Efficency – it makes no sense to have a highly skilled teacher spend half his or her time producing mediocre research so that he or she can get tenure or promotions. It makes no sense having a highly skilled researcher teach a class in order for him or her to keep his or her job. It makes no sense for either teacher or researcher to sit in from of a class while a test is being conducted, languish in the back of the room while a video is being shown, spends hours debating parking policies at a faculty staff meeting, and more. And when you have a hundred million graduate students to teach, then it makes no sense having eight students in a graduate class, no matter how good the exercise, becasue it means that most of those students will receive no graduate education at all. I'm surprised they haven't taken to the streets. I have only picked up on a few examples here, but it seems clear and obvious: if Cisco could produce an online learning system that was learner-centered, standards based, efficient, productive, and accountable, then people would abandon universities in droves, and more to the point, governments would be very hard pressed to justify spending a lot of money on the public system when the private system is doing the same job for more people and for less money. And even more to the point: we are already beginning to see signs of this today. Recently, DeVries was given accreditation in Alberta. This means that this private institution is now competing on a level playing field with the publicly funded colleges in our province. Should they prove more popular and more effective, our government will not be able to justify spending money on demonstrably inferior and more expensive alternatives. In Pensylvannia recently, a charter school offered classes online – thereby drawing the ire of proponents of the traditional public system but the praise of parents who found this particular alternative a giant leap forward in ease of use and efficiency. To the degree that universities and university professors drag their feet in becoming student-centered, efficient,

standards-based, accountable, and the like, that is the degree they are cutting the slender branch on which they all rest.

7.6 Quality and Control

Eskow continues,

> 3. Changing the Institutional Structures of Quality and "Control"
>
> Not "peer review" in the tradition of the professions, but "quality control" in the tradition of the factory system. Note John's models in his message" ISO 9000, Dunn and Bradstreet. Business organizations, industrial organizations as models for the university to emulate.

Eskow's loaded terminology displays his distaste for factories, accountants and perhaps industry as a whole. As I suggested above, I sincerely doubt that he would fly in an airplane evaluated solely by peer review, but that points not so much to the silliness of his argument as it does to a mistunderstanding of evaluation and review (and yet this guy is marking student paper... one wonders...)

Let me talk briefly about ISO and the 'quality' movement in general. What we have here is actually several things combined and sold as a package (as such it is a deeply flawed package, but it contains enough that is good to be marketable):

First, it embodies the idea that quality can be measured, and

Second, it embodies a business ethos which asserts that quality can always be improved, and

Third, it establishes a team-based structure of quality circles in order to impel and enforce these quality improvements

When I think about 'quality', my mind always to a picture printed about a decade ago in the Globe and Mail's Report on Business magazine (a nice, glossy, short-lived tribute to the corporate way) of a group of young and earnest looking Japanese workers, seated around a table, called the 'Paddington Bears,' whose sole objective in life (so the caption went) was to reduce the number of scratches in TV monitors from 8 per million to 1.

Now: reducing the number of scratches in TV monitors is good. We would complain if we bought a scratched TV, and we would complain if the cost of TVs were doubled because every second monitor must be discarded. But: spending all day reducing the number from 8 to 1 per million is foolish; and making it the basis of society is ridiculous.

What we want to do here is separate the concept of quality from the corporate ethos in which it has been packaged and marketed as 'total quality'. We want to keep the good: airplanes that fly reliably, food that is safe, water that is potable, education that is effective. And we want to discard the bad: individual subsumation to the wants and needs of the corporate entity, to the exclusion of all else.

Eskow, in deliberately conflating those three components of the quality movement, does his readers – and education in general – a disservice.

Focusing on quality only, we need to distinguish two types of quality. I have in previous emails referred to these as 'semantic' and 'syntactic' quality. One might think of them as 'qualitative' and 'quantitative' quality respectively. But I prefer 'semantic' and 'syntactic' to get away from the idea that the former consist only in touchy-feely emotions and that the latter consists only in cold-hearder mathematical calculations.

Now in the evaluation of student work, professors employ both forms of assessment on a regular basis. In the syntactic mode, they assess whether the student has his dates right, her facts straight, has correctly parsed a sentence, correctly applied a proof, use appropriate symbols in an engineering diagram, written a program that compiles, quoted Mill correctly, or successfully identified Shakespeare as English. In the semantic mode, they assess whether a historical description captures the mood of the times, whether a recitation of facts is relevant, whether a sentence flows, whether a proof is elegant, whether a diagram is neat and illustrative, whether a program is easy to use, whether Mill makes sense and whether Shakespeare's English is understood in context.

Obviously, no assessment of student work is complete without both the semantic and the syntactic mode of evaluation (though teachers are often criticized for ignoring grammar and spelling, even accuracy, in student essays, searching for that soft and fuzzy 'meaning' underlying the garbled scribble they see before them). So also it is with university instruction. Ignoring the syntactic misses the question of whether they are learning at all; ignoring the semantic ignores the question of how well they are learning. Ignoring the syntactic misses the question of whether a journal article follows correct procedure; ignoring the semantic ignores the question of whether it says anything worth reading. Two forms of assessment: and necessarily, two forms of evaluation.

Now the kicker: academics' evaluation of themselves, insofar as it occurs at all, is almost entirely semantic. Or to put the same point another way: there is almost no standards-based measurement of an academic's performance (except, perhaps, for adherence to the all-powerful (and mis-applied) bell curve).

Eskow identifies the 'peer review' as the traditional mode of academic evaluation. No doubt it is traditional, and widely practiced. But it is only half of a reasonable evaluation, and not even a very good half at that.

In my country, and no doubt in many others, we have a polite fiction called 'trial by your peers.' The idea is that in a jury trial, guilt or innocence will be determined by a panel of citizens similar to yourself. As I say, it's a polite fiction. I recently discovered that in Alberta (perhaps elsewhere), potential juries are selected from the set of people who have driver's licenses. This explains why I have never been selected for jury duty: I don't drive. But it also de-selects a certain, lower, stratum of society (one, oddly, corresponding with the set of 'peers' of many a convict, but I digress). Similar selection practices in other contries demonstrate a similar bias: selecting juries from the list of registered voters, for example, de-selects those people who, for one reason or another, are not registered to vote. Again, the weighting here is toward the upper stratum of society.

But there's more. When a particular individual is brought before the court, both the prosecution and the defense have the right to veto a certain number of jurors. Any number of criteria come into play: people are disqualified because of their race, gender, occupation, residence, and more. Often, they are disqualified because of their opinions. Because I am an opponent of the death penalty, for example, I would never be selected as a juror in a capital case in the United States (so I understand, anyways). Naturally, this predisposes the jury toward a panel that will opt for the death penality in such cases.

I have long wondered why gang members, homeless people, and other social outcasts never seem to be selected for juries. But of course, it's because the concept of 'trial by your peers' is a fiction. It really means, 'trial by your betters'. Or

at the very least, ’trial by people who think in the right sort of way.’

In popular opinion at least – and I am of the same view – the reliability of jury trials is questionable. Since the not so recent OJ trial, or the less vividly remember Klaus von Bulow trial, people have come to see jury trials not so much of an exercise in justice as in manipulation. Social activists will reel off a list of people wrongly convicted by juries on the scantiest of evidence. Jury trials, at least some of the time, are much less an exercise in justice than of prejudice.

Now imagine the same system, but without any standards at all. Without rule of law, to guide guity and innocent verdicts and appropriate penalties. Without rules of evidence to distinguish fact from fiction from hearsay. With no limits whatsoever on the biases, prejudices or qualifications of jurors. In essence, mob rule, with none of the standards that today (sort of) protect the innocent from wrongful incarceration, the guilty from dangerous liberty.

Such – in the academic world – is the essence of Steve Eskow’s position. Is it any wonder it draws a society-wide roll of the eyes?

’Peer review’ in academia is no such thing. Otherwise, we would see graduate students and even interested laypeople on academic review committees. No, journal review boards especially are populated with the academic elite, those whose publications and scholarly presentations have established their authority in the field. Nor is their selection random: constructivist journals do not select rabid anti-constructivists to review articles, Marxist journals do not recruit people from the Fraser Institute to edit their publications.

The actual review is secretive and closed-door. Nobody knows what process of reasoning, if any, occurs when professors are evaluating a colleague’s work. The results, at least from the eyes of the layperson, are less than impressive: reams of dime-a-dozen articles in unread academic journals, arcane dissertation topics suitable especially for ridicule by the national newspaper, forgotten theses read by an audience of three (and here I think of my own unlamented "Models and Modality"). Authors do not even know who their reviewers are, much less whether they are peers in any meaningful sense of the word. And woe betide the author who is not willing to acknowledge duly established Authority. ’Trial by people who think in the right sort of way’ indeed.

At least in the case of journal articles and publications, peer reviewers at least (we think) read the works they are reviewing. No such exposure to the actual product being reviewed occurs in the case of teaching. It is folly – and rightfully recognized as such – to dub the review of a professor by peers who have never seen him teach some sort of assessment. Such a review has everything to do with how the itinerant behaves in the Faculty Club and nothing to do with the sort of education he or she has left behind in the mind of his or her students.

Peer review has its place, as does any sort of qualitative assessment, but to make it the sole – even the primary – determinant of academic merit is beyond foolishness. It creates – quite rightly – in the minds of the public the image of a self-serving cadre of Old Boys who all think they are wonderful and who collectively exhibit wisdom so great that the word ’genius’ is an insult and a slur. Closed-door self-evaluation is as reliable in academia as it is in the airline industry or the food processing industry, which is to say, not effective at all.

Academia would do well to open its system of assessment and review to (a) quantifiable standards, and (b) an open review process. Something like a system of

standards – call them learning objectives, performance outcomes, whatever you will – should apply to graduates of a given class. Society should be able to know, without having to take Jo Blogg's word, that an A in logic 101 means that the student can recognize some basic logical fallacies and can string together a simple argument. That's not so hard: and there's even a standardized test for critical thinking.

And there is no reason why academic performance cannot be the subject of open and public review. There is no reason to restrict readers to a panel of three mysterious experts: works up for review should be publicly viewable and reviewed by anybody who cares to read them. Journals may even rely on those very reviewers, but the publication of a poor article by even a good journal will be widely recognized as such. And there is no reason why students cannot evaluate professors, and if the results cannot be posted on a website, then students should at least have the option of expressing their views by taking the same course from another professor or even another institution.

None of this infringes on the professor's ability to do as he or she sees fit: however, when a review process exposes poor and shoddy work, as it inevitably will, such perpetrators will invariably be held to account. Which is as it should be: in academia as much as in airplanes.

7.7 The Implications of Re-Form

Let's examine how Eskow concludes before concluding ourselves:

> John and his colleagues are in earnest. They want to Privatize, Massify, Marketize, Managerialize education. They want to change its vocabulary, and the new vocabulary brings with it the practices of industry--not the practices of the "new economy" or the "postindustrial society," but the older notions embodies in TQM and ISO 9000. Seonc Wave, if you will, not Third Wave. Each of us, I think, has to decide if this reading has merit, and if this is the intent of the new rhetoric and the new directions: conscious or unconscious intent. And perhaps more importantly, each of us has to decide what our position is on this drive to reform education-- re-form it in the image of business and industry. I think the future of distance education depends on whether this movement succeeds or fails, and I've chosen my side. And my hope is that those like John who do not yet see the full implications of their program of "re-form" will come to see things differently. They can become powerful allies for securing and improving the university, instead of dissecting and dissolving it.

As a learned academic, Steve Eskow should know better to frame the argument in such a false dilemma. He should know better than to use such loaded terminology, calculated as much to inflame as to argue. His are the tactics of a sourthern lawyer arguing toward a carefully preselected and predisposed jury: he wants to paint all advocate with the same brush, and he is not above quoting some carefully selected freudian mythology in order to drive his point home. Steve Eskow would have you believe that if I support John Hibbs on thsi point then I am the same as the Great satan, the corporate sellout, the soulless butcher who would cut the throat of a fair institution in a minute if only given the chance.

But: those people who are persuaded by Eskow's crusade are hastening that very act of homicide, sure as the Sun rises in the east. By perpetuating the idea that any change in academia is a knife in its back, Eskow is freezing the university system into an unsustainable stasis, ensuring that even the slightest attempt from the

corporate side of the house will be successful.

It is interesting – ironic, even – that Eskow paints two divides: the collegial university system, on the one hand, and the cold, calculating world of business and industry on the other. But there are not two solitudes here, there is only one; were Eskow to look about society around him, he would find that all manner of enterprises follow the dicta of client service, accountability, efficiency and relaince on standards. Not only industry, but sports, recreational travel, home repair, cooking, amateur astronomy... absolutely, utterly everything but education (and perhaps some handmade wooden crafts shops).

It turns out, in the wider world, that people do not want to spend their time and money (a) meeting someone else's needs, (b) paying for work that doesn't need to be done, (c) not knowing the results, (d) not knowing what is being produced, and (e) is more than they can afford anyways. If this is the picture of academia that Eskow is defending, then it is doomed, and if by falling it must fall into corporate hands, then Eskow's own logic has as its inevitable consequence the privatization of education.

And that would be a bad thing: but not simply because some academics don't like it (and not simply because it doesn't meet their arbitrary standards of quality (whatever that is (because, remember, they are opposed to standards)).

7.8 Intellectual Wealth and Society

At the beginning of this treatise I spoke of the privitization of knowledge. I would like to say here that if the university system (and the public education system in general) fails, then this will result in the privatization of knowledge. Now even that is in itself not a bad thing – I have already acknowledged that there ought to be latitude for ownership of knowledge, whether it be by virtue of copyright on an essay, ownership of a petent on an invention, or some similar claim to intellectual or emotional property.

But the market economy, as I also suggested, works only if there is an adequate supply of the commodity in question. Once a scarcity is achieved, the market breaks down: we move into a monopoly (or duopoly, etc) mode in which prices rise all out of proportion to the value of the commodity and in which a substantial portion of the population is forced to do without.

With the rise of the information economy we have seen not only a concerted attempt to privatize knowledge but concordantly an effort to create artificial shortages in knowledge. Where once books circulated freely, were shared and loaned, read by the thousands in libraries, sometimes photocopied, sometimes transcribed by hand, there is today a movement afoot to create the single-use book, an entity that may be viewed but never reproduced nor shared nor copied in any form. Where once academics freely circulated copies of their article abstracts, exchanged ideas at conferences and conventions, today we see sponsored research, per-user subscriptions to e-journals, non-disclosure agreements, and more.

Clearly this is damaging to the intellectual wealth of society as a whole, because not everybody can afford to pay $24.95 for each knowledge-product per annum, much less amass a permant and useful library of e-readings. Where once we could at least alleviate some of the strife in developing nations by sending them books and magazines, today we are told that such action constitutes a violation of copyright – it is not even legal to load our used copy of Windows 3.1 on used computers to send to East Timor, as some Australians found out.

But it is damaging also because it limits the voices we can hear. Just as top 40

radio streams consumers into a megastar mentality, so also dissenting voices disappear when knowledge is controlled by corporations and dispensed in pre-approved (and costly) allotments. We are all too aware of the Russian programmer, recently arrested in the United States for writing forbidden software, or the professor in (as I recall) Princeton who was ordered not to publish a decryption algorithm. But it is much more pernicious and much deeper than that. It is the expulsion of a boy who wore a Pepsi shirt to 'Coke Day' at his school. The forced apology and expulsion of a student who dissed a corporate sponsor.

The World Championships in Athletics are being held here in Edmonton as I write. The championships are sponsored by (in part), and thereby essentially owned by (in part), Nike. As a columnist in the Edmonton Journal observed, Nike's influence is pervasive: and at a press conference in which a renowned anti-doping athlete was asked to comment on the reinstatement of a competitor, the Nike spokesman intervened to assert that athletes would not be answering questions about doping. Bad for the image, you see.

We tend to think that the corporate control of information is about big things, like freedom of speech and the right to protest: and it is. But it is manifest in a deluge of little things, and bit by bit, our knowledge and our freedom are slowly eroded. And we're back to being the Paddington bears, not merely because we cannot utter any opposition to this ethos, but becasue we cannot conceive of one.

The fall of public education in this country and in this would would be a disaster of the greatest magnitude, resulting in the descent of a corporate curtain of ignorance. Failing to move, failing to respond to the need for a greater, more vigorous system of public knowledge than ever, is to silently, stupidly, acquiesce.

I cannot believe that educators today are not knowledge guerillas, silently and steathfully subverting through the covert education of as many people by whatever means possible. I cannot believe that academics today steadfastly defend their bastions of privilege, ignorant of the fact that they castles they so rigorously fortify will defend a totalarian regime that will upset their tottering rule. Academics must, in order to survive at all, obtain the support of the people, but they will not do this if they withhold from people the one thing they value.

I know that there are many open-source open-content academics in the community working hard to stem the advance. It is a race against time, creating public domain knowledge management systems, public domain encyclopedias, courseware, almanacs, maps and illustrations, literacy guides, media readers, free textbooks, trying like townspeople in the face of the invading army to hide as much of the community chest before the hordes descend to lay claim to everything they see. Hide, hide the knowledge where they'll never think to look for it: among the people.

Academics who defend their privilege in an Eskowish manner are like those who, citing the long standing tradition of ownership and privilege, sit on their treasure, thereby safeguarding it for the arrival of the invaders.

I really think that universities best protect themselves by doing the one thing they can do better than corporations: producing and distibuting knowledge. But they must do it in such a way that it remains better than the corporate alternative. This means mass education. This means a marketplace of educational opportunities. This means top-flight educational resources produced by teams of experts. This means a student focus. This means efficiency, accountability and

productivity. This means open standards and open evaluations. This means, above all, reform.

Academics are at the crossroads. They could, collectively, use new technology and new techniques to produce a flowering of human intellect the like of which has never been imagined. Or they can hunker down, cling to their privelege, and usher in the twenty-first century equivalent of wage labour and cut-throat knowledge capitalism.

Remember, the chains you most fear are the chains you forge yourselves.

I wrote this item in response to a request for an article, and so it conforms more to length restrictions than to any particular thread of discussion. It is included here because it expands on the idea of what I called in my introduction 'epistemological quanta' – the idea that what was once a monolithic thing – education – is breaking apart. And I try to argue, with reference to Minsky, that once we enter such a world the logic of connectionism and emergent properties begins to take hold. What we are seeing, I write, is learning as it evolves from a centrally controlled distribution of information from an expert to a consumer to an interactive, dynamic, user controlled set of information exchanges.

8 The Fragmentation of Learning

Written October 22, 2001. Published in Education Canada, Volume 41, No. 3, Fall, 2001, pp. 4-7.

In his groundbreaking work Marvin Minsky proposed that human intelligence is derived from what he called the Society of Mind (Minsky, 1988). The idea behind Minsky's theory is that the human mind consists of millions of task-oriented agents characterized by two major features: autonomy and ignorance. Autonomous in the sense that agents make decisions by themselves without direction from a central authority. And ignorant in the sense that agents never see the whole picture: they make their decision based on a limited set of inputs from other agents. From the autonomous actions of these agents, working in concert, human intelligence arises as an emergent property of the whole; though no individual agent could be said to be intelligent, the actions of the agents working collectively could be said to manifest intelligent thought.

It is ironic that Minsky called the collection of autonomous agents a society because society as it existed when he wrote The Society of Mind, and even as it exists today, does not operate as described by Minsky's theory. Despite platitudes to the contrary, contemporary society is not a society of individual action. It is a society of mass action: of nations and political parties, of labour unions, trade associations and professional groups, of mass media, broadcasting, assembly lines, religions, fan clubs, and professional sports. It is a society where taking direction is the norm, rather than the exception, and where activities in unison with others far outweigh the actions taken as autonomous agents. True, not all direction is autocratic, nor is all action compulsory, but nonetheless, without the mass, society as we know it would not exist today.

As a result of information and communications technology, this model of society is breaking down. New models of organization, new modes of action, and even new understandings of self are being increasingly defined by the individual as autonomous agent. By this I do not mean 'individualism' in the sense understood by, say, Ayn Rand. Autonomy does not entail competition with and disregard for the interests others. Rather, individualism, understood in the present sense, reflects the idea that individual decisions – to cooperate or to compete, to help or to hinder – are made autonomously, and not guided by a more or less autocratic direction from a third party. Autonomous individuals can and often will work cooperatively with others, but always at their own discretion, based on their needs, as they perceive them.

In her book The Future and its Enemies Virginia Postrel calls the defenders of centralized authority technocrats. Crossing political boundaries, technocrats favor stability and order (Postrel, 1998, p. 12). They are goal-directed. (p. 13) "By design, technocrats pick winners, establish standards, and impose a single set of values on the future.. There must be a single blueprint for everyone." (p.p. 17-

18) In contrast to the technocrats, Postrel describes dynamism as "fluidity, variety, competition, adaptation, learning, improvement, evolution and spontaneously emerging order." (p. 28) Similarly, "Dynamist policies do not try to rig results. They do not impose one best way." (p. 45)

Contemporary learning is designed along the technocratic model. It is designed to be stable and unchanging. It is characterized by rules, regulations, standards and policies. It is goal directed, and the same goals (high school, a bachelor's, a PhD) are desired for everyone. It is directed by a central authority – a ministry of education, a board of governors, a professor, a teacher – that provides direction, dictates outcomes, and establishes standards of performance. It is designed for the masses: a graduating class, a student body. And it, along with the rest of society is changing.

Society is changing because it can. Information and communications technologies, and especially the internet, have made mass individual communication and action possible. Where before people had access to perhaps a few thousand books, magazines and journals, the World Wide Web consists of at least a billion pages from millions of individual web sites (Sherman, 2001). This multiplicity of voices has resulted in a multiplicity of publishers; voters in the 2000 U.S. presidential election, for example, had thousands of points of view to consider, totaling 87 million pages, all within the reach of their desktop. (Alexa, 2001) Add to this the billions of emails, instant messages, and chat sessions exchanged each day. People today receive more information, from more sources, than ever before.

It is not possible for any individual to read all the information that is available. As a result, people are applying filters to the information they read: they collect, in bookmark files, sites they visit regularly. They create customized and personalized news feeds through services such as NewsIsFree. (Lasica, 2000) Media critics decry the lack of authority and credibility of these new news sources. (Lasica, 1999) And other pundits are concerned that individual filtering will narrow a person's view of the world, going so far as to suggest that they ought to be forced to view a centralized news source. (Shapiro 1999, p. 205) The distribution of news and information is becoming fragmented, and though waves of mass recognition sweep through the internet on a regular basis (the Princess Diana funeral, for example, or the "All your base are belong to us" craze), such incidents of mass awareness are emergent, rather than designed, phenomena.

Exactly the same trends are sweeping through education, though many practitioners are not aware of it (or deny that education will be affected in the same way). The number of educational opportunities available to an individual is increasing exponentially: not only may a person attend any of hundreds universities or colleges online (Office of Learning Technologies, 2001), they may take any of thousands of courses from any number of private and public providers in specialized disciplines, such as CyberU's Small Business Training Center.

Critics and pundits of education policy are raising the same concerns as media pundits. Writers such as David Noble suggest that new technology is resulting in the "commodification" of education, with a corresponding drop in quality (Noble, 1998). It is worth noting in passing that as much concern is raised regarding online credentials as regarding online learning. Nobody doubts, for example, that the information contained in C|Net's SmartPlanet online learning service is accurate. And it is worth asking "If face-to-face communications amongst students and teachers in classrooms are desirable for quality teaching and learning.why are

schools rapidly introducing the Internet into the classrooms and why is distance education via the Internet growing each year?" (Belam, et.al Despite the critics, online learning proliferates. And the readership at self-directed learning sites such as WebReference – more than 50,000 viewers daily – suggests that it is proliferating a lot more than is widely acknowledged.

Noble is absolutely correct in his main thesis, though. Education is becoming commodified. Attention, as he says, to "the production and inventorying of an assortment of fragmented 'course materials': syllabi, lectures, lessons, exams." (Noble 1999) The process is, as he says, "transforming courses into courseware, the activity of instruction itself into commercially viable proprietary products that can be owned and bought and sold in the market." (Noble 1998) And while Noble may rail against this trend, the same forces are at work in education as in information and communication generally: there is an ever increasing abundance of educational materials available, which is resulting in a narrowing of the field of opportunity for education providers.

What is pushing this abundance of educational material is a steadily increasing demand. While Noble may argue that "Education is a process that necessarily entails an interpersonal (not merely interactive) relationship between people." (Noble 1999) learners using the internet are disagreeing with their mouse clicks, determining for themselves that in at least some cases, the intermediation of the educator is not required. And while learners can – and do – request direction from experts in any of thousands of list servers on the Internet, many are learning everything from Tai Chi to Internet programming to flower arranging on their own.

Institutions are busily creating online learning opportunities, not because of some deep desire for control over faculty, as Noble suggests, but because they must: if universities and colleges don't develop online learning, private enterprises will tap into this enormous market ($7.1 billion dollars – Thompson 2001) by themselves. The distance learning market will increase from 5 percent of all higher education students to 15 percent by 2002. (Oblinger, et.al. 2001) One example of this is the corporate e-learning industry. Private sector companies earned $1.4 billion in revenues from the sector in 2000 and are projected to earn $4 billion by 2004 (Chen 2001) at minimum (Pastore 2001, Urdan 2001).

A large part of the reason for this is that learner expectations are changing. Just as news readers are filtering and customizing their news and information, online learners are learning to filter and customize their learning. This new mindset is neatly captured in a recent report distributed by EduCause (Oblinger, et.al. 2001, citing Frand 2001). In the "information age mindset," computers don't even count as technology, the internet is better than TV, and reality is no longer "real." The learner of the information age finds that doing is more important than knowing, expects to make mistakes and learn by trial-and-error, multitasks, and stays connected. There is zero tolerance for delays. And the traditional line between consumer and creator is blurring.

The new 'consumer' of online learning 'commodities' wants learning in a way never conceived by traditional colleges and universities. First, and most important, the new consumer wants it now. No applications for admissions or September start dates. The new learner wants it where they are because, as multitaskers, they are probably too busy to drop everything and take the bus downtown. And the learning must be customized to meet a specific need. No mandatory curriculum, no prerequisites: if there's something they don't understand, that simply creates a new education need. It's no wonder online learning is looking at portable,

wireless internet as its delivery mode of the future! (Stover 2000)

What we are seeing is learning as it evolves from a centrally controlled distribution of information from an expert to a consumer to an interactive, dynamic, user controlled set of information exchanges. And while most traditional efforts in online learning are directed toward putting 'classes' and 'courses' on the internet, as exemplified in Bates's standard model (Bates 2000), online learning of the future is looking a lot more like content management or knowledge management. (Morrissy, 1999) Even today, we are seeing a rise in interest for 'just-in-time' learning (and a corresponding rash of criticisms). But while the name 'just-in-time' suggests 'quick-and-dirty,' it is really an example of user-directed and demanded learning.

We reach, as an end point, what may be characterized as a "fragmentation" of learning, an environment where there are as many learning providers as there are web sites today, an environment where each learner picks and chooses from the array of opportunities. In such an environment, there is no centralized control of learning, no core curriculum, no universal set of standards and practices. Each person withdraws from the common pool that learning which is necessary for his or her interests and abilities. This suggests as well a multiplicity of delivery modes, from the 'how to rewire' instructions delivered via PDA to the top of a telephone pole to Tai Chi videos available on demand to informal seminars, conferences, and yes even, from time to time, old fashioned classes.

It is reasonable to ask – and many people will ask – whether as a result of this we emerge with a Minsky-style collective intelligence, or at the very least, a social level of education at least as high as exists today. But while such a question is probably of significant importance to the future of society, it is an empirical question. We are well down the path toward the fragmentation of learning and the question of asserting centralized control over our collective educational aspirations is well past moot. The age of chaos is upon us.

This article originally appeared in Education Canada, Volume 41, No. 3, Fall, 2001, pp. 4-7.

References

(dates in brackets indicate date viewed for this article (and by inference, the date the article was first written))

Alexa. Press Release. Library of Congress and Alexa Announce Election 2000 Collection. 22 June, 2001. (06 July, 2001)

Bates, Tony. 2000 Managing Technological Change. Jossey-Bass.

Belem, Robin, Lorber, Renata, Malcomson, Melanie and Peasgood, Sean. 1999. Education and the Internet. Chapters developed in STPP 4C03: "The Internet, Society, and Social Change", a fourth-year course Fall Term 1999 taught by Dr. Carl Cuneo, Technology, and Public Policy, McMaster University. (06 July, 2001)

Chen, E. Yegin. Corporate E-Learning ROI Scoreboard: Early Leaders Emerge. Eduventures. (07 July, 2001)

CyberU Small Business Training Center. Website. (06 July, 2001)

Lasica, J.D. 1999. The Media's Matt Drudge Syndrome. American Journalism Review, April, 1999. (06 July, 2001)

Lasica, J.D. 2000. The Daily Me: Personalization on the Web. Online Publication (to be a chapter in an unnamed McGraw Hill textbook). (06 July, 2001)

Minsky, Marvin. 1988. The Society of Mind. Simon & Schuster.

Morrissey, Charles. 1999. Enhancing Professional Education through Virtual Knowledge Networks. The Technology Source. (06 July, 2001)

NewsIsFree. Website. (06 July, 2001)

Noble, David. 1998. Digital Diploma Mills: The Automation of Higher Education. First Monday (reprint). (06 July, 2001)

Noble, David. 1999. Digital Diploma Mills: Rehearsal for the Revolution. (06 July, 2001)

Oblinger, Diana, Barone, Carole, and Hawkins, Brian. Distributed Education and its Challenges: An Overview. American Council on Education Centre for Policy Analysis. (06 July, 2000)

Office of Learning Technologies. 2001. Directories of Online and Distance Courses. (06 July, 2001)

Pastore, Michael. 2001. Companies, Universities Moving Toward E-Learning. Cyberatlas. (06 July, 2001)

Shapiro, Andrew L. The Control Revolution. Century Foundation, 1999.

Sherman, Chris. Google's Cool Billion. About.Com. 2001. (06 July, 2001)

Stover, Del. 2000 Hands-On learning. ElectronicSchool.Com. (06 July, 2000)

Thompson, Chris. 2001. The State of E-Learning in the States. National Governors Association. (06 July, 2001)

Urdan, Tracy. 2001. Corporate E-Learning: Exploring a New Frontier. W.R. Hambrecht. (06 July, 2001)

Webreference. Web site. (06 July, 2001)

This paper introduces a new thread to the discussion conducted thus far, the idea that students themselves are autonomous entities, and that this has implications on our educational system. I was fortunate to have had the chance to argue these points back and forth over several days with David Merrill, who himself has written on the topic. Basically, my point is this: it's not 'learner centered' if you tell the learners what to do, and this extends not only to learning style but to the selection of what to learn. In a more general sense: the system must be non-hierarchical and non-authoritarian all the way down.

9 In Practice...

Written January 28, 2002. Unpublished.

9.1 The Panel

SAN DIEGO – I am at the National Learning Infrastructure Initiative conference to take part in a panel discussion titled Learner-Centered by Practice: Applying What We Know About Learning and Cognition in Designing for the Online Environment.

With me were the panel moderator Helen Knibb, David Merrill from Utah State University, and Lynette Gillis, a consultant with Learning Designs Online. Like most quality panels, these members pushed some of my thinking into new directions. And I think as a result we need to restate clearly what we mean by learner centered learning.

So – what do we know about learning? Quite a lot, actually, and we've known it for a long time. Helen Knibb outlined some features:

Learning starts from what you already know

Learning provides usable knowledge

Learning involves learning to learn

Learning is community centered

Learning is addresses a "discipline base" of knowledge

All this is well and good, but as David Merrill argued, most teaching on campus, online or off, is "terrible." Of 60 online courses reviewed, he argued, only five had any educational value at all. At best, they do nothing more than provide information, but information is not instruction. The net result is what he called "pooled ignorance."

When we more to the concept of learner centered learning, it gets even harder. "Designing open learning is ten times as hard" as designing an online class, he said. Designers must go well beyond traditional "tell and ask" or "simon says" modes of learning.

But that said, good teaching does not change over time. The principles we have understood about learning apply in an online environment. "Technology is just another way to deliver stuff," he said. But the rules don't change: good teaching:

facilitates learning

apply to any system

apply to any architecture

are design oriented

Merrill's five principles could form the basis of the Commonplace Book of Learning:

Good learning is problem centered

It activates previous experience and knowledge

It relies more on demonstration than on telling

Learners should be required to use their new knowledge of skill to solve problems

And it should integrate new knowledge or skills into everyday life

None of the panelists had any disagreement with either Knibb's or Merrill's descriptions of good instruction. But on examination of this concept, especially in a web-based environment, we find ourselves drawn further and further from the traditional model of learning.

Lynette Gillis puts this intuition into concrete form. She described two projects undertaken by her group, one in which call center operators are taught new cell phone features, and another in which staff in a hospital are taught computer systems.

In neither case is the resulting learning structured along the lines of a traditional thirty-nine hour seminar and test. The cell phone application is essentially a virtual cell phone with two modes: show and try. The hospital learning was a combination of computer lab, expert coaching, self-study and documentation.

In preparing for the hospital training, Gillis's group undertook a large scale study looking at preferred training methods based on where participants were in their training. The focus groups suggested that people's preferences might change as they became more familiar with the material, and the study confirmed this suggestion.

When first introduced to a new application, participants preferred a short session in a computer lab. But for follow-up training and support, the vast majority of people opted for experts in the work area to coach them – no more classes. As mastery improved, they began to opt for self-study courses and documentation.

9.2 The Issue

As I learn about and experience more and more online learning, I found myself drifting toward a less and less popular position.

Well – that's not completely accurate: the position is one I sketched in my 1998 Future of Online Learning. But while intellectually I felt that it was a good prediction, I am acquiring more of a conviction that it was a good prediction (I won't try to explain the difference; just trust me on this).

The prediction is this: as online learning takes hold, fewer and fewer people will opt for traditional courses and classes, opting instead for less formal learner driven forms of learning.

In my talk I progressed through four phases of increasing evidence for this conviction. Experience – my own experience as an online learner is compelling. I never take classes, and yet have learned most of what I know today – from CGI programming to instructional design to Roman history – in informal, non-structured learning. On reflection, I find myself going through four stages:

Theory – I read the background (including even books) or theoretical basis for the discipline

Example – I look at examples of what I am trying to study, deconstructing the work, finding out what part does what

Practice – I write software, articles (like this one), or create web pages

Community – I distribute part of what I create, soliciting feedback, engaging in dialogue, participating in the discipline community

Observations: At dinner yesterday David Merrill suggested that my theory amounts to me wanting everyone to learn the way I do. I think there are worse ways to do it, but not, that's not it.

Having basically lived on the internet since the early 90s, I have had ample opportunity to observe how people who are strongly connected to the internet learn.

these people are not the sort of people who are studied by university instructors evaluating the effectiveness of online learning. For the most part, the university professors never even see this group, much less evaluate them.

For a significant number of people, if they want to learn about something new, they do not sign up for a university course or program, they turn to Google. They find out what information there is about the subject. The internet community expects this of each other (from this expectation comes the expression "RTFM" (Read the, um, manual) before any sort of 'instruction' occurs.

Their second source of information comes from the mailing lists, discussion boards, Usenet posts and other sources of exchange on the net. People new to a field are expected to "lurk" – that is, to listen to the discussion for a bit before jumping in with questions and comments.

Through these discussions learners are exposed to the practice of the discipline. How this works varies with the discipline – programmers and designers are exposed to actual code or designs, writers are exposed to writing and journalists to articles, historians are exposed to the back-and-forth interplay between theorists, and instructional designers to pedagogy.

If the learners need more detailed instruction – and many do – they can opt for it in a variety of ways. In almost every discipline, some sites are dedicated specifically to providing instruction. One example I use a lot in this content is WebReference, a site accessed more than 50,000 times daily by web programmers and designers. Project Cool is another. My own Guide to the Logical fallacies is yet another.

At last, the new learner is expected to put the learning into practice. A new web page design, some sample code, an email question or comment, an article – the essence of the web consists in 'putting it out there.' Seldom does an endeavour fail to elicit comment. And the exchange that occurs between new learner and seasoned professionals integrates the learner into the community.

Research: I could probably refer to Lynette Gillis's work and leave it there. But I am compelled to mention the Open University study that I mentioned last week.

According to the study, the number of young people studying at a distance is increasing rapidly. What's interesting isn't the number of people, it's the reasons they give. They find it cheaper, according to the release, they want to start working right away, and they find the university a bad place to study. "Some of our students have come to the OU having tried studying at other universities, where they have found the lifestyle, including the lack of a strong work ethic that some of them perceive, not for them."

What are we to conclude from this? Minimally, that when given the option, many young people choose not to study in a traditional university environment. That as people become more connected, they choose alternative forms of learning. And that

their choices are based in the fact that, on balance, the university isn't really a very good place to learn things, certainly not if you are not prepared to stop everything else in your life for four or more years. And for continued professional development – a rapidly growing area of learning as people become used to living in an era of constant change – a university course or program is simply unreasonable.

Practice: OLDaily is my attempt to instantiate my thinking in practice. Although it is called a newsletter, is designed to be a learning environment. My learners are professional course designers and people interested in online learning. My resources are the materials to which I link every day. My role as an 'instructor' (the word just doesn't fit) consists in my selection of materials and commentary surrounding that selection.

But OLDaily, the newsletter, is just a front. It is the most visible part of what is supposed to be the learning environment as a whole. For as materials are added to OLDaily, they are stored in a knowledge base. the knowledge base is intended to be used as a tool by people working in the field.

My theory of learner centered learning, in short form, amounts to this: learning ought to be created by the learner.

Now let me emphasize that I do not mean 'created' in some sort of constructivist way. What I mean is more like this: where in traditional learning (and traditional online learning) the selection and sequencing of the learning materials is a task performed by the instructor, in learning centered learning the selection and sequencing of the learning materials is a task performed by the student.

The system behind OLDaily, then, is a means of enabling this to happen. When completed (there is still quite a bit of work to do) it should be possible to work the knowledge base in a relatively intuitive way to create a sequence of learning activities directly related to the learner's area of interest.

In the first instance, only one sequence is possible: a chronological sequence. Type something into the search field (or make selections in the advanced search) and a list of resources will be displayed, most recent first (why search engines don't at least attempt to sort by date is beyond me).

The search terms are not matched to the content of the entries in the knowledge base; they are matched to the content of the commentaries. That's important – it allows me to weave threads of thought through the knowledge base by employing a consistent vocabulary. Now what's interesting (in my view) is that I, as the commentator, do not know what threads of thought are being woven into the data. This is something discoverable only by the student.

From today's newsletter, I wrote:

The Parsimony of the Explicit I think there's something important happening in this article but I'm not quite sure I can put a finger on it. Elearningpost summarized it as follows: "David Weinberger: Most Web designers try to control the users? experience. Some try to shape it. And a precious few try to become that experience." That's not bad, but it misses the fact that most of the article is a response against a view proposed by the W3C's Charles Munat to the effect that, "a web site is data, relationships among data, and transformations that may be applied to that data. These are all abstract. For us to interact with a web site, the data/relationships/transformations must become concrete. In an ideal world, the user would have complete control over how this process of um, reification, for want of a better word, occurs." Weinberger's response is, essentially, that if you

separate the content from the manner in which the content is presented – that is, if you separate the medium from the message – something important is lost. Now here's the important part: I think that both Weinberger and Munat is correct. The reader must create the relationships and the presentation, and yet, these must also be created by the designer. How how how? Solve this, and you've solved the fundamental problem of learner centered learning.

I think that something like OLDaily is that meeting point. The learning sequence that results is, in a certain way, a dialogue between the instructor and the learner, but one in which the learner doesn't know what the instructor has to say until he asks for it, and the instructor doesn't know what he said until it is asked for.

I've been working with code to take this a step further – the automatic book writer. Now perhaps one day I will write a book, but not soon. One of my major problems with writing a book is that I don't know where to begin and where to end. A book is a linear representation of a three dimensional topic, and thus is never better than a certain facet. And a book is almost always representative of a facet distinct from the learner's interests (which is why the whole concept of sequencing learning objects is odd).

If you would like to try it out, here is my book generator (if you create a book you like, feel free to publish it and give me the royalties – heh).

OLDaily will also include some further features intended to further express this concept. While today there exists only the [Refer] link after each item, I will be adding two more: a [Research] link, which will give the reader an opportunity to create a sequence of resources based on associations with the current resource, and a [Reflect] link, which will give the reader an opportunity to create a running commentary that will be shared with other members of the reading community.

What I am after here is a concept of learning where what is created is an environment, and where learning occurs through working within the environment. It is learner centered not merely in the sense that it is, as David Merrill would say, "open-ended", it is learner centered in the sense that the learning is created by the learner.

Will it work?

It is already working.

9.3 Epilogue

I saw this coming but I did it anyways.

In our discussion group here in San Diego – a group consisting mostly of university educators and administrators – I advanced the idea that learners, and not instructors, ought to design their own learning.

Now let me be clear – what I mean by that is learner created learning in the sense I have been describing above. It is not learner centered learning in the sense where we put them into a library and tell them to read. Experts are available. Instruction is available. But the organizing of the learning is undertaken by the learner.

Now in our discussions of learner centered learning the question of learning styles came up – inevitably. Interestingly, to none of us was learner centered learning a matter of designing for different learning styles. Not because we think that people don't have different learning styles (though Merrill hummed and hawed

on the point), but because designing for different learning styles doesn't capture the essence of learner centered learning. A good designer can design for different learning styles while at the same time being instructor or institution focused.

But Merrill did raise a division of learning styles advanced by one writer – I forget the name and can't find it on the net (when will people learn to publish their stuff on the net, where people can read it!). Essentially, the author proposed that learning styles can be defined according to the learner's attitude toward learning:

Transforming – they engage and transform the learning materials

Performing – they do what they need to get an A

Conforming – they do what they need to get by

Resisting – they do not want to be in this learning situation

There are different ways to approach this thesis. One way is to suggest that you need to design for each of these four groups. Another, volunteered by a member of the audience, is that the objective of instruction ought to be to convert every learner to the transforming mode of learning.

To me, however, the fact that there are people in a class who are not transformers – the fact that there are people in the class who only want a grade, who only want to get by, who don't want to be there at all – is evidence that the learning in question is not learner centered. Or to put the same point another way: learning centered learning means that each and every learner wants to be learning whatever they are learning.

In an important sense, learner centered design is a misnomer. Once we start making decisions for the learner – even if they are in the learner's best interests – we have moved from the realm of learner centered learning to the realm of instructor centered learning.

As I said, I saw this coming. Not exactly Valdy's chorus of boos, but a chorus of comments to the effect that learners are not able to structure their own learning.

I reject that proposition.

Look at it this way: according to David Merrill, the instructors can't teach. And according to Merrill and many others in the audience, the students can't learn on their own. The two sentences can't both be true, because it follows from them that no learning is happening at all, a proposition that is manifestly false.

Now in fact a little of both is happening.

It turns out that university professors, even if they have no schooling in the finer points of pedagogy, are able to convey some knowledge of their field to (at least) those people with a genuine desire to learn about it.

And it turns out that university students, even if faced with "terrible" teaching, are able to organize themselves sufficiently well to be able to learn (and as many commentators point out, a lot of this learning is explicitly informal – study groups, sessions in the pubs, practice in student clubs).

Indeed, it seems to me that most of the evidence and most of the argument against learner centered learning is based on bias in the questions and bias in the practice. Bias in the question, in the sense that self-learners never seem to be included in studies of the effectiveness of learning online. And bias in practice, in the sense that (in a university environment especially) only learning accomplished through formal instructor-centered learning is recognized as

legitimate.

I think it would be a useful exercise to develop some tools and some processes to support learner centered learning. To design, for example, a learning management system that sits on the student’s desktop and is operated by the student, not the instructor. To compare, in a neutral field, the intellectual achievements of self-managed versus institutionally managed learners. To see what learners can do, if we’d only let them try.

I have always admired the energy of Arun Tripathi and his constant dredging of relevant philosophical works has been of value to me. And his summary of a book by Hubert Dreyfus (who I remembered as a cybersceptic way back when I was studying connectionism at the University of Alberta) prompted this reply. It's a complex paper. Dreyfus is essentially arguing that there can be no virtual experience. My reply is a combination of McLuhan and Hume. As extensions of our perceptions, virtual realities are no less immediate than physical realities, I argue, because the perceiver is real. And such perceptions can (and do) form the basis of our experiences, and this is sufficient for the transfer of knowledge. I write, the key to success in teaching is in being able to connect abstract thought with concrete experience, to represent new knowledge and new information (and practices and skills) in a way that connects with the student's accumulated body of experience. I never published this paper because it quotes Tripathi exclusively, and not the original work of Dreyfus.

10 Education and Embodiment

Written April 26, 2002. Unpublished.

We hear over and over again the idea that our online experiences are impersonal, disembodied experiences, and that therefore a full education, which necessarily involves bodily experiences, is impossible online. It is impossible because we are in an important sense detached from the physical world, and therefore unable to distinguish between the significant and the trivial, because there is no causal connection between ourselves and cyberspace, no consequences to our actions.

This view of learning and living online not only misrepresents cyberspace, it also misrepresents how we experience and how we learn generally. For we are not disembodied entities when we interact online; we remain feeling and breathing beings. Our experience of cyberspace is not one of disembodied transportation; it is one of embodied sensation. We experience virtual phenomena through the same mechanisms that we experience physical phenomena, by integrating such phenomena into a personal ontology. And by interpreting both virtual and physical phenomena from similar points of view, we are able to experience virtual phenomena in the same way that we experience physical phenomena.

As a consequence, in a fundamental sense the causal impact of virtual phenomena is the same as the causal impact of physical phenomena: we cry real tears at the death of fictional characters. The objective of education – and the root of foundational aspects of education, such as cultural awareness and expert learning – is not based in rote imitations of the teacher or master. It is not based, and could not be based, in physical phenomena alone. It lies rather in our being able to relate new knowledge to our personal ontologies, to root significance and importance in the construction of our personal experience rather than the ministrations of a master. Thus, the experiences necessary for education are available from virtual, as well as physical, sources.

10.1 1. The Experience of Cyberspace

Let me begin with the bit of Hubert Dreyfus, as quoted by Arun Tripathi:

"When we enter cyberspace and leave behind our animal-shaped, emotional, intuitive, situated, vulnerable, embodied selves, and thereby gain a remarkable new freedom never before available to human beings, we might, at the same time, necessarily lose our ability to distinguish relevant from irrelevant information, lack a sense of the seriousness of success and failure necessary for learning, lose our sense of being causally embedded in the world and, along with it, our

sense of reality, and, finally, be tempted to avoid the risk of genuine commitment, and so lose our sense of what is significant or meaningful in our lives."

There is, as there is with most of Dreyfus's writing, a lot packed into a short paragraph. But importantly, what is here asserted is inconsistent with what one actually experiences on the Internet. Now one assumes that Dreyfus has worked on the internet, but one wonders whether he has really lived on the internet, lived with the internet, and come to see it as anything more than disembodied text. My own experience is that the Internet is a warm, rich, lush environment, so far removed from Dreyfus's barren characterization as to suggest that we write about two different realities.

At the core of Dreyfus's argument is the idea that we do not causally interact with the world when we interact with the world through the Internet. The Internet, therefore, in an important sense, keeps the consequences of our actions separate from ourselves, and conversely, keeps our selves causally insulated from whatever happens in the world. The idea here is that if we do not have a direct physical connection with whatever it is we interact, we cannot have a genuine experience of that interaction. Hence we are unable to distinguish the irrelevant from the relevant, the real from the unreal, the significant from the insignificant.

This is the sort of conclusion one would expect to hear from someone who is observing the interactions that occur online from afar, from one who sees and even uses the Internet but does not engage. It is the sort of reaction one would expect from a person who, while watching a person read a book would conclude that the experience of reading must be barren and artificial because there can be no direct connection between the author and the reader. Yet we know from our own experiences that the act of reading can transport us to rich and engaging worlds, and that the interaction between reader and author can be as intimate and as detailed as many experience in the world outside books.

It seems to me that Dreyfus overlooks entirely the possibility that the mind can engage, through intermediate sources, the reality that lies at the other end of the interface. Indeed, it should be remarked that this is a natural and normal function of the mind, and that this is something that we must do every instant of our lives. All experience is, to a degree, mediated, either though the waves of light and sound that interact with our senses, or even through the nerve impulses that carry the impact of a physical event in our toes to our brain. And what is normal and natural for the human mind is that it creates a story around these interactions: where there are gaps in our knowledge, it fills them; where there should be feelings, it supplies them. The tapestry that is our daily experience of the world is no less a product of our mind than it is of the world.

The entire scope of media would be impossible were this not the case. Observe the people at a movie theatre during a sad movie – or feel it for yourself – and watch as they are genuinely moved by the presentation of images and sound, images and sound that are, moreover, known to be fictional. Any person who cries at a movie would deny that there could be no causal impact from the movie to the self: there is indeed a causal impact, manifest by the tears or the laughter or the cheers. And yet there is no causal connection between the characters on the screen and the patron in the audience, no causal connection because there could not be: the characters themselves do not exist, being nothing more than roles played by actors; the actors, meanwhile, are sunning themselves at a resort in Tunisia, oblivious to any emotions their past performance may be at this moment evoking.

And just so with the Internet. It is a channel for information about the world, no more or no less than any other channel: like a book, like a movie, like our direct perception of the world around us. And when we receive information from the Internet, or when we interact with people on a discussion board or chat line, our senses engage in much the same manner. We reach out and we create the world with which we are interacting: that creation becomes what we understand as reality, and it is that reality that has a causal connection with our thoughts and emotions: a causal connection, because the reality and the reaction are in the same place, in our own minds.

We do not – as Dreyfus implies – ever leave our own bodies when we interact online. We do not "leave behind our animal-shaped, emotional, intuitive, situated, vulnerable, embodied selves." Quite the contrary; we remain all of these things: we remain firmly rooted in out chair, head and senses remain mere inches away from pancreas and liver, and the experience of our online engagement remains a physical one. What happens in cyberspace is not a transportation of the self, but an extension of the self: we do not ’go out’ into cyberspace, but rather, we ’look out’ or ’reach out’ into cyberspace. Dreyfus’s picture is like one who, on looking through a telescope, would imagine that he is actually on the moon: but of course that is absurd. We have merely extended our capacity to gather information from a remote location and transport it back to ourselves. Astral projection though cyberspace is an impossibility, and it makes no sense to draw conclusions about cyberspace based on the suggestion that astral projection though cyberspace is a representation of what really happens.

10.2 2. Meaning and Experience

Dreyfus argues that "if our body goes, so does relevance, skill, reality, and meaning." The theory is that our experiences of the world occur not merely in the mind, but are a function of an entire bodily awareness. Nowhere is this more evident than in the domain of skills, where in order to do something with expertise – such as to, say, throw a dart – we must train not merely our mind, but also our body: and indeed, even, the actions of the mind, if applied to the throwing of a dart, actually inhibit our capacity to make the throw. And so, similarly with what appear to be higher level cognitive functions: our ability to determine which information matters to us and which information does not. For our relation with information is, in the end, necessarily a physical relation, defining our physical interaction with the world, and a physical interaction with the world is necessary in order to understand what information would have an impact on that interaction, and what information would not.

But the question of whether a bodily interaction with information is necessary for learning is moot. This is, more precisely, the question of whether stimuli that have their origin in the world and which enter our system through the mediation of the bodily senses are necessary for learning. We can stipulate for the sake of argument that the creation of meaning, relevance, skill and reality are impossible without the bodily acquisition of information. To be more accurate, it is probably more precise to say that, without the body, our creation of meaning, relevance, skill and reality would be very different from what actually is the case. But let us agree that at least some bodily experience is necessary. We are left with two important questions:

1. Is bodily experience *sufficient* for the creation of meaning, relevance, skill and reality? In particular, is their creation the result of some sort of *deduction* from physical experience? Obviously not: for we know that our

conception of reality extends beyond that which we directly experience. Even our development of a skill transcends experience, for otherwise we could never throw a dart better than we have before. I have shot three 180s in a single evening; there was a time when I could not have done this, and had no experience of having done so: and yet I produced in myself the skill to do it, to go beyond my personal experience. It is clear that experience, while necessary, is not sufficient. Something else happens.

Philosophers have of course suggested numerous theories as to what it is that we do. Some suggest that we have innate linguistic structures that organize experience in certain meaning-bearing ways. Others suggest that we have some sort of innate knowledge, such as the knowledge of our own existence, around which all our experiences are centered. For my own part, I aver that we do not add linguistic structure or additional knowledge to experience, but that we manipulate information, through a process of selective filtering and recombination, in such a way as to achieve balance or equilibrium. What is important here is that our personal ontology – our view of the world and our place in it – is generated through both experience and our manipulation of that experience. It is as though, as I suggested above, we fill in the gaps of experience: as though we add what is needed to the sketch we are provided to create the causally complex set of objects and interactions we call the world. But our personal ontology is wholly contained within the mind: it is not 'out there' somewhere, and though we may feel and believe that what is in our head resembles (or represents, depending on your theory) what is out there, the centre, for us, of meaning, relevance, reality and even skill is in our mind, in our head, and not external to us.

2. If bodily experience necessary for *every* instance in which meaning, relevance, skill and reality are created? Or, put another way, in *some* cases, could we, on the receipt of information only, with no bodily intermediation, attach to that information meaning or relevance? Or learn a skill? Or even say whether what we have experienced is real or not? Now of course, in all cases, we are going to receive information through *some* bodily channel, direct perception being evidently beyond our capacities. But the question is more like this: can information, accessed through purely cognitive means – such as by, say, reading – be given meaning, be identified as relevant, be designated as real, or teach a skill?

All experience comes to us through the body, for it wouldn't be experience otherwise. But our experience of some things is direct: we are in the physical presence of the thing, and can see it, touch it, or if necessary, taste it. Our experience of other things is indirect: we are not in the physical presence of that thing, but learn about through some form of communication, through some writing, for example, or a video broadcast. What we see, touch or taste is not the physical object itself, but some representation or image of the object. Viewed in these terms, the question just posed is thus as follows: can we learn about something even though our experience of it is indirect? Can we assess what we have experienced, identify it was relevant, give it meaning or place it into context?

And of course, the obvious and intuitive answer is: of course it can. You are in the act of reading this sentence. As you read it, you are posing (and answering) questions to yourself, questions such as, "Is he right?" and "Why do I care?" If you have a background in philosophy then as you read the paragraphs above you may have been saying, "Oh he means Chomsky" or "This sounds like Descartes." As with the previous discussion, it is harder to imagine one learning a skill on the basis

of information alone, but there is no denying that the provision of information can help someone learn a skill: otherwise, when we want a person to learn how to, say, operate a radial arm saw, we would simply provide them with a saw and a lot of wood and tell them to get some experience. But we don't: we preface this practice with some sound, sensible, *cognitive* advice, from little things (like, "wear goggles") to big things (like, "don't put your finger in front of the blade").

The fact that our personal ontology is internal to the mind is what allows us to relate and to learn about things in an indirect manner, and indeed, teaching and learning would be impossible otherwise. Through experience and training we come to associate a visual perception of a cat and the words "a cat" in much the same way: while on one level knowing that a visual perception of a cat is not the same as someone uttering the words, we nonetheless draw the same inferences ("four legs, a tail and meows"), contemplate the same actions ("feed, pet, toy with a laser pointer") and feel the same emotion ("awwwwwww"). When someone describes the cuteness of their cat over the chat room window, my experience is analogous to a direct experience of that cat, analogous because I understand that the word "cat", a picture of a cat, and the direct perception of a cat, signify the same *thing* in my internal ontology, and thus, that what pertains to the real cat also pertains to the described cat.

Yes, I understand that the words on the screen merely represent the cat, and yet I understand that words are distinct from – and therefore have different properties than – cats, and yes, the cat in Romania is not rubbing on my leg. But being told, "The cat is rubbing my leg," signifies the same thing to me as a cat actually rubbing my leg, because I am able to understand that the words signify the physical event. And indeed if I am sufficiently engaged, if I am focused on nothing other than the words and what they signify, I can feel the cat rubbing on my leg, feel the pleasant sensations this evokes, feel warmth and fondness for the cat, even though cognitively I am aware that the cat is half way around the world.

Of course, if you have never owned a cat, much less seen a cat, it would be much more difficult for you to reproduce that experience from indirect experience alone. You would need to stretch the analogy a bit further, imagining it, perhaps, to be similar to your dog rubbing against your leg. Or perhaps, had you no experience with pets at all, you would imagine it to be similar to other exhibitions of fondness you have experienced in your life. Your direct experiences form the raw material from which you construct your personal ontology, and are essential in the beginning, but as you acquire a richer set of experiences, direct experience becomes less necessary in order for you to perceive an indirect, or virtual, experience in the same way you would perceive a direct, or physical, experience.

It may be argued that it is not possible to actually have a sensation of an event if the event is not occurring, but this again seems intuitively and obviously false. For if it were true, we would never 'hear' voices or music in our head (and yet, I can 'listen' to a tune endlessly in my head, so much so that the real problem is that I can't get it out of my head). Some people, such as myself, when they read or write, actually 'hear' the voice they are reading or writing. Visual perception works in the same way: it is possible for some people to visualize an object so clearly that the perception appears real. My guess is that it takes practice to attain such levels of visualization, but my own personal experience is that it is possible. When I dream, I have a sensation of being in a situation and

feeling interactions: this sensation appears real to me, and it actually takes some degree of reasoning to understand that it was just a dream. When a person is hypnotized (and if hypnotists are genuine) then people can have experiences based on the hypnotist's suggestion alone.

Our capacity to have experiences without being in the physical situation that produces the experience is manifest: this capacity is based on our ability to comprehend physically different phenomena as though they were the same phenomena, in turn caused by our capacity to represent any given phenomena to our internal ontology in a way of our choosing. In an important sense, whether the experience is direct or indirect, we first construct the experience from both our sensory input and by analogy with previous experiences. Both direct and indirect experience of the same entity are thus remarkable similar, and as a consequence, produce similar effects in the mind. This is the basis on which entire industries are founded: the publication of books, the showing of movies, the playing of music, teaching and learning, simulations and role-plays, and numerous more artifices beyond mention, so many indeed that any suggestion that bodily experience is necessary in every case appears absurd and misguided.

10.3 3. The Web and the World

Whole ranges of cultural phenomena are completely inexplicable if Dreyfus is right. The shock and distress felt throughout the United States when Kennedy was shot. The triumph and exaltation we felt watching Neil Armstrong step on the moon. The joy and pride felt by Canadians when Henderson scored the winning goal in 1972. The disbelief felt by so many when Elvis died. The outburst of sorrow when Princess Diana was killed. The horror felt around the world during the World Trade center disaster. The anger people feel as they watch the continuing violence in the Middle East. None of these are events that happened *to us* - we experienced them only virtually - and yet they have a deep and continuing impact. In October of last year I was in Sydney and a large aircraft flew low over downtown to avoid a thunderstorm. I looked up and *felt* *fear* - a tangible emotion caused not by the low flying aircraft but by my virtual experience of the terrorist attacks half a world away.

Dreyfus (and others similar in thought) depict in their minds a scenario in which nobody ever enters the world. Arun Tripathi quotes William Bennett talking about, "A school where students never enter a classroom. Where their math and science lessons are done in cyberspace, from home. Where their teachers sit in front of a computer instead of a chalkboard, and communicate with them by phone or e-mail. And where parents act as academic coaches, guiding their children through it all." And Dreyfus writes, "E. M. Forster envisioned and deplored an age in which people would be able to sit in their rooms all their lives, keeping in touch with the world electronically. Now we have almost arrived at this stage of our culture. We can keep up on the latest events in the universe, shop, do research, communicate with our family, friends and colleagues, meet new people, play games, and control remote robots all without leaving our rooms. When we are engaged in such activities, our bodies seem irrelevant and our minds seem to be present wherever our interest takes us. As we have seen, some enthusiasts rejoice that, thanks to progress in achieving such telepresence, we are on the way to sloughing off our situated bodies and becoming ubiquitous and, ultimately, immortal."

And yet - as I write this item I am eating three bagels and a pear alongside my coffee. This is necessary because I felt hungry. I have been typing for about an hour now and I feel a little bit tired. I pause, and look out my window to the

green forest. An air conditioner drones on the wall beside me, needlessly in today’s cool April weather, and annoying as vibrations echo through the room. I’m in my office, here not because of the dictates of connectivity, or even of employment, but because I like to say "Hi" to Sophie in the morning, to exchange jibes with Rod out on the patio, to feel the wind in my hair as I cycle down the hill, to select my environment. Solitary? Disengaged? I am never alone! If I really wanted to get out of touch with the world I would put down my keyboard and take a walk through the trees.

Dreyfus and others depict a world in which cyber interaction replaces all physical interaction and yet it is a world that exists nowhere but in their own mind. No proponent of online learning (save, perhaps, Bennett, assuming (which I doubt) that he is a proponent) proposes that students be locked in their rooms to interact online only. Why that would be as absurd as forcing them to travel to a special "learning room" completely isolated from the rest of the world! But even more significantly – even at the *very* *time* I am online – such as now – I am intimately connected to my body and to the world. It is not as though my physical experiences cease simply because I am typing and reading a computer screen. No, hardly. Nobody escapes their body, and it is disingenuous and dishonest to suggest otherwise. Every minute of every day, one is in contact with his or her body, and as a consequence, is aware on a minute by minute basis of its feelings, its needs and its foibles. Yes, perhaps there may one day be a way to transfer our minds completely into a computer: but until and unless that happens – or even comes close to happening – concerns such as Dreyfus alleges are moot.

Indeed, I would take this even a step further. Douglas Rushkoff, in Cyberia, pointed out that people who work on the internet find themselves traveling a lot more than they used to. This is certainly my own experience: in the last couple of years, even, I have traveled further and wider than I could have imagined in any former life. I have also met more people, attended more conferences, and touched tangibly more and more people. As I connect with people around the world there is a *pull* that draws me from my desk and into the world. *This* is the genuine experience of cyberspace: we are drawn *closer* to the people around us, not separated from them by a wired degree of separation.

The advent of wireless and mobile internet – a development no doubt feared by some – frees me even from this office, keeps me connected even when I am in the forest, allows me to visit friends in Australia without losing touch of my correspondents in Argentina, allows me to react, to *be* *in* the world in a way I could never be before, changing my one-dimensional and merely physical interaction to a rich multilayered set of interactions with others, adding to and enhancing, not replacing, my physical presence in the world. The suggestion that the internet would somehow replace physical experience can only be a suggestion made by a person with no degree of involvement with the online world:. It is as though they are suggesting that seeing in colour somehow diminishes our ability to see in black and white. But they have never seen in colour and cannot know that only by seeing in colour are we able to appreciate the unique and valuable nature of the monochrome art form.

10.4 4. Culture and Telepresence

Arun Tripathi quotes Dreyfus, "Like embodied commonsense understanding, cultural style is too embodied to be captured in a theory, and passed on in courses. It is simply passed on silently from body to body, yet it is what makes us human beings and provides the background against which all other learning is possible. It is

only by being an apprentice to ones parents and teachers that one gains what Aristotle calls practical wisdom -- the general ability to do the appropriate thing, at the appropriate time, in the appropriate way. To the extent that we were able to leave our bodies behind and live in cyberspace and chose to do so, nurturing children and passing on ones variation of ones cultural style to them."

Dreyfus's point here is that there is some knowledge that cannot be passed from person to person through teaching, that it can only be acquired through a process of direct interaction. He captures this point through invocation of what he calls "cultural style," an example that seems more appealing for its vagueness than its basis in fact.

No doubt Dreyfus is a cultured and cultivated man (I have never met him), popular at parties, the one with the nod and a smile at exactly the right moment, the one people want to be like, or to be seen with. My own experience differs. Culture is something I never really acquired: my experiences in high school and university form a mélange of one awkward social encounter after the other. I was a nerd and a geek in the classic sense of nerdiness and geekiness, more likely to be the one in the corner of the room weighing arguments with a pint or porter than the one exercising witty repartees with the cultured elite. I am still that way to a some extent, and yet this is due to no lack of social engagements: I have had many examples to emulate (and have even tried from time to time, but it appears that culture, like a suit (in which I am also uncomfortable) is much more a personal matter than a socially shared set of conventions).

I am by no means alone, though perhaps you need to escape from the centre of the party to see this: in my storied history of bowling allies and biker bars, malls, neighborhoods and back allies, classrooms and clubrooms, Legion halls and living rooms I have seen more than my share of cultural inappropriateness. Indeed, if anything seems to be the rule, it seems that people do *not* learn culture except via explicit instruction: that the mannered are well schooled in their manners, that breeding – as they say – shows. I am all too aware that what passes for being "passed on silently from body to body" to Dreyfus is experienced by many as "passed on from bully to victim" in the schoolyard: that our cultural awareness, if indeed any such thing exists, is nothing more than a complex set of coping mechanisms designed to enable a smooth (or at the very least pain-free) passage through the rites of childhood and adolescence. Even in my thirties and forties I find myself being given explicit cultural lessons ("You don't announce how many games you've won," I was told, after summarizing the results of a particularly good set. "Since when?" I replied.)

Children are explicitly told to mind their manners, to sit up straight, to chew with their mouths closed, and more. Everybody knows the day they were instructed in the use of the salad fork (and how it differed from the dessert fork), for no amount of observation and emulation seemed to riddle this mystical morass. In my own life, the moments of cultural awareness came as the result of explicit – and pointedly non bodily – instruction: I learned how to communicate in a corporate environment by taking a video course called "On the Way Up" while I was with Texas Instruments in Austin; I learned how to speak well publicly by reading a book (I believe) by Keith Spicer called "Winging It," and I learned how to be the popular (and humble) man I am today by studying Dale Carnegie's "How to Win Friends and Influence People."

On looking back on my education, I think that I would have been much better prepared for the world around me had I experienced less of Dreyfus's body to body

interactions and more explicit instruction. It would have been very helpful to me to be able to enter simulations or practice environments where I could 'try on' social personas for cultural fit, to find a style I found comfortable and which did not offend the neighbors. I do not know what it is about cultural knowledge that Dreyfus feels can be transmitted only through body to body knowledge, but I am quite sure that as we press for the details we will find that this knowledge is very rarely passed except via explicit instruction, whether it occurs in the schoolyard, the classroom, or in very many cases (to judge by the self-help section in the bookstore), via books and audio tape.

10.5 5. Embodiment and Education

Nobody believes that (as quoted by Arun Tripathi) "the development of the Internet will solve all the problems within education." Nobody. To put it in philosophers' language: the Internet is a necessary, but not sufficient, means for providing a quality education for all.

The quote continues, "If the development goes in the right direction, they maintain, first class education will be available for everyone – in so far as they master the information technology. Thus the problems posed by too many students and too few universities as well as the serious problem of access to the good but expensive universities will be solved."

Again, more misinformation. Educators and designers work toward educational systems where it is not necessary to "master the technology". And online learning isn't about providing everyone with "access to good but expensive universities" – it is not about that at all. To many people – myself included – it is about replacing the *need* for "good but expensive" universities, about making education that is cheap, easy to use, accessible and applicable, making it available to everyone regardless of their income. The idea of online learning isn't to solve the problems that beset an array of nineteenth century institutions but to come up with a twenty-first century replacement for that array.

It will take much more than techno-dollars. Much more. We have to rethink in some fundamental ways just how we go about teaching people, how we *can* go about teaching people now that we have some real tools (as opposed to cave-people chalk-and-slate media) for teaching.

Here is the Dreyfus take on distance learning, as described by Arun Tripathi: "the imitation of the example of the teacher is a crucially important element in education at all levels. In many areas, the student can only learn to be an expert by imitating the day by day responses to specific situations of someone who is already an expert, or ideally, a master, and only by working closely with students in a shared situation can teachers pass on their passion and skill to their students. As the shared situation included community practices as part of what is learned and sometimes it will not, but in any case the actual presence of the coach or master is essential. So, in general, in so far as we want to teach skill in particular domains and practical wisdom in life, which we certainly do, we finally run up against the limits of the World Wide Web. As far as we can see, learning by apprenticeship can work only in the shared situations of the production sites of the crafts, or in the nearness of the classroom and laboratory; never in cyberspace. Thus the use of the Internet represents an impoverishment, not an improvement, of education. It can facilitate a kind of mass education, but it will only teach the students the rules and facts that can make them competent."

I think that people *do* learn by imitating, but I'm not sure it's always appropriate, and I'm not nearly convinced that learning by imitation in a classroom environment is the way to go.

I believe that people learn by imitating. I can recall numerous instances when I have done it. Most of my recollections, though, are of hilariously bad imitations. The day when our Scout Troop was camping by the shores of the St. Lawrence river and we found a particularly good wood for burning (called "punk") that the assistant Scoutmaster called "den-o-mite". Well, he probably said "dynamite," but I heard what I heard. Of course, when I repeated the word a few hours later, it was deemed ridiculously inappropriate by my fellow scouts. Or when I heard my boss Bob Avila say of Carly Simon, "She has the perfect life - a hit song and married to James Taylor" - I learned the hard way that only James Taylor fans would find being married to him anywhere near perfect. That's the thing with learning by imitation: it's hit and miss. When you learn by imitation, you fail to make those subtle distinctions between the imitations that will be genuinely useful in later life and the imitations that demonstrate clearly that you have *not* mastered the intricacies of cultural awareness.

In the classroom environment though my teachers were without exception well intended, they did not always make the best models to imitate. Over time I learned as every schoolboy learns that teachers are exactly *not* the people to imitate, for they embody everything (well, almost everything) that a person does *not* want to be. A know-it-all. A ruthless despot and enforcer of order. Sometimes inclined to allow personal preference to outweigh justice and fairness. Completely culturally unaware (one teacher of mine had the affront to confiscate my collection of hockey cards - the most serious and deeply disturbing action anyone could undertake on the playground, and action that would be, except for the imbalance of power evident in every classroom, a mortal sin).

In any case, if imitation is the source of learning, then teachers are (and always will be) vastly outnumbered. Even in special cases where the master takes the apprentice under his wing, there are numerous outside influences. In today's mass media environment, people are much more likely to imitate Bart Simpson than their math teacher (they see more of Bart Simpson, and in any case, he's popular). People hear the cultural wisdom of 'NSync dozens of times a day; no teacher could hope to match that. One's older brother is a much more pervasive influence than one's once-a-week geography teacher. And in the kingdom of the schoolyard and the locker room, as everyone knows, the teacher holds no sway whatsoever. What we find in contemporary education is that the influence (and imitation of) the teacher runs counter to many of the prevailing trends, so much so that there are concerted efforts to provide children with positive role models on the television screen, in cinema, in music and sport, and yes, even on the internet (through such sites as PBS's recent "It's My Life").

It is fortunate indeed that educators can depend on a much wider array of teaching tools than rote imitation. It is fortunate indeed that the primary role of the educator is to teach the student to move beyond mere imitation and into higher levels of cognitive awareness. One of the major arts the teacher practices is to enable the student to reason - and learn - at what might be called higher cognitive levels, to not depend on what embodiment for instruction but to be able to reason and integrate learning more abstractly. Yes, to a large degree, what we learn will be necessarily concrete, but from the first day of school the teacher - through the process of teaching language, art, mathematics and music - is teaching

the child how to represent perceptual experiences in non-perceptual form, so that (for example) the student can learn about Spain without actually having to be taken to Spain.

The key to success in teaching is in being able to connect abstract thought with concrete experience, to represent new knowledge and new information (and practices and skills) in a way that connects with the student's accumulated body of experience. With a mathematical formula to evoke *this* reaction, with a turn of a phrase to evoke *that* sensation. Connecting language with experience is probably the most difficult form of learning possible, but it is also the most effective, for language is pliant in a way that experience could never be, and if I could, with words, *describe* the moons of Jupiter in such a way that you actually, physically, shiver, then this transcendence has been achieved and mere imitation is no longer necessary.

For what the teacher wants to do over the course of an education is to facilitate in the child what I have been calling a personal ontology, a world view full of causally connected experiences, populated with entities (necessarily constructs of the mind), where rules and principles apply, some by nature, some by society, where the consequences of actions and interactions in the mind resemble in important ways the consequences of actions and interactions in the world: no instructor can ever transplant his or her personal ontology into the mind of a child, for no child has the instructor's experiences and sensations, but through the generation of experiences and interactions, and through the provision of cognitive tools useful in creating a framework and interface layer, the instructor can foster the child's development of his or own personal ontology and the tools needed to use it as a means of understanding and interpreting all manners of information, real or virtual, in the future.

10.6 6. The Promise of Indirect Experience

So in one, trivial, sense Dreyfus is correct. Of course we need direct experiences in order to learn. Direct experiences are the raw material from which we construct our personal theory about what's in the world and what it does, our personal ontology. The entities we touch, see and taste form the prototypes against which we evaluate, understand and assess subsequent experience. Were we to have no direct experiences, we would have no basis, no foundation, on which to learn anything at all.

Fortunately for the teacher's art, none of us is without direct experiences even for an instant. We are at all times connected to our body, at all times amassing and assessing a constant flow of sensory input. Even when we are watching television or surfing the Internet, the body's productions continue endlessly. The data we collect from the video terminal forms only one part and arguably even a small part of the experience of the moment.

True, what we experience on the Internet is only an indirect experience, but the words and images displayed on that tiny screen can produce a powerful impact. They can take us to distant nations, to introduce us to people from different cultures, expose us to difficult ideas. The power of the Internet is directly derived from the fact that it is indirect: a galaxy of worlds and ideas beyond our experience, beyond what we could experience, is presented to us. And yet, because we have in our mind a cognitive basis for the assessment of indirect experience, a personal ontology against which we can weight this wealth of information, we can experience these worlds and ideas for ourselves.

Of course the world on the computer screen is a virtual world. But the experience of that world is real, and in the end, that’s all that matters.

In a doctrine that has been popularized this year in the paper World of Ends, good network design involves making the network stupid and the applications smart. The idea is that the network should not establish anything more than a minimal constraint (and as I say elsewhere, only a syntactical constraint) over its contents. The reason for this is that the purpose of the network is to enable unfettered (and unaltered) communication from one node to the author, and each constraint acts against this. It is with this in mind that I reacted to the idea of smart luggage that never gets lost, and so I floated the idea that the smartness of the network could be obtained not only by making the ends intelligent, it could be obtained by making the objects intelligent.

11 Smart Learning Objects

Written May 4, 2002. Published in Learning Place, May, 2002.

A recent article [1] about airline luggage prompted me to think about learning objects in a new way. The premise of the article was that that airlines would lose luggage less frequently if the luggage were equipped with more intelligence. In other words, 'what if we managed bags like we managed people?' After all, 'passengers are smart entities traversing a stupid network, whereas pieces of luggage are very stupid entities traversing a marginally smarter network.' Wouldn't it be better if 'a suitcase could check itself into airplanes, order transportation, track news about delays or cancellations, and make sure, in case of unforeseen changes, that it will be booked on the next flight or sent back home again?'

My first thought was that we could revolutionize education by treating students more like people and less like luggage. Students, for example, could pick their own learning 'destination' (though we may want to provide them with learning 'travel agents' to help them). They could choose their own time and mode of travel, paying for first class if they need the extra assistance and a pillow under their seat or economy if they just needed to get to there from here in a hurry. We could depend on students to find their own way to the learning 'gates' that would lead them to their destination. And while it may be a little more expensive to provide students with the infrastructure they need to make their own choices, the results would be better: they would be far more likely to arrive at their destination of choice instead of, say, Latvia.

But as much as I like my first thought – and I do like my first thought – I like my second thought better. Over the last five years or so we have all been struggling with the need to deliver educational content in chunks through a distributed learning network, an effort very similar to the airlines' handling of luggage. Our conception of a chunk of learning content – or a learning object, as the current jargon has it – is that it is about as intelligent as a piece of luggage. And like a piece of luggage, it sits there in a repository, waiting to be found, waiting to be directed, assembled, placed in its seat in the airliner of learning, waiting to be delivered to eagerly waiting students clumped together around the carousel on the concourse.

Learning objects are like luggage. I don't know if I've seen that exact analogy employed in any of the many papers on learning objects, but the descriptions I have seen suggest the metaphor. 'Digital images or photos, live data feeds (like stock tickers), live or prerecorded video or audio snippets, small bits of text, animations, and smaller web-delivered applications, like a Java calculator' [2] – doesn't this sound like 'socks and shorts, a toothbrush, camera and electric razor?' And doesn't 'webpages that combine text, images and other media or

applications to deliver complete experiences, such as a complete instructional event' sound like 'suitcase, carry-on bag and backpack?'

Gerard's 1969 description of how 'curricular units can be made smaller and combined, like standardized Meccano [mechanical building set] parts, into a great variety of particular programs custom-made for each learner,' now cited with approval in contemporary learning object literature [3] points to the fundamentally stupid nature of learning objects as currently conceived, objects so stupid that they could not possibly function without a professional instructional designer and a six-figure LCMS in order to be of any use. 'They must be seen in terms of their place in an architectural hierarchy capable of finding, comparing, and selecting them and then joining them together to perform an orchestrated instructional function that requires more than a single object can accomplish unless it is a self-contained instructional product.'

It seems like a lot of work. In order to use a learning object you must first build an educational environment in which they can function. You need, somehow, to locate these objects, and then arrange them in their proper order, according to their design and function. In certain cases – as for example when the object is a Flash animation or a chunk of streaming media – you must arrange for the installation and configuration of appropriate viewing software. And you're still not finished: the objects must now be delivered in some sort of instructionally appropriate context – a problem solving environment, say, with expert models. Sure, it's a lot easier to do all this with learning objects (indeed, we wouldn't dream of doing it without them), but can't we manage learning without having to build the online equivalent of an international airport? Sure we can. We just need smarter learning objects.

So what would it take? We need to stop thinking of learning objects as chunks of instructional content and to start thinking of them as small, self-reliant computer programs. This means more than giving a learning object some sort of functionality, more than writing Java calculators or interactive animations. When we think of a learning object we need to think of it as a small computer program that is aware of and can interact with its environment. This is the purpose of what the authors of SCORM call a 'wrapper,' a set of 'functions that encapsulate the functionality that an AU might use to communicate with the LMS (Learning Management System).' [4] The idea here is that a learning object is more than just data being pushed around by an LMS, but rather, a piece of computer code that plugs into and actually works with the LMS – or whatever environment it may find itself in.

So how do we move from the concept of a wrapper – which is still pretty primitive – to the concept of a smart learning object? 'Know thyself,' according to Socrates, is the first step on the road to intelligence. As a learning object is being created, it should be created in such a way that the wrapper can learn for itself who the author is, what company or institute the author works for, what day it is, what its name or title is, what format it is, and more. No author should have to type endless fields of metadata; a wrapper, when first created, should initialize by detecting its environment and the nature of its contents. Just as Microsoft Word automatically embeds the authors name, institution, date and other information into each document it creates, so also the wrapper should obtain similar data and create some metadata tags on its own.

When launched, the learning object advances to the philosophy of Rene Descartes, asking, 'what kind of being am I?' A quick scan of the first few lines of content

should be sufficient to tell the object that it is a PDF file, or a GIF image, or a Flash animation. It should also be able to find its own size (or dimensions) as the case may be and write a few more lines of metadata. In an ideal world, the learning object would then search through its environment for an appropriate player. If it is a PDF reader, for example, it locates an Acrobat Reader and bundles it with itself. In the same way that a piece of software can be stored as a self-extracting zip archive that automatically installs itself when activated, so also can a learning object become a self-contained executable file that contains everything it needs to run properly.

And so, thus equipped, the learning object is cast into the air (or ether, as it were). From a short list provided, it seeks out the nearest learning object repository. Each learning object repository indexes data its own way, of course (people long having since abandoned the foolish notion of there being one and only one set of metadata standards for the entire world). It identifies the fields required by the repository and then consults the recommended application profile to learn what possible values could fill those fields. The learning object isn't so smart that it can know what those fields actually mean – after all, what is the Dewey Decimal System to a small chunk of code? – but it knows that if it submits to a scan by the library association's auto-summarizer it can obtain legal values for the repository's DD field. It does this for each field in turn – check the field, check the parameters, then access a web service to fill in the appropriate for itself. After a long, laborious process (almost a whole minute, an eternity for a learning object), it generates a new metadata file, checks a copy of itself into the repository, and then moves on to the next repository in the system where the process is repeated.

Once the learning object has checked itself into a few major repositories it goes to sleep. The metadata it generated sits in the repository, ready to generate a response to a request. Since the learning object repositories are networked, a request from anywhere in the system will eventually reach the learning object. Because the learning object generated detailed metadata, it will be exactly what the searcher was looking for. Nudged into wakefulness by the repository, the learning object sends a message to the searcher volunteering its services. It also checks the repository to see if any reviews have been written or whether it has achieved certification, and if so, offers that information to the searcher as well. It makes some enquiries of the user's system. Does the user have any credit with the appropriate commercial broker (if not, then propose a commercial transaction)? Does the user have the appropriate prior learning (if not, generate another request and then go back to sleep).

Serendipity! The learning object is exactly what the user wanted, and the user is qualified to download and run the learning object. Without waiting for any further instructions, it fires a copy of itself to the user's computer. A quick scan: what language is the user using? Spanish? Quick, access an auto-translation service and reconfigure. Next, obtain the user's preferred font styles and sizes, colours and other environment parameters. Then look at the parameters of the request: send your metadata, says the request, to the following program address. Here's what I am! The learning object proudly announces to the system, and when you need me to start up, just send the following command!

Obviously in this short article I have glossed over many of the details. But the main point of this article is to show that, if we made learning objects a little smarter, they could perform many of the tasks we now envision the hiring of

minions of baggage handlers to accomplish. There is no in principle reason why we could not develop smart learning objects: if we can write self-extracting executables, self-pacing audio streams, or applications that report back to Microsoft when you use bad words, then we can write learning object wrappers that perform basic self-analytical tasks, scan the web for web services, and learn about their new environment.

What's more, once we develop an architecture of smart learning objects, we are no longer constrained by the bounds of 'supported' data formats. Should a developer want to deploy a previously unused 3-D multimedia file format, the developer need not wait until the learning management system has built in support or until the user has downloaded a plug-in: everything that's needed ships with the learning object. Should developers become dissatisfied with Adobe's or Microsoft's file formats (or pricing structures), they can simply write their own presentation application.

Smart learning objects. So much more than learning luggage.

[1] Attendre le suitcase

http://www.acm.org/ubiquity/views/e_andersen_2.html"

By Espen Andersen, Ubiquity, April, 2002.

[2] Connecting learning objects to instructional design theory: A definition, a metaphor, and a taxonomy

http://www.reusability.org/read/chapters/wiley.doc

David A. Wiley, in The Instructional Use of Learning Objects, David A. Wiley, ed.

[3] The Nature and Origin of Instructional Objects

http://www.reusability.org/read/chapters/gibbons.doc

Andrew S. Gibbons, Jon Nelson and Robert Richards, in The Instructional Use of Learning Objects, David A. Wiley, ed.

[4] Sample learning Management System

http://www.flextraining.com/

Version 2.0. Advanced Distributed Learning.

In the philosophy of mind there is an old saw known as the homunculus theory, which posits that a little man sitting inside our brain handles all our thoughts and beliefs. The problem with the theory, aside from being physically impossible, is that it does not solve the problem of cognition, it just pushes it back a level. In the larger world, though, not only is it physically possible to have little men at the controls, they actually do seem to solve the problems. That's why we have ministers and mentors, teachers and truckers. But the irony is, in a world where we obtain a sufficient degree of connectivity, these homunculi are no longer needed. What's more, we find that they don't actually solve the problem of rationality in a network, they simply push the problem to a specific human brain. Because of this, they are actually counterproductive, substituting authority for rationality. Thus it is with the role of the teacher in the information age, and thus it is that I cast them in a new role, described below.

12 Aggregators, Assimilators, Analysts and Advisors

Written June 14, 2002. Unpublished.

In his newsletter today, Elliott Masie wrote, "I never balked at paying for a print newspaper, yet this was a bridge to cross before I could feel good about paying for an e-newspaper. Ironically, the paper tells me they have gotten way fewer subscribers than they thought! So, they are playing with their pricing and business models. I am intrigued about readers' thoughts about your own response to paying for e-content, whether it be a daily newspaper or knowledge flows." What follows is my response to him.

The whole question of subscription fees for online newspapers is a useful area of enquiry. The analogy with e-learning is significant, and many of the lessons providers of online learning content are about to learn have already been learned in the online media community.

That said: there are many ways we could approach this whole question, but I will take the simplest and most direct. Why are you paying $18 per month for something online when you could get it mostly for free? Because there are so many news sources out there, most of what is in your newspaper is freely available. Vin Crosbie has observed on numerous occasions, and I agree, that local content is the only real value a local newspaper has to offer online.

It would make sense to pay a subscription fee only if you could not obtain your local news online. With three city newspapers, you have plenty of choice. But perhaps one day all three will start charging subscriptions. You now face what I sometimes call an artificial scarcity: there is no shortage of supply, however, withholding (often via a cartel, an approach Steve Outing frequently recommends), creates one.

That all aside, is $18 a fair price for what you are receiving? First, let's not forget that you have had to purchase the reading device (your computer) and access to the distribution system. $18 a month is comparable to the cost of receiving each issue on your front doorstep. And yet the newspaper is able to deliver this item to you without paying for the newsprint and delivery costs. You should ask why it costs the same for you to purchase a product that costs the newspaper a lot less to produce. Particularly when only a small percentage of the content is useful to you.

This same issue arises in other domains. The RIAA is asking for a rate of a dollar a song for web music downloads, which amounts to about $12 a CD, which that same organization reports is the average price of a CD in the United States. Why should

you pay the same price when the producer does not incur the cost of manufacturing the CD and case? By and large, consumers have rejected the proposed pricing with respect to music, and if the evidence that parades through the sources above (Steve Outing's ministrations notwithstanding) is any indication, consumers have essentially rejected that pricing model for online media.

Juliette Adams, in a recent article, articulated the issue nicely. She wrote, "The American market has spoken: an article is worth $1–5, a 10 hour–course $50–100, a full–length course $500–2000. These prices have been set by the aforementioned publishers, journals and academic institutions, my favorite clients withstanding. But if these prices are 'right', then why aren't there more 'eLearners'?"

In fact, you do not have to pay your local newspaper in order to access local news. Leaving aside the Google news search – for my area, New Brunswick – there is a wealth of alternative sources available online. For people living in larger cities, especially, the gamuit of press releases, announcements, web logs, activist pages, and similar sources gives you as detailed a picture as any newspaper. What you are missing, true, is someone to filter all this information and present it to you in a nice format. But the more people are looking at these alternative sources the less people are satisfied with the filtering and writing offered by professional journalists.

A discussion that has of late occupied a lot of ink on the Online News discussion list is the whole area of alternative news sources. Eric Meyer chimes in regularly with the assertion that they are parasites. But there has been a fair number of words devoted to the idea that web logs (blogs) are not only link lists but actually first hand accounts of news events, and therefore, in an important sense, more authentic. Paul Andrews documents this at length. "Aided by the Internet and personal–computer software, online communities with their own publishing tools and networks are redefining news in the 21st Century."

Educators play the same sort of role in society as journalists. They are aggregators, assimilators, analysts and advisors. They are middle links in an ecosystem, or as John Hiler puts it, parasites on information produced by others. And they are being impacted by alternative formas of learning in much the same way, for much the same reasons. "By adding to the diversity of original content, weblogs have added a whole new layer to the Media Food chain. That puts weblogs at the base of the food chain, generating the sort of grassroots journalism that the new Media Ecosystem has grown increasingly dependent upon. Because bloggers are closer to a story, they'll often pick up the sort of things that traditional Journalists miss."

One of the things that attracts people to weblogs is the diversity of their content. Everybody knows what to expect from traditional journalists. There's a fairly predictable story style, a fairly predictable political tone, a fairly predictable range of coverage. Weblogs draw from a much wider range of content, style and opinion. Loyalty is bought by bringing a point of view, a perspective, to the aggregation. By not merely listing news articles but describing how they fit into an overall pattern, by expressing an opinion based on a certain set of assumptions. This is not cheap and easy; it requires expertise and commitment.

It's one thing to talk about our being used to paying for printed newspapers, and to talk about it only being fair to compensate editors and writers for plying their craft, but in fact with virtually free global syndication, the substantially reduced cost of publication, and an increasing capacity on the part of the public to speak for itself, such productions and such professionals are not needed in

nearly the quantity they were formerly. When we look at what is possible with new media and internet technology, it makes less and less sense to be paying print era prices for online reproductions of industrial age products.

Indeed, not only are their new modes of news production and distribution, it is arguable that in such a new environment traditional barriers like subscription fees do not create wealth, they hinder it. Most alternative sources of news are free, despite the fact that they represent hours of time and effort on the part of their authors. My own work is a case in point. Would there be a net gain in wealth were I to charge subscriptions for this material or even to run advertising in this material?

My answer to this is no: subscription fees would mean that many thousands of people who could benefit from the material on this site would not benefit. Even the use of advertising makes the distribution of the site in many contexts (such as schools and religious institutions) problematic (assuming the advertisers are willing to support material that undercuts their methodology). Could I make a little money? Yes, maybe even a lot of money. But only at the cost of removing a valuable resource from society. And only, moreover, at the cost of cutting off my own access to similar resources.

What goes around, they say, comes around. I have in turn made use of the extensive body of free information available on the world wide web. By making use of this information, and by sharing freely on my website, I have been able to gain expertise in some fields (one of which, I like to think, includes online publishing). This expertise has helped me to obtain a progression of positions in the online learning and online resource sector. It is true, I do not actually get paid for any of my online work. And true, I must provide actual services in order to earn my salary. But the salary I earn, the positions I hold, would not have been possible to attain without the free sharing of information, both on my part as a gift to others, and by others as a gift to me.

It is astonishing to me that there are some writers, indeed, some entire communities of individuals, who are unable to imagine any form of compensation other than direct payment for services rendered, and who are indeed not even to imagine the possibility that one may gain more from freely sharing information than by hoarding it for oneself. But this is my experience and the experience of not a few, but hundreds, even thousands, of individuals living and working on the web and in related industries. Look around and find any 'guru' or 'expert' or even 'consultant' in any field and you will find a wealth of information distributed for free. These people are by and large earning more money than writers, based largely on their writing, and yet are not paid as writers.

In the old economy, scarcity was wealth. But we do not live in the old economy any more.

Indeed, the very idea of placing a newspaper on the web presupposes access to a wide range of free services. Imagine what it would be like in a subscription-happy world were newspapers to pay the full cost of putting out a web edition. They would have to pay Tim Berners-Lee a whack of money for spending 8+ years of his like providing a free system (HTTP+HTML), pay the U.S. military and dozens of nameless programmers who developed the underlying TCP/IP code, pay large royalties to people (like me, say) who through trial and error demonstrated the feasibility of using the web as a publishing medium at all, pay royalties for each email (POP+SMTP) they send, and more.

While people talk about paying fair compensation for value received, perhaps we ought to examine the conduct of newspapers themselves. Imagine what the news would be like if you had to pay royalties to every accident victim you covered for the use of name, story and photos, pay and obtain permission from politicians and other flacks who issue press releases, pay and obtain permission to run stories about strikes, lockouts, and other labour conflicts, and so on. Newspapers depend on an environment of free information. Indeed, their essential function is to pick up information that other people have produced, sometimes at great cost to themselves, repackage it, and distribute it in a bundle along with some advertising. "Crash kills 5. Eat at McDonalds."

This is why it is dangerous for newspapers to take the subscription route. The companies most likely to be damaged by putting restrictions on the free flow of information are those companies that earn their livings from the free flow of information. If information becomes a commodity, as some of you are suggesting, then why should I, as a newsworthy (and humble) person, allow you free access to any of it?

The significant issue here, one that is obscured by the day-to-day question of whether we should pay subscription fees for newspapers, music or online learning, is the manner in which the online content industry is being warped in order to protect these industrial mode forms of commerce. The most frequently voice argument you hear, no matter what the domain, is that the author|musician|artist should be fairly compensated for their work, and that this new mode of commerce – whether it be file swapping, online used book sales, or free online academic journals – is endangering that revenue.

But uppose you were prohibited from selling your vehicle as a used car because your sale would cut into the earnings of those who build new cars. That argument seems pretty ridiculous, but it is essentially the same one being advanced by authors opposed to Amazon selling used books. "Amazon’s practice does damage to the publishing industry, decreasing royalty payments to authors and profits to publishers," the guild wrote in its message. "There’s no good reason for authors to be complicit in undermining their own sales."

I think that a lot of such lobbying is being done by people who have no idea how commerce works. Take me, for example. I pick up a book by John Brunner for a quarter in a used book shop (the real investment, of course, is the time it will take to read the book). I read it, I like it, I pick up a few more used Brunner books, then I start scouring Chapter’s for his latest release. That’s how it works. Cut off used book sales and it’s like you’ve cut off the oxygen. The same logic applies to most content, online or offline. The software I buy is the software I’ve been using for free for a while. The NY Times when I’m south of the border I buy because I’ve become used to reading it for free online. The text I recommend for my class is the one a colleague loaned me over the summer. I don’t know what authors and publishers think will replace the churn of ideas that constitutes a free information exchange, but I can tell you this: if you kill off that churn, you kill off the fuel that drives the information economy.

This paper introduces in an overt way another thread in the discussion: the declining value (and therefore cost) of information. In a connected world, our access to sourcs of information increases dramatically. Where once we were limited by scare radio frequencies and the cost of print publication, in the information age we are limited only by our cognitive capacies. We have gone in a decade from an era of scarcity to an era of surplus, and it makes no sense to continue to pay the same price in an era of surplus that we pai when information was scarce. This paper runs the numbers and arrives at the conclusion that the value (and therefore cost) of information will be (all other things being equal) two times an order of magnitude less than it was before online communications I sent the article to Dave Pell. He cut off my subscription. I moved on to other sources. You can't argue with surplus..

13 Five Choices: Or, Why I Won't Give Dave Pell Twelve Dollars

Written August 21, 2002. Unpublished.

Last week at the New Directions forum in Wisconsin (summarized here) I suggested that online content providers should get used to the idea that they will not make a lot of money from online content.

"Content is of diminishing value," I said. "We can only keep royalties up by creating artificial shortages."

And I received the usual arguments back. People deserve to be paid. People won't produce content unless they're paid. And my suggestion that people get another job – that's preposterous!

But I look now at the trillions of dollars lost by investors when the dot com bubble burst. Much of this money was lost by companies – like Salon, say, or the Industry Standard, or any of the dozens (hundreds?) of companies that went belly up or nearly so – who thought they could make a go of it by offering quality content on the internet.

Content didn't sell on the net, and it won't sell on the net, because the premises of the rejoinder are false: people sometimes don't deserve to be paid, even if they work hard, because the market doesn't work that way. And people will produce content, even if they don't get paid for the content. The online evidence for this is overwhelming, as two billion free web pages will attest.

But still these business plans were launched on the belief that, somehow, this content or that subscription model is, somehow, different. But the interenet is a harsh place, and people learn to become humble in a hurry. Believe me, I know. And so they lost their shirts.

I would like very much to see e-learning avoid this error. This means dashing some hopes (and maybe even a few business plans). It means saying harsh things, sometimes. But if education as an industry loses a trillion dollars, it is not unwitting investors who pay the price. It is colleges, universities, the governments that support them and the students who attend them who must pay the cost.

So, this article. In my email today I received my daily issue of NextDraft, a high quality news summary authored by Dave Pell. Dave has been producing this newsletter for free for some time now, and he would like to make a business of it. I wrote back saying, in essence, the numbers just won't support it.

What I want to suggest is that exactly the same argument applies to online educational content. It's harsh, but it has to be said: you think your content is

unique, but it’s not.

So read on. The numbers are real.

Hiya Dave,

You wrote,

"So I’m looking for places to take NextDraft in the future. One thing I know is that sooner or later, I’ve got to make this more of a business and less of a hobby. It’s too much work for the latter and someone’s got to pay for my satellite television addiction. One idea would be to try to get NextDraft picked up as part of an existing publication. Another would be to include ads.

"And here’s another. What if NextDraft became a pay service? I’m thinking of something on the very low-end, say about a buck a month. There might ultimately be two versions of the newsletter, one paid and one free (the free on would either be much shorter, have no links or come weekly – or something like that. Any ideas on that would be quite appreciated)."

Dave, I won’t be paying 12 dollars a year and I voted "no" on the poll. Not because I don’t like NextDraft. Not even because I think all content should be free. But because I am not willing or able to support this model of paid content. It’s not personal or even professional. And it has nothing to do with me, really.

I subscribe to maybe 100 newsletters. A number of them are dailies, like yours. Others are weekly or even monthly. At $12 a pop (a low end figure; most people seem to expect more than $12) I’m looking at $1200 per year, which in Canadian dollars is more like $1800 per year. Or, for me, $150 per month.

That’s not bad, but if you add to that my access costs of $50 per month, plus the cost of my own comuter and equipment – I just spent another $700 for my wireless home network yesterday – plus odds and ends, and now I’m looking at a pretty hefty bill. Plus the fact that I would have to go out and type my credit card number 100 times (at least I have a credit card; until a couple months ago I didn’t – talk about making online purchasing difficult).

All this for significantly less content than I get in my newspaper (for one year, $172.90) or even the 60 or so channels on my television (one year, $600). And the newspaper doesn’t require me to purchase additional reading devices (though I should include the cost of the coffee I drink with the paper).

As I said, I like your column. I like it about as much as a well written editorial, a story and box score of yesterday’s Montreal Expos game, a column by Alan Fortheringham or Gynne Dyer. In my newspaper I’ll get about 40 or 50 of these items a day (it’s not much of a newspaper). Say 40. That means any given column costs me $4.30 per year.

So if we consider the $12 a year you’re asking for (which is, again, compared to similar subscription requests, very reasonable), it is three times the cost of the equivalent cost of a newspaper article. Not counting the cost of my computer equipment and access. And not even considering the fact that you don’t have to print it on paper and physically deliver it to my door.

This is my problem. Even leaving aside the inconvenience factor, your content would cost me more than other, more traditional content that is of roughly equal value to me. Three times more. And yet it seems to me that the cost of information online should be substantially less than what I am paying for my paper-based content.

How much lower? In various talks I have talked about the target of "two-times order of magnitude." That means that if a column cost me $4.30 per year under the old model, I should be looking at 4.3 cents under the new model. Now you might think that this is absurd, but look even at the situation now. If you look at the average cost of everything I read, say, fifteen years ago and compare it to the average cost of what I read today, you can see that this cost reduction has already been achieved.

As I said, I receive about 100 newsletters. Every one of them is free. I also access dozens of web pages daily, read a number of journal articles, and read about 200 emails (tossing out another 200 spams). I also read three newspapers, which I pay for, the odd magazine, part of a book (on a daily average basis). I subscribe to cable, but my radio access (using an old Rockport radio my mother gave me) is free.

Fifteen years ago, radio (on my old Rockport radio) was free, but everything else cost me money. Oh sure, cable was half what it is now, but I got less than half the channels. Newspapers cost about the same. Fifteen years ago I was averaging more than a book a day, some obtained from the library, but most paid for out of my own pocket (according to the movers, I now own about two tons of books). Free content? Once or twice a month I would get a letter (as opposed to a bill or a flyer). I suppose if I had paid the $20 cents per to send mail I would have received more back.

It should be clear that I pay much less per item today than I did 15 years ago. Easily two-times order of magnitude. Probably more so. So I think I'm on good grounds here.

Now you may say that the information I get is of much lower quality than the information I paid for fifteen years ago. But I would beg to differ. Your own newsletter is a case in point. It is easily of as high a quality as anything in the newspaper now or then. The same can't be said for all the newsletters I receive, but then again, I was never really impressed by Dear Abby yet paid my 2 cents daily for it (fifteen years ago) each day (sometimes two or three times, depending on whch papers I read). I read some discussion from Curtis Bonk yesterday which would certainly equal the quality of the same material were he to put it in a book, but with the added advantage that he said it yesterday, not three years ago.

No, I don't think I need to concede anything on the quality side of things. Television fifteen years ago was even more of a wasteland than it is today. The purchase of books was always risky, especially when you shelled out as much money as I did for cheap science fiction paperbacks. Even supposedly authoritative texts, refereed, audited and selected by a university professor, would contain a certain amount of dreck – so much so, in fact, that we would often skip entire chapters in our study of the work.

So, no. It's not that today's content is of significantly poorer quality. Most of it comes from the same people I would have read in print fifteen years ago, saying the same things (only more recently), in much the same way. I live with a certain amount of poor content, which i dismiss quickly, and I spend most of my days poring over very high quality content. The difference is not the quality. It's that it is cheaper to produce, it's easier to access, and there's so much more of it. That's what drives the cost down.

So finally we turn to the real reason you would like $1 per month for NextDraft:

your desire to make it less of a hobby and more of a business. As you say, it’s a lot of work, and like me, you have expenses. I would certainly agree that you’re underpaid – earning nothing is underpaid by anyone’s scale. You certainly deserve to make a few dollars out of NextDraft, maybe even a living. I know the work that you put into your publication, since I do eactly the same thing for more than 1000 readers with my own newsletter.

You may say, "What’s a dollar – 5 cents an issue? Surely it’s a fair compensation for all the work that I put into NextDraft." But my point is that it’s not fair compensation. The value of your column on the open marketplace is not determined by the amount of work you put into it or even the quality of the content. It is determined by what people are willing to pay, more accurately, what people actually pay, for the column. And columns on the internet average much less than five cents per issue. If we apply the two-times order of magnitude rule, it’s more like five cents per year.

The real issue here for you, of course, is not how much I pay per issue or per year. It’s how much you earn per issue or per year. Let’s run some numbers that would tell us, given the current market for online content, what success would look like. Suppose you wanted to make a living off your column. That’s, say, $50,000 per year (I picked a number that makes everything add up nicely; adjust according to your lifestyle preferences). To earn $50,000 per year, you need one million subscribers. Probably more, because it’s going to cost you a few dollars to send a million emails a day (though, as spammers know, it’s not nearly as expensive as you’d think).

One million sounds like a lot of readers when you have a few thousand subscribers. But ask yourself, how many newspaper columnists have a few thousand readers? Not a one. Daily circulation for newspapers is in the hundreds of thousands. Syndicated columnists can count on a million regular readers. And if we look at a potential worldwide audience of a billion or so people online (give or take 500 million), you can see that you need to reach only a miniscule 1 percent of them to make a living from your column. So it’s not unreasonable that you could make your living off your column. But, with market demand for online columns being what it is, and with your cuirculation being what it is, you have to overcharge to even hope to pay for your satellite dish.

This is the thing. You can say that "you ought to pay" or that you "desrve" to be paid as much as you want, but I am not forced to pay. Should you charge money, I can always decline to subscribe – an option I exercised by selecting "no" in the opinion poll – and to read a cheaper (or free) publication instead. The point is (and I hate to state it so harshly) is that I don’t owe you a living, for the simply reason that I could not possibly pay all the people I would "owe a living" under such conditions. No, the relation between you and I has nothing to do with morality, no matter what the advocates of paid online services say. It is a purely market transaction: you offer to sell me a service at a price, I consider my options, and accept or decline.

The other harsh reality is that, if you do go ahead with your subscription model, you should expect to lose roughly 98 percent of your readers (figures vary depending on who you read). So say you have 10,000 subscribers (which would make yours a very successful internet newsletter, as these things go). You would be left with 200 subscribers. At $12 a pop, you’re looking at $2400 per year. That’s not bad: it will pay for your satellite system.

But don’t forget, that’s a percentage based on the churn your free newsletter has

already generated. People passing it to their friends. People linking to your home page or to your articles. You have to expect this to drop off once you enter the subscription mode. Oh sure, you will still generate some churn your your abbreviated free version, but a lot less, because it’s simply not as good as it was. And sure, people may still link, but many fewer, because people don’t linke to link to a sign-up screen.

You may generate enough new subscriptions to offset the inevitable attrition. Hard to say. But it seems likely, unless your newsletter is so much better than the free content that it becomes a "must-read," that you will be fighting a never ending battle to obtain subscriptions. This means advertising and other expenses that you may not have counted on. Just to maintain your current 200 subscribers you may find yourself eating into more and more of your $2400 annual income. And how are you going to advertise to a market that spans the globe?

Your major issue isn’t the fact that I won’t pay. It is that you are by no means alone.There are hundreds of thousands of blog writers (half a million, according to a recent (free) MSNBC article. On top of that, hundreds of thousands more authors of various sorts, including university professors (each of whom thinks he has the one best way to teach calculus, and that the would ought to pay for it), politicians (who will now and always write for free), sports fans, pundits and consultants, and more. heck, there are even software programs out there that will do much of what you already do - gather relevant news headlines and display them on a page. I know you add a lot of valuable commentary to what you write. So do I - and yet one of my competitors (or it would be, if I were a commercial service) simply harvests topical headlines from PR newswire and has five times more subscribers than I do. It may not be fair that your competition numbers in the thousands, uses automated tools, and produces a lower quality product, but that’s life on the internet.

So. You want to make a living with your hobby. In my view, you have limited options.

1. Lower your production costs. That’s what some services do. Use a reasonably well-programmed harvester to do most of the work, then take a half hour to add some comments. That way your hobby is more like a hobby again and you don’t need to worry about making a living from it.

2. Increase your volume. You may not get a million people to read your one newsletter, but if you managed to get 10,000 people to read 100 newsletters you are obtaining the same result. Unfortunately, you’ve just increased your workload by a factor of 100, which probably not what you had in mind. You will have to lower your workload by the same amount.

3. Get a million subscribers. For you, probably not a real possibility unless you were to join some sort of larger service that really does have the capacity to reach a million people. You could get a million subscribers by becoming a columnist for the New York Times, for example. For most of us, this probably isn’t a live possibility. But it will work for some people.

4. Create higher value content. Let’s face it, you are putting out the equivant of a newspaper article. It’s good, it’s useful, but it’s not a "must-read". You can’t charge more for it because you wouldn’t be able to charge more for it in the traditional print-and-paper world. But if you were able to generate absolutely unqique content that nobody else could offer, you would then have effectively eliminated your competition and thereby increased your market value. That’s how

companies get away with charging money for stock quotes, for example (of course, having an artificial monopoly on that content helps). Consultants such as Forrester charge more research results that only they have (because they did the research themsleves). For most of us, that too isn't really an option: few of us own stock markets or research institutes.

5. Or – as I suggested (to much derision) at a conference the other day, get another job. The fact is, online content production doesn't pay the bills. But it can act as a loss-leader for the provision of other services. By reaching a wide audience with your free online content you are able to display – almost without cost – your unique expertise or skills. You may be able to obtain employment based on these credentials. Or secure consulting gigs or speaking fees. Or you may do some writing for hire for a firm that could use your easy touch with a typewriter. If you have sufficient expertise and credentials, teach an online class.

To wrap up my discussion, let me sketch a small analogy that, in my view, nails it down.

I play darts. I work very hard at my dart game; I practise for a couple of hours a day. I am actually very good at darts and can walk into most pubs and whip the locals. I have invested a lot of my time, energy and sweat into becoming a good dart player. Just like the hundreds of thousands of people who play baseball, basketball or hockey. We do it for fun, and there are enough of us who are good enough to say that we could be on the verge of turning pro.

We're good, but it turns out that there just isn't enough of a market for all of us. No matter how much I practise playing darts, I'm only going to make a few hundred dollars a year, no matter how much I deserve more for all the work I put into it. It's not fair, but sports isn't fair, and professional sports are even less fair.

So my choices as a dart player are much like yours above. I could practise less (of course, my game might not be as good). I could play in more tournaments – though probably I could never play in enough to earn any sort of reasonable money. I could convince more people to watch darts – but I would need to own ESPN in order to pull that off. I could improve my game and maybe win the big prizes. That means being one of the top ten players in the world, though. Or I could write my darts off as a hobby and derive any value I can from it – the contacts that I make, the line on my resume, the improvement in my character, or the use of darts as a great analogy from time to time as I pursue my day job.

It's not a question of right or wrong, fair or unfair. It's just that there are too many dart players and too few people interested in paying to watch darts. There's nothing I can do about it, but I can say this as certainly as I've said anything else: if I charged $12 a year for people to watch me play darts, nothing's going to change. I still won't get paid any more than I'm making now. And I will have succeeded only in annoying my friends.

Epilogue

See, it is easy to stand at the front of the room in a (sparsely attended) forum and denounce my suggestions as preposterous. It is easy to say that content authors deserve more money, and that anyone who thinks otherwise just isn't in touch with reality.

But the hard numbers don't support the case. The only way to raise the price of online content is to severely restrict the supply. That's why so many LMS and LCMS vendors are signing "exclusive deals" with publishing companies. They know you

won’t pay several hundred dollars a pop for some B-grade online learning content unless (a) it’s the only material available on the subject, and (b) you need it.

It won’t last. And even if it does last, I want all of you who are would-be educational content authors to run through the following calculation:

How much did you earn from publishing journal articles and books last year? Take one percent of that? That’s how much you will earn from online content.

Now, for those of you creating your online magnum opus, you may be thinking something like this: but I’m already making more than that. I got a $40,000 grant or project or dispensation to create this content.

Well yes. You got $40,000. But the content is unlikely to earn $40,000 back. Not in a free and open market. Do some real calculations: outside of your own course, how many institutions and professors used your online content? How much did they pay for it? How much did you actually put in the bank over and above your own salary?

See, everybody is in the "loss-leader" frenzy right now. The golden era before the bubble bursts. They’re drawing a salary, but very few are really selling content. Oh sure, the commercial vendors have established a (temporary) corporate market for (custom) content. But the colleges and universities?

Like it or not, your five options are the same as David Pell’s.

1. You can lower your production costs by employing content authoring tools, reusable learning objects, and low paid (graduate student) labour. But this impacts the quality of what you offer.

2. You can dramatically increase your production of courses. This means lowering production costs. And even then, you probably won’t be able to lower production costs enough.

3. You can get a million students (if you’re the Open University or the University of Phoenix) per course.

4. You can create higher value content, content so good and so unique that people will have to pay for it. But fair warning: it had better be really good – better than MIT’s, which is already online for free.

5. Or you can give up on the dream of making money from content and get back to your real job, providing an education. Your content will get people in the door. And it will make your job of providing a service easier. Cheaper for students. But it won’t pay the bills.

In May, 2000, I wote my paper Learning Objects. It was a summation of the work that I had started with The Assiniboine Model and The Future of Online Learning. The paper had a wide readership and was arguably one of my most important publications. It appeared in IRRODL the following year. Some time later, Maxim Jean-Louis of Contact North asked me for an updated version of the paper. Many months later, I sent him this work, renamed and expanded to reflect developments in the field and a growing awareness of the network since I had published the original piece. I have no idea what the actual publication ate was, or even the actual completion date; it became The Project That Would Never End. I presented it at NAWeb in October, 2002 and finally decided I was finished writing it. Though a much more internally consistent and complete treatment of the same topic as Learning Objects, it achieved nothing like the same readership. It's all in the timing, I guess. It is included here as a survey of the field, a summary, if you will, of 'what is known' about learning objects and their distribution. It's almost impossible tounderstand the material that follows without this basis, so though long, it forms a central and essential component of this book.

14 The Learning Object Economy

Written October, 2002. Published by Contact North, 2003.

Learning (lûrnng)

n.

The act, process, or experience of gaining knowledge or skill.

Knowledge or skill gained through schooling or study. See Synonyms at knowledge.

Psychology. Behavioral modification especially through experience or conditioning.

Object (bjkt, -jkt)

n.

Something perceptible by one or more of the senses, especially by vision or touch; a material thing.

A focus of attention, feeling, thought, or action: an object of contempt.

The purpose, aim, or goal of a specific action or effort: the object of the game.

Grammar.

A noun, pronoun, or noun phrase that receives or is affected by the action of a verb within a sentence.

A noun or substantive governed by a preposition.

Philosophy. Something intelligible or perceptible by the mind.

Computer Science. A discrete item that can be selected and maneuvered, such as an onscreen graphic. In object-oriented programming, objects include data and the procedures necessary to operate on that data

Economy (-kn-m)

n. pl. e·con·o·mies

Careful, thrifty management of resources, such as money, materials, or labor: learned to practice economy in making out the household budget. An example or result of such management; a saving.

The system or range of economic activity in a country, region, or community: Effects of inflation were felt at every level of the economy. A specific type of economic system: an industrial economy; a planned economy.

An orderly, functional arrangement of parts; an organized system: "the sense that there is a moral economy in the world, that good is rewarded and evil is punished" (George F. Will).

Efficient, sparing, or conservative use: wrote with an economy of language.

The least expensive class of accommodations, especially on an airplane.

Theology. The method of God’s government of and activity within the world

14.1 Abstract

The intent of this essay is to provide a comprehensive overview of learning objects and related topics for the non-technical reader.

First, some common arguments identifying a need for learning objects are described and through an examination of the problems identified a description of learning objects is obtained.

Second, the development of learning objects is placed into a theoretical context by identifying the underlying concepts in computer science, standards initiatives and distance learning theory from which they evolved.

Third, learning objects are looked at from a practical point of view. Tools and techniques for creating learning objects are described. The method of preparing learning objects for reuse is outlined. Then the delivery of the objects in a learning environment is described.

Fourth, the learning object economy as a whole is developed. This is the system of learning object repositories, distribution systems and rights management. A variety of initiatives and technologies are introduced.

14.2 A. The Need for and Nature of Learning Objects

This section describes the need to design online courses in such a way as to reduce costs without diminishing the value of a university education. We need to do this by extracting what these courses have in common and by making these common elements available online as learning objects. This section describes the need for learning objects and then offers a definition of learning objects drawn from the description of that need.

14.2.1 i. The Idea of Learning Objects

There is no consensus on the definition of learning objects. Definitions abound and numerous analogies are employed to elucidate the concept. The basic idea, by

virtue of its simplicity, allows wide latitude for interpretation. Learning objects are intended to support online learning. They are intended to be created once and used numerous times. Because they are delivered online, they are intended to be digital objects. And because they are used in learning, they are intended to have an educational component.

The Lego Analogy. A popular metaphor for learning objects is a set of Lego blocks. Though the use of a simple design, Lego blocks may be reused and combined into a variety of different combinations. As Wiley writes, the idea is that any Lego block can be combined with any other Lego block, Lego blocks can be assembled in any manner you choose, and Lego blocks are so simple anyone can use them. (Wiley, 1999)

The Atom Analogy. As Wiley notes, the Lego analogy is to broad for many conceptions of the learning object. Some learning objects do not fit well together. The nature of learning objects may restrict what can be created with them. And learning objects are probably more difficult to use than Lego blocks. Accordingly, Wiley proposes that learning objects are more like atoms than Lego blocks. Learning objects are the basic components of, say, online courses. But different learning objects function in different ways.

Consensus? Digital reusable objects to support learning. Most discussions of learning objects agree that learning objects are digital, reusable, and are intended to support learning. The IEEE (2002) defines learning objects as "any entity, digital or non-digital, which can be used, re-used or referenced during technology supported learning." Wiley (2000) settles on a definition of a learning object as "any digital resource that can be reused to support learning." Even so, as Wiley comments, "the definition is broad enough to include the estimated 15 terabytes of information available on the publicly accessible Internet."

A functional definition. Probably no definition of learning objects will ever be sufficient; there will always be those that say the definition allows too much or too little. Part of the purpose of this paper is to approach the subject of learning objects from a different direction: rather than to say what they are, this paper attempts to show the problems learning objects are intended to solve and the manner in which they are used. In other words, this paper is intended in part to provide a functional definition of learning objects.

The Learning Object Economy. The functional definition of a learning object offered at this juncture is that a learning object is anything that is exchanged in what may be called the learning object economy. As this paper will show, the learning object economy is a complex of networks and systems intended to support learning, a vast learning mill. Learning objects, whatever they are, are the grist that circulates through this mill. Learning objects are the raw material used to support learning; the learning object economy is the network designed to produce and distribute that raw material, and the eventual definition of learning objects will depend on what people want to receive at the output end of that mill.

14.2.2 ii. The Case for Online Learning

We need accessible and affordable learning. The need for and usefulness of online learning is today no longer in question, but to understand the need for learning objects it is useful to reflect on the factors that led to the development of online learning. And though the availability of the technology was a key factor, the primary driver behind the development was a widespread need for accessible and affordable learning.

Dimensions of accessibility. In a world where many or most people have access to the internet, online learning promises to make learning more accessible. Accessibility has numerous dimensions. Among the most frequently cited (as in DOI Online, 2002) are timeliness (online learning may be used any time of the day or night), accessibility (online learning many be access from almost anywhere), and flexibility (online learners can proceed at their own pace).

Accessibility as choice. To a large degree, accessibility may be defined as choice. As Vail (2001) writes, "Students turn to online classes and schools for varied reasons, but they have one thing in common: They all want or need something that's not easily available in the traditional brick-and-mortar school building. Students in rural communities can take classes such as Latin or AP calculus that their schools are too small or too poor to offer. Sick or hospitalized students can finish their class work without falling behind. Gifted students, students who have problems in the regular classroom, students traveling with their parents."

Accessibility as lifelong learning. Accessibility may also be defined as having the opportunity to continue learning while employed. An Open University study noted for example that "an increasing number of people do not want to study for three years before embarking on a career. Instead they want to combine starting a career and studying as soon as they finish school." (Major, 2002) As Human Resources Development Canada (HRDC 2002) noted "With 50 percent of the workforce of 2015 already in the labour market, and a smaller projected youth cohort, Canada must take action now to ensure we can meet the skills needs of the economy.. A broad-based, accessible and comprehensive adult learning system must be a prominent feature of the country's learning infrastructure."

14.2.3 iii. The Cost of Online Learning

Traditional courses are typically created by a single artisan. Though instructors in traditional classrooms use common course materials such as textbooks and journal articles, each time a course is offered by a school, college or university, it is created from scratch. And although instructors sometimes use core curricula and often use the same course outline one year to the next, these are adapted and localized on a case by case basis. The task of creating a course in the traditional classroom, therefore, resembles what may be described as a cottage craft industry: it depends on and reflects the skills and inclinations of an individual artisan.

Online courses are also typically created by an individual artisan. Modern schools, colleges and universities developing courses for online delivery have migrated this strategy into their internet offerings. Although supported by teams of designers and web specialists, courses are essentially the product of individual teachers or professors. And though common materials, such as course packs or other online learning resources, may be used, the online course is essentially created from scratch each time it is delivered. Like traditional teaching, online teaching today is labour intensive, and therefore, expensive.

Example: Tony Bates. Examples abound, but this process is typified in Managing Technological Change by Tony Bates. According to Bates, a course costs $24,4000 to develop, taking 30 days of a subject expert's time seven days of an internet specialist's time, and additional time for copyright review, academic approval and administration. Delivery costs an additional $13,161 per year. To cover these costs, students in Bates's course pay $463 or $695 in course fees, plus an additional $177 for required readings.

Example: A generous estimate of 5 man-hours per manual page would result in 1000 man-hours for a 200 page manual, which could be covered in 40 hours. At a burdened rate of $60/hour, the in-house development cost would be: $60/hour development cost x 1000 hours = $60,000. If an outside consulting firm did the job, it would cost: $120/hour x 1000 hours = $120,000. (Kurtus, 2001)

Example: Using the 50 percent time-reduction estimate specified in step 1.4, it is assumed that the CD-ROM training would only require seven student hours. Therefore, estimated cost of development is 7 hours x $50,000 per hour = $350,000. (Kruse, 2002)

Online courses are therefore at least as expensive to develop as traditional courses. Almost all online course developers use the design model Bates describes. It involves a course being developed from scratch, using nothing more than a traditional university course or a good textbook as a guide. The course author typically authors all the content, including examples and demonstrations, quizzes and tests. Because of the cost of development, there is little use of course specific software or multimedia. The course is then offered to a small number of students over a limited time, resulting in course fees that are comparable, if not greater than, traditional university course fees.

We can do so much better than this. We need to design online courses – even university courses – in such a way as to reduce these costs without diminishing the value of a university education. We need to do this by extracting what these courses have in common and by making these common elements available online.

We can create better online course materials. Consider the Teacher's Guide to the Holocaust. This site consists of dozens of resources on the Holocaust may be used and reused by any teacher approaching the subject. Each of the 'class activities' could be treated as an individual learning object. The Holocaust is a very large subject – much larger than sine waves – and is appropriately divided into many components. But it is far easier, and of far greater quality, to assemble a lesson or series of lessons from these materials, than to create something from scratch.

Example: Hamlet. There is not of course one single description of Hamlet, but there is only one text of the play Hamlet and it is not a stretch to envision a definitive online multimedia edition. A course specializing in Hamlet would employ the digital Hamlet as a central resource, and incorporate as well essays, discussions and articles from scholars around the world. Such an edition would not only contain the text, it would also contain video clips, audio clips, commentary from selected sources, pop-up glossaries, and more.

We can lower the cost of learning: It is not a stretch to imagine a multimedia company spending a million dollars on such a production. Assume that Hamlet is taught in 10,000 schools, colleges or universities around the world (hardly a stretch). Assume 20 students per class (an underestimate, to be sure!). At $5 per student, the company would make it's million back in one year! The economics are very good, and this excellent resource would be cheaper than even the book alone.

14.2.4 iv. The Argument for Learning Objects

Define learning objects by defining the problems they solve. What are learning objects? To answer this question is best approach is to describe what problems learning objects are intended to solve, and thereby, to describe what learning objects are designed to do. In this first section, then, we begin by describing the problems. In later sections, we will look at approaches to solving these problems. This will in turn allow us to extract a list of defining features of

learning objects.

The problem: online educational content is expensive to produce. Online educational content is not cheap. Even a plain web page, authored by a mathematics professor, can cost hundreds of dollars if you take into account server costs and the professor's salary. Include graphics and a little animation and the price can double. Add an interactive exercise and the price can be quadrupled.

If each institution produces its own materials the cost multiplies. Suppose that one description of the sine wave function is produced. A high quality and fully interactive piece of learning material could be produced for, say, a thousand dollars. If a thousand institutions share this one item, the cost is a dollar per institution. But if each of a thousand institutions produces a similar item, then each institution must pay a thousand dollars, or the institutions, collectively, must pay a million dollars. For one lesson. In one course.

The cost is reduced by sharing similar learning materials among institutions. The economics are relentless. It makes no financial sense to spend millions of dollars producing multiple versions of similar learning objects when single versions of the same objects could be shared at a much lower cost per institution. There will be sharing, because no institution producing its own materials on its own could compete with institutions sharing learning materials.

To solve the problem of cost, learning materials will be shared. Economics, then, dictate that we need to be able to share learning materials between institutions over the internet. But this raises a host of issues. What sort of materials can be shared? How might they be created? How do we account for the content that does change from one institution to the next? And from the millions of objects on the internet, how can we find the one item we need for a particular course at a particular time? These issues and more need to be resolved, and so we need to look at the problem of sharing more closely.

http://www.careo.org/vision/TomCarey.ppt Tom Carey, A Vision for Online Learning Objects

at the University of Waterloo, Building a Vision for Sharing Education Objects in Alberta, 2001?

14.2.5 v. Courses? No, Not Courses

The problem: what will be shared? If we accept the premise that institutions will share learning materials, then we need to ask, what will they share? What size will they be? This is sometimes known as the problem of the granularity of learning objects.

Postulation: we will share courses. The answer that intuitively offers itself is: courses. Existing listings of online learning materials, say TeleCampus (see below), list only courses. Good listings, they are divided into subject areas, where each subject page contains a list of similar courses offered by different institutions.

Courses are what students purchase from institutions. These directories are directed at potential consumers of learning material, that is, students. Students are typically motivated by an interest in a topic and select courses from the list of offerings in that topic. Moreover, students are typically offered learning materials in course-sized units, and attempt to complete degree of diploma programs defined as sets of related courses.

Institutions already share courses to a degree. Why, then, would institutions not

share these courses? To a certain degree, they already do so. Most colleges and universities define course articulation policies, whereby a course completed at one institution is accepted for credit at another institution. A good example is the Baccalaureate Core Course Equivalency defined by Oregon State University for courses at thirteen regional community colleges. http://www.orst.edu/Dept/admindb/arttable/scr1140_arttab.htm.

Course articulation is an example of sharing courses. Course articulations are the result of complex negotiations between teams of academics. Consider, for example, the information contained in the Illinois Mathematics and Computer Science Articulation Guide. http://www.imacc.org/articulation/ To count as equivalent credit for, say, a trigonometry course, a candidate course must require certain pre-requisites and contain material covering a certain set of topics.

However, course articulation is complex and regional. Because of the regional nature of course articulations – it is notable that Oregon State University has made no attempt to articulate courses offered by, say, community colleges in Florida – and because of the detailed topic-by-topic definition or articulation agreements, course sharing between institutions is difficult to define and maintain. It is unlikely that any course could be shared by any significant number of institutions in different states or different nations.

Courses offered by institutions vary widely. We see this disparity reflected on online course listings. Returning to the Telecampus guide we find twenty separate history courses listed. No two of the courses share the same name. And though a number of courses focus on the same region and time period, no two of the courses share the same contents. This is true to more or less a degree across all subjects and across all institutions. Although courses may share elements in common, it is rare to find two courses from two institutions that share the same, and only the same, set of elements.

So we will not share courses. Courses themselves are not suitable candidates for sharing. Yet the dominant form of online educational today is the course. So it should come as no surprise that there is very little sharing of educational resources, even online resources, despite the tremendous cost savings.

Conclusion: we will share parts of courses. What needs to be shared may be best described as parts of courses, or more accurately, course components. From this it follows that we need not only collections of course components but also some mechanism of assembling course components into complete courses. This may be thought of as the problem of packaging learning objects. In the sections that follow below we will first explore the idea of what sort of things can be shared as course components, and then we will look at the problem of packaging.

14.2.6 vi. Sharing the Old Way

We best understand sharing by looking at existing examples of sharing. To best understand the concept of sharing course components, and go get an intuitive understanding of what may constitute a learning object, it is useful to look at how and why learning materials are shared in traditional classrooms. It is important to review the "old ways" of sharing resources not only to show that resource sharing is an established fact in today's classrooms, but also to point to some of the elements of resource sharing already in place.

Today's classrooms already share learning materials. If we describe 'sharing' as meaning 'one centrally produced resource used by many,' then today's classroom is already an example of extensive resource sharing. Various publishers and content

producers produce resources centrally and distribute them to classes around the world. And while many of these resources are distributed for free, the majority of shared resources in classrooms are purchased from their respective producers or intermediaries.

An example of sharing: the textbook. The clearest example of resource sharing "the old way" in today's classrooms occurs through the use of textbooks. These resources bear all the hallmarks of sharing: they are centrally produced and obtained as needed by classroom instructors around the world. In many cases, the information in textbooks is so commonly used the work becomes standard.

An example of sharing: classroom displays. But textbooks are just one type of item among many that are shared by classes around the world. No K-12 school is complete without a set of wall maps in geography classes, periodical tables of the elements in science classes, and sets of large block letters for the early years. A rich and useful set of classroom displays is distributed by organizations as varied as astronomical societies, museums, and publishing companies.

An example of sharing: multimedia. In the area of multimedia, teachers employ a wide variety of centrally produced materials including filmstrips and videos, CD-ROMs and other software, presentation graphics and even complete learning resources, such as are produced by Plato.

Sharing today involves the buying and selling of learning materials. Neither the producers nor the consumers of those resources would describe the distribution of textbooks, classroom displays or multimedia "sharing." Textbook publishing and sales, especially, is a lucrative industry. The National Association of College Stores estimates U.S. / Canadian college store sales to be $8.959 billion for the 1998-99 academic year. Nonetheless these classes are sharing resources as defined. They are produced centrally and used by many institutions.

Sharing today involves decomposition. Instructors frequently employ only components of purchased learning materials in their classes. Many course syllabi require that students obtain more than one textbook. They may use, for example, only a few chapters out of a textbook. In class, they reassemble these selected materials in a way that meets their instructional goals. (Wiley, 2000)

Sharing today involves sharing parts of courses. In many cases, the resources sold by publishers and distributed to classrooms are not entire courses, but rather, components of courses. This is most clearly the case for classroom aids such as wall maps and posters. Sometimes also students purchase only parts of courses, such as lecture notes or workbooks. And students frequently photocopy only parts of books (or parts of journals) in their research and reading.

14.2.7 vii. Contemporary Online Sharing

Many agencies offer educational materials for sharing, but problems exist. In the traditional classroom, course components such as textbooks, classroom aids and multimedia are bought and sold and then combined by teachers and students to support classroom instruction. On the internet, though most educational institutions offer complete courses only, many other agencies have started offering smaller, more portable learning materials. These materials fall short of what we will later define as 'learning objects', but they offer some insight as to the direction and potential of online resources.

Canada's SchoolNet. In Canada, the leading learning resources portal is probably Canada's SchoolNet. A list of resources is displayed, each with a short description and a link to an external website. SchoolNet also provides information

about each site and provides an "advanced search" using metadata. Each resource in the "curriculum" area is approved by a professional "pagemaster". For the most part, however, SchoolNet links to institutional home pages, and not to learning resources per se. Teachers using the SchoolNet service must still search through these sites in order to locate suitable materials. http://www.schoolnet.ca/

Merlot. Linking directly to learning resources themselves is a site based in the United States and maintained by the Educational Object Economy Foundation. Merlot currently lists more than 2,000 learning applications that can be accessed via the world wide web. These applications are specific materials on specific topics; for example, Merlot lists such items as "Chaucer", "The Great 1906 Earthquake and Fire" and RSPT Expansion (Perturbation Theory). Materials are sorted into category and subcategory and have been contributed by educators from around the world. Educators attempting to use Merlot's resources, though, will still experience frustration. Although the topic hierarchy is more detailed than SchoolNet's and although much more focused resources are listed, educators would still have to spend quite a bit of time searching for materials. http://www.merlot.org

MarcoPolo. MarcoPolo is a compilation of teaching resources from six educational institutions which provide free internet content for K-12 education. What the six partners have in common, and what makes this an important and interesting development in online learning, is an adherence to national curriculum and evaluation standards in the subject areas. Material is categorized by grade level and individual items are matched to individual learning topics. Despite its strengths, however, MarcoPolo is a closed project; only the six member institutions contribute content. There is no centralized search facility and no metadata listings for the resources. http://marcopolo.worldcom.com/

XanEdu. Xanedu is a learning resource site that collects articles from journals, magazines and other resource providers. Instructors may compile 'course packs' consisting of collections of these materials; students who subscribe to XanEdu may access these course packs. The materials are sorted by category and may also be located using a search mechanism. Like MarcoPolo, however, XanEdu is a closed project. It draws materials only from selected publishers. And while it allows subscribed students to browse through its materials, the vast bulk of resources available on the internet cannot be found through XanEdu. http://www.xanedu.com/

Problem: it is difficult to locate relevant learning materials. The internet contains a wealth of learning materials. But even with the help of portals, these learning materials are hard to find and hard to use. The portals need more robust mechanisms for updating and submissions. They need much better systems of categorization and searching. They need to be tied more closely to learning objectives, but in such a way as not to be tied to a specific curriculum. This would allow materials directly relevant to a given course topic to be quickly located. (Schatz, 2001)

Problem: existing portals offer access to only a fraction of available materials. Though the resources offered by learning materials portals are very good, and in some cases, very comprehensive, no portal offers more than a fraction of the materials available on the internet. Materials available from one portal are not available from other portals. And because publishers sign exclusive agreements with certain portals, they are blocked from wider access except through that portal.

Problem: there is no consistency in the materials offered. An even greater weakness appears when we look at the collective set of learning resources (or

applications, as Merlot calls them) offered by these. It is almost not possible to identify consistency in format, scope, methodology, educational level or presentations. Some resources include lesson plans, but many others do not. Some are authored in Java, others in HTML, and others in a hybrid mixture known only to the author. Some involve ten minutes of student time, others would occupy an entire day. And there is no structured means for an instructor to know which is which.

14.2.8 viii. What We Need

Learning objects are defined by the problems they solve. What would we need to implement the sharing of course components online? We would need something similar to the initiatives described in the previous section, but something that addresses the weaknesses of those initiatives (and in fairness, each of these initiatives is taking steps to address these weaknesses). The description of an online entity that addresses these problems forms the basis for a definition of learning objects.

Learning objects are sharable. By ’sharable’ what we mean is that a learning object may be produced centrally and used in many different courses. Sometimes people speak of this criterion by saying that learning objects must be reusable. This is accurate to the degree that it means that learning objects may be used over and over again. But equally important is the idea that they are used by different educational institutions.

Learning objects are digital. By ’digital’ what we mean is that they can be distributed using the internet. While for the sake of argument some people could talk of physical entities (such as textbooks or maps) as learning objects, such objects cannot be used online and therefore are not part of an online course.

Learning objects are modular. A learning object is not an entire course, it is a part of a course. Therefore, in order to create an online course, learning objects must be assembled or packaged into a larger entity. This what we mean by ’modular’ – that collections of learning objects may be assembled into a single, larger unit. This in turn means that, as Longmire asserts, learning objects must be ’free standing, nonsequential, coherent and unitary." (Longmire, 2000)

Learning objects are interoperable. By ’interoperable’ we mean that learning objects produced by different publishers, or available through different repositories, may be packaged together into a single course. An instructor creating a course using learning objects must be able to select from all available learning objects, not merely a selected subset of proprietary learning materials offered by a single provider. Or as Singh writes, "the...framework must allow content and their data to be exchanged and shared by separate tools and systems connected via the internet." (Singh, 2000)

Learning objects are discoverable. By ’discoverable’ we mean that the appropriate learning object for any given instructional application can be located in a reasonable amount of time by a person who is not necessarily an expert at searching the internet. Just as an average person could go into a library and, using the catalogue system, locate a particularly useful book, so also an average person should be able to go online and locate a particularly useful educational resource.

In conclusion, learning objects are digital materials used to create online courses where these materials are sharable, modular, interoperable and discoverable.

14.3 B. Learning Objects from a Theoretical Perspective

The design of learning objects is similar to the design of software objects in computer programming. In this section, we look at some of the theoretical assumptions underlying modern software programming.

14.3.1 ix. Course Construction and RAD

Today's online courses are like old computer programs. People typically think of an online course as being similar to a textbook, or at best, a classroom where a course is being delivered. But from the standpoint of online course design, it makes more sense to think of an online course as being similar to a computer program. This is especially evident when the problems facing early computer programmers and computer users are compared to the problems facing today's online course designers. Early computer programs were written from scratch. They were expensive and time-consuming to create. Moreover, they didn't work with other programs: a document crated by one program could not be read by another program.

Modern programmers use rapid application design (RAD). Software engineers have long since learned that it is inefficient to design applications from scratch. Educators need to learn design techniques learned by the software industry long ago, and in particular, they need to learn a concept called 'Rapid Application Design' (RAD). Rapid Application Design is a process which allows software engineers to develop products more quickly and of higher quality. RAD involves several components, including a greater emphasis on client consulting, prototyping, and more informal communications.

Modern programmers use programming environments. In modern programming, the engineers' re-use of software components within the context of a CASE (computer-aided software engineering) environment. The idea of RAD for software development is that a designer can select and apply a set of pre-defined subroutines from a menu or selection within a programming environment. A good example of this sort of environment is Microsoft's Visual Basic, a programming environment that lets an engineer design a page or flow of logic by dragging program elements from a toolbox.

Analogies: the well prepared chef and mechanic. Similar methodologies exist for a wide variety of creative or constructive tasks. A professional chef, for example, will carefully design a kitchen environment so that when he is called upon to create Crepes Suzette, the essential ingredients – including pre-mixed recipe ingredients. Auto mechanics also work in a dedicated environment and also have at hand every tool and component they may need to fix anything from a Lada to a Lamborghini.

RAD will be applied to course design. Online course developers, pressed for time and unable to sustain large development costs, will begin to employ similar methodologies. An online course, viewed as a piece of software, may be seen as a collection of re-usable subroutines and applications. An online course, viewed as a collection of learning objectives, may be seen as a collection of re-usable learning materials. The heart – and essence – of a online course design is the merging of these two concepts, of viewing re-usable learning materials as re-usable subroutines, applications and documents assembled by application specialists in a computer assisted software environment.

RAD is being used in corporate learning already. Educators in the corporate and software communities have known about this concept for some time. As Wayne Wiesler, an author working with Cisco Systems, writes, "Reusable content in the

form of objects stored in a database has become the Holy Grail in the e-learning and knowledge management communities."

14.3.2 x. Object-Oriented Design

Object-oriented programming is essential to RAD. At the core of Rapid Application Design, and therefore central to construction and organization of learning objects, is another concept from computer programming, object-oriented programming (OOP). The idea behind object-oriented programming is that bits of software common to many computer programs are designed as self-contained entities (or 'objects') when are then used by different computer programs. It is these objects that are assembled by an application specialist.

Example: Windows task-bars. Any person who used Windows is familiar with objects. At the top of Windows programs is a solid bar, called the 'task bar', in which the name of the program or document is displayed and which contains the 'maximize' and 'minimize' buttons. When a software engineer designs a program screen for a Windows application, the engineer does not write dozens of lines of programming specifying the task bar's location, dimensions and colour. He or she, working within a visual environment, simply selects and drags the 'task bar' icon onto the page being designed.

Example: JavaScript alert box. In a similar manner, a person using Javascript to design a web page application does not write detailed programming specifying the size, location and colour of the alert box that pops up on web pages. The JavaScript programmer simply writes a single line of code creating the alert box and giving it some text to display.

Example: the Student Object. The task bar and the alert box are examples of objects. In a similar manner, software objects can be used in online courses. Suppose, for example, a course designer wanted some educational text to refer to a student by name. When creating the document, the designer would first create a 'student object.' When created, the student object automatically retrieves information about the student, for example, the student's name, and inserts it into the document text.

Objects are defined as prototypes. To generate a student object, a programmer designs a prototypical student and for it properties common to all students. Many properties of the prototypical student would be undefined, however, such as the student's name, age, or phone number. These unknowns would be given placeholder values (or 'defaults') until they are defined. When a program needs to work with a student, it refers to the prototype and 'clones' a copy of the prototype in the computer's memory (it's actually called 'cloning' in computer science – in perl the prototype is cloned and 'blessed' to reserve its place in memory). The newly cloned prototype is given a name, and then values or attributes are assigned to it. For example:

Objects are used by designers working within the programming environment. When a designer needs to refer to a student, the designer refers to the prototype and 'clones' a copy of the prototype in the computer's memory (it's actually called 'cloning' in computer science – in perl the prototype is cloned and 'blessed' to reserve its place in memory). For example, the designer may click and drag the 'student' icon onto the page being designed. The newly cloned prototype is given a name, and then values or attributes are assigned to it by the program.

Objects can perform functions. The course designer can make cloned objects do things by referring to pre-defined functions in the object (or, in computer

terminology, 'methods'). To have Fred Smith register in a course, for example, we would execute a command that tells the student object to perform a function called register(). The course into which Fred is registering is itself another object. When the function register() is executed in Fred, the Fred-object communicates with the course-object and executes a related function in the course object, add_student().

Objects interact with each other. Objects may interact – or more generally, be related to each other, in many ways. The most useful and common form of interaction is the containing interaction. Just as Fred may contain various other objects (such as a heart or a liver, most obviously, but also $4.95 in change, a six inch ruler and a pager), one object may in general contain one or more other objects. A course may contain students, for example. Or a course may contain units or modules. A unit may contain a test. Each of these items is an object, defined from a prototype, which may interact with other objects in predefined ways. In a course which contained both a unit test and a grade book, for example, the unit test could interact with the grade book. What would happen is that Fred (the 'student' object) would interact with the test (the 'test' object'), which in turn would interact with the grade book (a 'grade book' object).

14.3.3 xi. Open Standards

Open standards are like common languages. A third major concept drawn from the world of computing science – and especially from the recent emergence of internet technologies – is the use of open standards in course construction. An open standard is like a language understood and used by everyone. Just as, for example, the meanings of such terms as 'Paris', 'the capital of France', and 'European' are understood by almost all speakers of English, so also in an open standard are the meanings of terms and definitions widely understood and shared.

Example: HTML. The open standard with which most online educators are familiar is Hypertext Mark-up Language, or HTML. This language is a shared vocabulary for all people wishing to read or write internet documents. The term '<h1>' is commonly understood as a header tag; the term '<I>' denotes italics.

Open standards are contrasted with proprietary standards. Open standards may be contrasted with proprietary, or closed standards. Consider a document written in an older version of MS Word, for example. This word processing program used a special set of notation to define italics, bold face, and a wide variety of other features. Because other software manufacturers did not know these standards, only people using MS Word could read a document written in MS Word.

Open standards enable programs to interact with each other. The purpose of open standards is to allow engineers from various software or hardware companies develop devices and programs that operate in harmony. A document saved in an open standard could be read, printed or transmitted by any number of programs and devices.

Three major types of open standards. There are three major types of open standards. The transport protocol defines how digital material is transported over the internet. The internet is based on transport protocols such as HTTP and FTP. The second major type of open standard is the mark-up language, which defines how parts of documents should be identified and displayed. HTML is a type of mark-up language. The third major type of standard is the program interface. The program interface defines what functions (or methods) can be called by one program in another. A web browser uses a program interface, for example, when it displays a

plug-in such as a Flash animation or a Java program.

Learning objects require mark-up languages and program interfaces. Insofar as online learning is delivered using the internet, it can use common transport protocols such as HTTP. However, the documents and programs used in online learning are unique to online learning, and therefore, there is no existing set of mark-up languages and program interfaces for learning objects. Therefore, before we can use learning objects, we need to define each of these. The next few sections will discuss efforts to establish common standards for online learning specifically.

14.3.4 xii. A Common Mark-up language

Online learning uses XML. The common language adopted by online learning designers is - and being adopted by database programmers, librarians and designers around the world - is the eXtensibe Markup Language, or XML, developed by the World Wide Web Consortium.

XML represents documents according to their structure. XML is a means of representing documents according to their internal structure. Each element of the document structure is denoted with some standardized script, known as a 'tag'. For example, in a book, the chapters, chapter titles, and paragraphs would beach be denoted with individual tags. The collection of tags used in a document is known as the document's 'mark-up'.

XML tags. XML tags are very simple. The beginning of the document element - say, a chapter - is identified by the name of the element in angle brackets, like this: <chapter>. The end of the chapter is identified by a second tag, this time with a forward slash included. Like this: </chapter>. Any text between these two tags is a part of the chapter.

XML tags are nested. XML tags can be nested, which means that the content between two XML tags - between, say, <chapter> and </chapter> - can contain additional XML tags. For example, a chapter may have a chapter name and a block of text. Thus, the chapter would be defined as follows:

```
<chapter>
   <name>Chapter One: A Day in the Shire</name>
   <body>Once upon a time there was a land...</body>
</chapter>
```

In a similar manner, a course containing units, modules and exercises may also be represented in XML.

```
<course>
   <name>Introductory Psychology</name>
   <lesson>
      <name>Lesson One: Freud and his Ego</name>
      <body>Sigmund Freud, an Austrian psychologist...</body>
   </lesson>
</course>
```

XML tags are defined by schemas. In order for tags to be used by many developers, there must be some common understanding of what tags are allowed, what they may contain, and what they mean. Otherwise, what one developer understands to be a

course may by another developer be understood as being only a lesson. Essentially, there needs to be a dictionary of tags. Thus, designers using XML refer to documents called schemas to define the tags they are using.

XML tags can be used to describe things. Some types of documents cannot contain XML tags. An image file, created using the GIF format, for example, cannot contain XML tags. However, the image can be described using XML. This is useful because it allows designers represent information about the image which may not be stored in the image itself such as the photographer, the date the image was created, the copyright holder, and more. The XML file used to describe the image contains this information and then points to the location of the image on the internet. Any document – including XML documents – may be described in this way.

Metadata is data about a document. We use the term 'metadata' to denote data that is about a document as opposed to data that is a part of the document. Thus, for example, the metadata for the book titled 'Moby Dick' includes the book's title, its author and its date of publication. The document itself is the text that begins after the words, "Call me Ishmael..." Metadata is extremely useful. Metadata may be stored in one location while the document may be stored in another location, thus allowing for centralized directories of objects distributed all over the internet. And metadata can be used to describe more than just documents: it can be used to describe people, computer programs, classrooms, AV equipment, and indeed, anything that can be described.

14.3.5 xiii. Common Program Interfaces

Program interfaces allow programs to work together. Suppose a designer is creating an online course and wants to use an innovative interactive quiz provided by an educational publisher. Using the course design environment, the designer would click and drag the quiz into the course, just as she would any other learning object. In order for the quiz to work properly with the course, however, it must be able to communicate with other elements of the course. For example, it must be able to find out which student is taking the quiz and to be able to report the results of the quiz back to the course. What the quiz and the course need is a program interface: a way of speaking to each other.

Three types of program interfaces: Launch, API and Data Model. There are three major types of program interfaces: first, a means of telling when the program has started and stopped; second, a way to communicate directly with the other program; and third, a way of defining the type of data that can be exchanged between the programs. These are called the Launch, API (Application Program Interface), and Data Model respectively. The Launch reports when the object has started and stopped. The API reports on the internal state of the program (reporting errors, for example). And the Data Model is used to track the status of the content of the learning object.

Program interfaces are vendor-specific. There are many ways for a learning object to communicate with a course. Learning objects may be written in a variety of computer languages, such as C++ or Java, or they may be program specific files, such as Flash animations. Depending on how the learning object was created, it may interact with the course using Java applets, Microsoft's active server pages and Common Object Model (COM) components, CGI programs such as perl, or Common Object Request Broker (CORBA) components. Commonly, the program interface is accessed through JavaScript calls that ensure that the content is "wrapped" with the means to establish communications with the course once it begins to execute.

14.3.6 xiv. Standards and Standards Based Initiatives

Standards initiatives define schemas and program interfaces. In order to enable the sharing of learning objects, a variety of standards initiatives have been undertaken. These standards initiatives have two major tasks. The first is to define the names of the XML tags, their allowable values, and their meanings in online learning. And the second is to define the launch, API and data models used by program interfaces. Not all standards initiative define all of these elements, and some standards initiatives 'piggyback' on others, extending or more clearly defining elements of the standards in question.

The purpose of standards initiatives is to enable interoperability. Initiatives, such as the IMS Consortium (see below), are intended to promote the widespread adoption of specifications that will allow distributed learning environments and content from multiple authors to work together (in technical parlance, "interoperate").

By "distributed learning environments and content", the standards authors mean different sets of learning materials, authored in different programming languages using different programs and located on different computers around the world.

Interoperability amounts to programs being able to interact. This is an elusive goal. It amounts to enabling content produced using Blackboard and stored on a computer in Istanbul – an interactive atlas, say – to be used in a course authored in WebCT and located in Long Island, New York. And by 'used' what is meant in this context is that the two elements – the atlas and the course – could interact with each other; the atlas, for example, might report to the course how long a give student spend studying cloud formations, and the course might instruct the atlas to display the appropriate university logo and links to discussion boards.

In order to interact, programs must use common definitions of objects. In order for this to work, the atlas in Turkey and the course in the United States must define similar objects in a similar manner. For example, both programs must understand what was meant by 'course', or 'institution', or even 'logo'. Thus there is a need to obtain a common definition of the objects and properties used by the two separate systems. Thus, the core of the IMS specification involves the definition of prototype objects (or more accurately, descriptions of prototype objects, since they would be defined differently using different computer languages). The IMS Enterprise Information Model, for example, defines a 'Person Data Object', a 'Group Data Object', and a 'Membership Data Object'.

The difference between standards and specifications. Not all standards initiatives are the same. Some initiatives, referred to as specifications, are intended to capture a rough consensus among practitioners. They are descriptive of current practice in the field, often incomplete, and often the work of an ad hoc consortium. (Lim 2001) By contrast, standards are regulatory principles formally endorsed by a standards body such as IEEE or the World Wide Web Consortium (W3C). Typically, we would think of specifications as leading to standards, but this is by no means always the case.

The difference between schemas and application profiles. As mentioned above, standards initiatives define, among other things, schemas. Schemas – also sometimes called namespaces (Bray, et.al., 1999) – define the XML tags that can be used in a given type of XML document. In practice, however, any given enterprise is likely to use a number of schemas for specific purposes. For example, an e-learning company may use both an e-learning schema and an e-commerce schema to

create its own customized schema. This customized schema is called an application profile (Heery and Patel, 2000). Application profiles may draw on one or more existing schemas, but cannot introduce new data elements. They may specify permitted values and they can refine standard definitions.

14.3.7 xv. E-Learning Standards Based Initiatives

Following is a list of some major e-learning standards based initiatives, sorted by category. It should be noted that there is some room for interpretation as to whether an initiative is, say, a specification or an application profile.

Specifications

Dublin Core: The Dublin Core is a metadata element set (or schema) specification intended to facilitate the discovery of electronic resources. It lists common elements of electronic documents, such as the document title, author and publisher. The original workshop for the initiative was held in Dublin, Ohio in 1995. Hence the term "Dublin Core" in the name of the initiative. The Dublin Core forms the basis for many metadata schemas, including those used by librarians and those used in online learning. http://dublincore.org/

IMS. The Instructional Management System Global Learning Consortium, Inc. (IMS) is a consortium of educational institutions and private corporations developing and promoting open specifications for online learning activities. IMS standards include standards for locating and using educational content, tracking learner progress, reporting learner performance, and exchanging student records between administrative systems. The standards include content packaging specifications, learning content metadata, and some program interface specifications. http://www.imsproject.org/

ARIADNE. ARIADNE (Alliance of Remote Instructional Authoring and Distribution Networks for Europe) and ARIADNE II are European Union research and technology development projects in "Telematics for Education and Training." The projects focused on the development of tools and methodologies for producing, managing and reusing "computer-based pedagogical elements and telematics supported training curricula" (learning objects and computer assisted learning). Along with IMS, ARIADNE was one of the major contributors to the IEEE Learning Object Metadata Standard (see below). http://www.ariadne-eu.org/

Standards

IEEE P1484 . The Institute of Electrical and Electronics Engineers (IEEE) is an association of engineering and electronics professionals that set electronics and computing standards. IEEE's Learning Technology Standards Committee (LTSC) oversees a set of standards for learning objects that may be identified collectively as P1484. The most important of these is P1484.12 (Learning Objects Metadata, known generally as IEEE-LOM). Other standards – such as those measuring student competencies – are incomplete or exist only as proposed projects. http://ltsc.ieee.org/

Application Profiles

AICC. The Aviation Industry CBT (Computer Based Training) Committee (AICC) has developed a set of nine Guidelines & Recommendations (AGR's) defining various aspects of computer based training. The AICC standard is important is provides a test for AGR-010, "web-based computer managed instruction". The certification test verifies that the learning object can communicate launch and program information to the course. http://www.aicc.org/

SCORM. Created by ADL (Advanced Distributed Learning), a partnership of commercial, educational and government organizations, the Sharable Content Object Reference Model (SCORM) is a set of standards intended for use by organizations providing learning to the U.S. military (and therefore, to military organizations in allied nations). ADL's SCORM draws on the IMS protocols and extends them, defining course metadata and program interface specifications. ADL also provides tests for interoperability, called 'PlugFests.' http://www.adlnet.org/

CanCore. A Canadian metadata standards initiative, CanCore is a "stripped down" version of the IMS standard. The CanCore specification takes a middle-ground approach between the minimalism of the 15-element Dublin Core Metadata Initiative (DCMI) and the massive 70-element IMS specification. The 34-element CanCore specification also provides a set of guidelines for the purpose and use of each element. http://www.cancore.ca

14.3.8 xvi. Classifications of Learning Objects

Classification helps in retrieval. Many of the issues related to learning objects address issues surrounding how to classify learning objects. The reason classification is important is that classification defines features that will be important when searching and retrieving the objects. For example, a Grade 3 teacher looking for an illustration of a frog to show for a few seconds to her class will not find it useful to retrieve a college-level week-long seminar on frogs and frog behavior. Classifying both the time of use and the level of difficulty would ensure that the teacher retrieves only appropriate objects.

Metadata is used to classify learning objects. The specifications, standards and application profiles described above are intended to aid in the classification of objects. The IEEE draft standard for learning object metadata (IEEE, 2001), for example, includes fields for catalogue entries, semantic density, aggregation level (or granularity), and more. These fields are intended for use by automatic filtering programs that aid in document searches.

Extending the Standard. The IEEE-LOM and similar standards should not be seen as limiting the classifications and descriptions available to metadata authors. IEEE is explicitly committed to the extensibility of the standard, which means that additional data elements may be added as needed though the classification and other fields. (Dublin Core Metadata Initiative,2000)

Extended classification using the classification field. Both the IMS specification and the IEEE standards also include a 'classification' field that is intended to provide "elaboration and extension of the meta-data." (Wason, 2000) The classification field provides descriptions of the learning object that might not be provided for in the rest of the metadata. Several such types of descriptions (or purposes) have been identified: discipline, idea, prerequisite, educational objective, accessibility restrictions, educational level, skill level and security level.

Classification values are determined by taxonomies. A taxonomy is an ordered classification of information using a controlled vocabulary of words and phrases. (Wason, 2000a) Taxonomies are typically defined by a library (or library association), such as the Library of Congress, or by a subject-specific professional association, such as the American Mathematics Metadata Task Force. To define the classification of a learning object, both the source of the taxonomy and the object's classification within the taxonomy must be stated.

Some common taxonomies (from Wason, 2000) are listed below. This partial list is

intended to provide examples of the range and origin of taxonomies available for use in learning objects.

LCSH: Library of Congress Subject Headings. Originally by the Library of Congress designed as a controlled vocabulary for representing the subject and form of the books and serials, it has become a tool for subject indexing of library catalogs in general, including a number of online bibliographic databases. http://www.tlcdelivers.com/tlc/crs/shed0014.htm

Yahoo. Although Yahoo is thought of as a search engine, it is comprised of an extensive set of classifications, though (as Wason points out), not necessarily in terms of sub-disciplines. http://www.yahoo.com

McRel. Content is organized in terms of educational standards for K-12 educational materials. At the top level, McRel uses a subject-based taxonomies consisting of 14 items. http://www.mcrel.org/standards-benchmarks/

Taxonomy of Educational Technology. This taxonomy describe a new way of classifying uses of educational technologies based on a four-part division suggested by John Dewey: inquiry, communication, construction, and expression. http://www.lis.uiuc.edu/%7Echip/pubs/taxonomy/index.html

2000 Mathematics Subject Classification. A complex classification of common topics in mathematics. http://www.ams.org/msc/

American Mathematics Metadata Task Force. Proposed subject classifications for school and college mathematics. http://mathmetadata.org/ammtf/taxonomies/

Medical Subject Headings. The National Library of Medicines's controlled vocabulary used for indexing articles, for cataloging books and other holdings. http://www.nlm.nih.gov/mesh/filelist.html

Other types of classification within the Standard. In addition to the classification field, other fields in the IEEE-LOM can also refer to external taxonomies. For example, the 'Learning.Resource.Type' field and point to an external classification of learning resource types, such as the Gateway to Educational Materials (GEM: See http://www.thegateway.org/) and specify a value, such as 'Curriculum', from that taxonomy. (Suthers, 2001)

14.4 C. The Practical Application of Learning Objects

A non-technical description of how individual learning objects are created and how courses are created from collections of learning objects

14.4.1 xvii. Creating Learning Object Content

We start by describing how individual learning objects are authored. While today most guides and references discuss online course authoring, since courses are composed of learning objects, the proper stating point is to discuss the authoring of individual learning objects. As defined above, learning objects may be any digital object, which means they may be any element – a map, a web page, an interactive application, an online video – that might be contained inside a course.

Two major components of learning objects: the content, and the wrapper. Learning objects consist of the learning object content, which is the actual instructional material in the learning object, and what may be called the wrapper, which consists of metadata describing the object and standards compliant program interfaces. We might think of authoring learning objects as akin to authoring pieces of a puzzle, in which case the content is the image or picture on the

surface of the piece, while the wrapper is the shape of the piece itself which allows it to fit snugly with the other pieces.

HTML is not a suitable medium for content because it is not portable. Today by far the most common medium for content is hypertext mark-up language (HTML). The problem with these HTML pages is that they're not portable. A web page designed for one course at one university will contain course and university specific information: the name of the course, the name of the university, and even a colour scheme. To be used or adapted by another course, the pages need to be redesigned. Moreover, HTML pages do not display well in multiple formats. A separate version must be created if, say, the page needs to be delivered over wireless access protocol (WAP) or if it is input as data for analysis by a Javascript or CGI process.

Portability requires separation of content and presentation information. In order to be portable, a document's content must be, first, structured, and second, separated from presentation information. A significant step in the right direction is to create course materials not in HTML, but rather, in a structured mark-up language such as XML. This goal is accomplished by XML, which uses tags to structure information and which refers presentation information to a separate document entirely, an XSL file. An XML file with identify the chapter titles in a document; an XSL file will say that chapter titles should be printed in 24pt red Helvetica text.

Content is authored using content-specific XML authoring tools. Rather than use a single tool, such as a generic XML editor, content authors use tools designed for specific purposes. For example, an author will use one tool to author a journal article, another tool to author a multiple-choice quiz, and another tool to author a simulation. These tools will generate XML output specific to the type of document being authored. Different versions of XML are used for different parts of courses, such as Math Mark-Up Language (MML), Scalable Vector Graphics Language (SVG) and the Synchronized Multimedia Integration Language (SMIL).

Hot Potatoes. Already, we have seen some content-specific authoring tools, one of the most popular being Half-Baked Software's Hot Potatoes, a tool for designing online quizzes. Authors using Hot Potatoes select the type of interactive exercise they wish to design, then fill in fields in a form with things like the question text, answer options, feedback content, and other elements of a quiz or exercise. The output code is generated automatically by the software as a stand-along piece of learning content ready to be inserted into an online course.
http://web.uvic.ca/hrd/halfbaked/

Virtual Cell Modeling and Simulation. The Virtual Cell program is used for modeling and simulating cell processes. Users can access biological and mathematical frameworks are encompassed within a single graphical interface, run the simulation, and output the results. Simulations can be exported in XML format.
http://www.nrcam.uchc.edu/

Many vendors offer suites of content authoring tools. It is unusual to see content development tools dedicated to a specific form of content. More common are content authoring tool suites, such as offered by Macromedia, Dazzler, Trainersoft or x.hlp. Authors working with a content authoring suite first select the type of content they wish to author, then enter content information in the space provided. Authors can also import elements, such as graphics or multimedia, created using other tools. Content authoring suites typically also offer means of linking different content components together so that an author could, for example, write

a lesson containing some text, an animation and an online quiz. http://www.e-learningcentre.co.uk/eclipse/toolbox/authoring.htm

14.4.2 xviii. Creating Learning Content

Learning objects meet instructional objectives. While some writers (such as Downes, 2001) depict learning objects as instructionally neutral, other writers (such as Friesen, 2000) argue that learning objects must contain a pedagogical component, that is, they must be directed toward some instructional objective. The purpose of a learning object, on this view, is to teach something.

Creating learning objects is similar to creating instructional content. Though the exact process varies from institution to institution, the creation of learning objects follows a pattern similar to that used to create instructional content. Almost all learning object authors follow roughly the same sequence of steps, characterized by Cisco Systems (2000) as: design, develop, deliver, and evaluate.

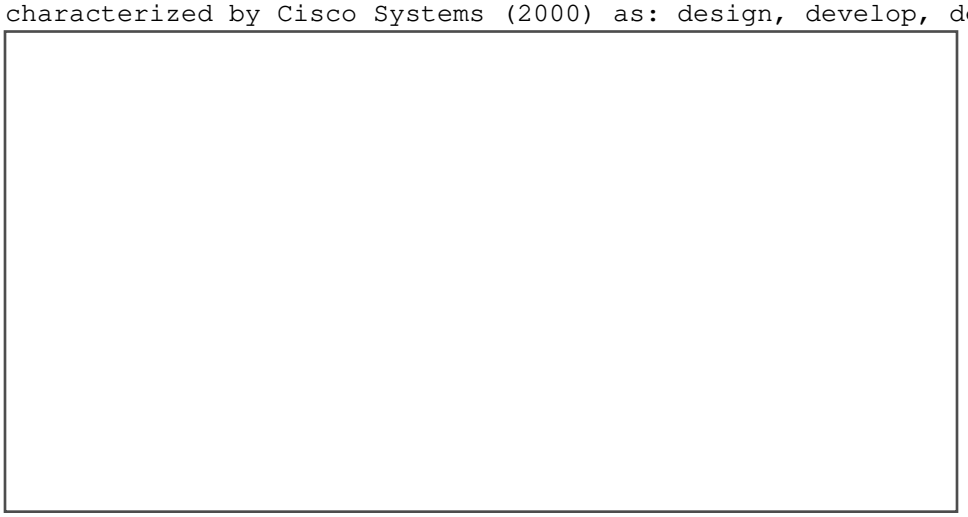

Design is based on learning needs. Though the method may vary from instance to instance, the design phase begins with an analysis of learning needs or desired learning outcomes. In a corporate setting this may be defined by a task analysis and an assessment of the knowledge an employee must have at each stage of the job. In a traditional instructional setting this may be determined by curriculum and educational standards bodies. The outcome of the design process is a list of concepts to be learned, each of which may be associated with various content elements, exercises and assessments (RIOs, in the diagram above).

Develop according to design specifications. The development of a learning object involves the writing, creating or assembling of educational content (which may include other, smaller, learning objects). The internal structure of learning objects varies, depending on the developers design model. Cisco (2000), for example, proposes that a learning object is composed of about seven information objects where each information object consists of content items, practice items and assessment items.

Deliver using multiple delivery media. Learning objects may be delivered as web-based courses, packaged as instructional materials on a CD-ROM, or employed by instructors as support materials for a traditional classroom lecture. With the advent of robust wireless internet access, delivery in remote locations will be enabled. And with the advent of web services, delivery through other applications (such as, say, a customer relations management (CRM) tool) will be possible.

Assess at multiple levels. Most institutions use some variation of Kirkpatrick’s

1959 description of four levels of assessment (Kirkpatrick, 1998). The first level, survey, measures whether the learners liked the instruction. Assessment, the second level, determines whether the learner met the learning outcomes defined in the design stage. The third level, transfer, examines whether the learner uses the new skills on the job or in practice. Finally, the fourth level measures the impact of the learning in a wider context. Companies with a strong eye to the bottom line may consider employing a fifth level , one which explicitly measures the return on investment (ROI). (Setaro, 2001)

14.4.3 xix. Creating Metadata and Wrappers

HTML defines metadata with meta tags. Most HTML authors are familiar with an earlier version of the wrapper, the meta tag. Located in the head of an HTML page, a meta tag is used to provide information about the document. Commonly used meta tags include the keyword tag and the description tag. While once commonly used by search engines, the abuse of meta tags by unscrupulous web site owners seeking to increase traffic has rendered them generally unusable.

Metadata is usually generated automatically by the content authoring tool. Most content authors will not author metadata directly. Rather, content authors using a tool such as Macromedia's Dreamweaver can generate metadata automatically using the metadata editor command. http://www.eduworks.com/LOTT/macromedia.htm

Metadata may also be stored in files external to the learning object. Although metadata is closely associated with the learning content, it may be stored in a separate file. For some types of files – such as a .gif image – this is a necessity, since .gif images cannot contain XML tags. And while other documents – such as XML pages or MS-Word document – may contain metadata expressed in XML, for reasons of size it may be more convenient to store the metadata separately from the document itself.

Metadata for legacy content is generated using conversion tools. A lot of educational content exists as plain HTML, MS-Word or databases. These were created before metadata assumed an important role in online learning. To convert this content into sharable learning objects, metadata must be created. While metadata could be created by hand, it is more efficient to use a metadata tool that reads legacy documents and automatically generates metadata. (Haven't been able to find one so far, though)

Learning objects require wrappers. To be compliant with all learning object standards, learning objects also require wrappers. A wrapper is a set of code, usually written in JavaScript, that tells the course where to find information about the learning object's program interfaces. For most content, the program interfaces are very simple. For example, if the content consists only of a document, the course needs to know only that the document has been opened or that the document has been closed. These can be handled by simple and standard Java applets, and hence, a very simple JavaScript pointer to these applets can be employed as a wrapper. http://www.eduworks.com/LOTT/Mm-LOTT/SCORM_RTI_help.htm

Wrappers can be generated automatically by the authoring system. Most content authors will not create their own wrappers. Though relatively new, wrapper generation tools are becoming available with educational content creation software. For example, Macromedia has just (December, 2001) released an extension for DreamWeaver that automatically creates SCORM 1.2 compliant wrappers. There's no good URL. Go to http://www.macromedia.com/resources/elearning/ and search for 'SCORM wrapper'

14.4.4 xix. Creating Packages

Courses are packages of learning objects. Learning objects created using content authoring tools are typically small, consisting of no more than the equivalent of an hour or two of instructional time (there is some debate as to how small a learning object may be and whether educational content must contain pedagogical features, such as a statement of learning objectives, in order to qualify as a learning object). Most educational institutions deliver larger chunks of instruction, called courses. To create a course, therefore, a set of learning objects must be assembled into a package.

Packages organize learning objects sequentially. In order to create a course out of, say, a dozen lessons, where each lesson is a separate learning object, a course author arranges these lessons into a sequence (it is worth noting that sequences are not defined within individual learning objects; they are defined by course authors). In some cases, where the learning objects are smaller units, course designers may need to create lessons composed of a sequence of individual modules, then the course as a whole out of the sequence of lessons. However created, the sequence of objects is used to define course-specific entities as the course outline or table of contents.

Learning objects are nested. When a collection of learning objects is organized into a sequence, the resulting product is also a learning object. Hence, a lesson composed of individual modules may be a learning object and would need its own metadata and wrapper. A course composed of lessons is similarly a learning object with its own metadata and wrapper. Each of these objects is created and stored in a database. The contents of this database are available to course authors. Some databases may be available over the internet, while other databases will be available only internally.

Packages are defined using manifests. The package is described in XML. IMS defines a specific set of XML tags used to describe a manifest. The manifest is like the shipping label for the package, detailing the contents of the package. The table of contents is an ordered representation of the titles of each item. The metadata for learning contents themselves may be actually contained in the package, or pointed to by a line in the page. Similarly, resources themselves may be contained in the package, or pointed to by a line in the package (obviously, non-textual resources, such as images, must be pointed to).

Packages are created using a Learning Content Management System. While a course author could locate and assemble learning objects by hand, it would be tedious and unproductive to do so. Courses created using learning objects are typically created using a development environment called a Learning Content Management System (LCMS). The LCMS performs two major functions: it provides authors with a means of locating learning objects, and it assembles them into standards compliant learning packages (or courses).
http://www.learningcircuits.org/2001/aug2001/ttools.html

Essential components of an LCMS. Though many types of LCMS are available, the typical LCMS will contain four essential features: an authoring application similar to the computer assisted software environment (CASE) described above, a collection of learning objects (called a repository), a means of sending the completed course to a delivery system (called a delivery interface), and administration tools.
http://www.internettime.com/itimegroup/lcms/IDCLCMSWhitePaper.pdf http://www.intern
ettime.com/itimegroup/lcms/

Using an LCMS. Using an LCMS, a course author defines major features of the course: its topic area, say, or its grade level. The author then instructs the LCMS to search through the learning object repository for relevant resources (because the data is in XML, the search can be very precise). From the search results, the author may review a learning object or select it for inclusion in the course. The LCMS retrieves the object metadata from the repository and inserts it into the course package. The LCMS automatically adds institution-specific formatting (such as the preferred colours, typestyle or wordmark) and prepares the package for delivery.

The LCMS is essentially a CASE. The learning content management system contains the essential features of the computer assisted software engineering environment described above. Learning objects, as defined by the various standards bodies, function in the same way as software objects (in an important sense, they are software objects). Thus, the use of learning objects within an LCMS allows course designers to emulate software engineers using rapid application design. The time and effort required to create learning materials is reduced substantially.

An LCMS usually also contains content authoring tools. The LCMSs offered by commercial vendors today do much more than merely assemble learning objects into courses. They are also bundled with one or more content authoring tools or may be used in conjunction with an authoring tool such as Dreamweaver. Thus, working within a single development environment a content author can create a repository of learning objects and then assemble those learning objects into complete courses.

14.4.5 xx. Delivering Courses

Courses are packages delivered by a learning management system (LMS). Once assembled, courses are delivered to a learning management system (LMS). The purpose of a learning management system is to provide a student direct access to course materials. Learning management systems typically restrict access to registered students, thus, courses offered via an LMS is usually locked behind some sort of password or authentication process.

An LMS integrates educational activities. In addition to providing access to course content, an LMS will provide access to a number of learning support tools. The most common of these include discussion areas and message boards, chat rooms, online quizzes and tests. The LMS provides a course navigation structure for the student, including a table of contents (or lesson plan), links to discussion areas, and access to repositories of files and supplementary materials.

An LMS provides student tracking and gradebook services. The LMS is frequently preferred by instructors because it automates many administrative functions. The LMS will track student progress through the course and will collect and display a students grades through a gradebook facility. The LMS manages the student's registration in the course and sometimes interacts with other campus services, such as PeopleSoft, to coordinate student registration and grade recording.

WebCT. One of the most popular learning management systems, WebCT (or Web Course Tools) is used by thousands of institutions worldwide. WebCT was popularized by its low price (though this is changing) and its practice of negotiating province or state-wide licenses covering dozens of institutions. http://www.webct.com

Blackboard. Blackboard is similar in form and function to WebCT and create a niche in the educational market by offering free course hosting as well as LMS services. This allowed educational institutions, especially smaller ones, to offer courses

without having to set up their own web server. Blackboard it is also available for purchase. http://www.blackboard.com

Hundreds of LMSs available. There are hundreds of LMSs available.

14.4.6 xxi. Course Format and Display

Courses need to be localized. When a course is delivered in a particular institution, such as at university or within a corporate training environment, the institution frequently desires to customize course content. At the simplest level, this may involve ensuring that the course displays the institution logo. At the most complex level, an institution may wish to customize a wide range of aspects of the course. Altering the course at the point of delivery is known as localization.

14.4.7 xxii. Pedagogical Issues

Learning Objects are used in an educational context. The delivery of instructional material seldom occurs in isolation. Courses are built, for example, on learning objectives or competencies. A single bit of learning forms a part of a larger picture. This creates a problem in the use of learning objects: how to arrange or sequence learning objects in such a way as to meet pedagogical objectives.

Context has multiple dimensions. It may be tempting to think of context merely as the location of a learning object in a sequence of learning objects. However the context of use involves numerous factors. In addition to content sequencing, these include the domain of discipline in which learning occurs (the difference, for example, between Engineers studying logic and English majors studying logic), cultural environment, organizational goals, and individual learning preferences or learning styles.

One object, multiple contexts. The advantage of using learning objects is that the same piece of learning material may be used in multiple contexts. Context selection may even be viewed as "the second path for personalization of objects (after adaptive selection of appropriate objects based on individual needs)." (Longmire, 2000) Several approaches to contextualizing learning objects are available to course and content developers (the list and terminology is from Longmire, 2000).

Tailored wrappers. As described above, wrappers define how a learning object interacts with a learning management system. One object can have multiple wrappers. Each wrapper then provides a different way to place the object into a specific learning context. While authoring a course using learning objects, an instructional designer might create custom wrappers to be used when the learner accesses that part of the course where the object would be used.

Tailored context frames. Think of a context frame as a generic wrapper describing a particular educational situation. The generic wrapper defines preferences for, say, learning styles, language, or educational background. Or a generic wrapper is created for a particular pedagogical model: one type of wrapper is used for a constructivist approach, another for a problem-based approach. The generic wrapper is applied to all objects being used in a particular course and can, as Longmire (2001) suggests, personal them with such techniques as humor, visual or linguistic themes, or explanations that relate it to a specific body of knowledge.

Adding context links to objects. This is like shaping learning objects so that they interlock like pieces in a jigsaw puzzle. Links are added to the learning objects to point to other learning objects (thus creating a sequence) or to

external context (thus linking instruction to, say, educational activities or quizzes).

Pattern templates. This is like placing learning objects in the correct location on a checkerboard. Templates based on instructional models and domain-specific sequences provide a data structure for the content of an online course. The selected learning objects are then deployed in the correct location through a process of matching the objects' properties (as defined by its metadata) to the requirements for each position in the course. Thus, for example, a learning object covering topic A2 and employing learning strategy G1 would fit into the A2-G1 location specified by the template.

14.4.8 xxiii. Beyond Courses

Learning objects enable custom learning. In the preceding few sections we looked at the use of learning objects to construct courses. Assuming that the promise of learning objects is achieved, and that course construction therefore becomes a relatively simple process, it becomes possible to envision the creation of a unique course for each learner. Learning objects, therefore, hold the promise of the creation of custom courses.

Toward learner-centered learning. Many advocates of online learning are proposing that new learning technology be used with a learner-centered approach to learning. This model replaces what could be called instructor-centered learning. Instead of having each learner dependent on an instructor for the arrangement and provision of learning resources, a learner-centered approach depicts the learner as selecting from a set of available resources.

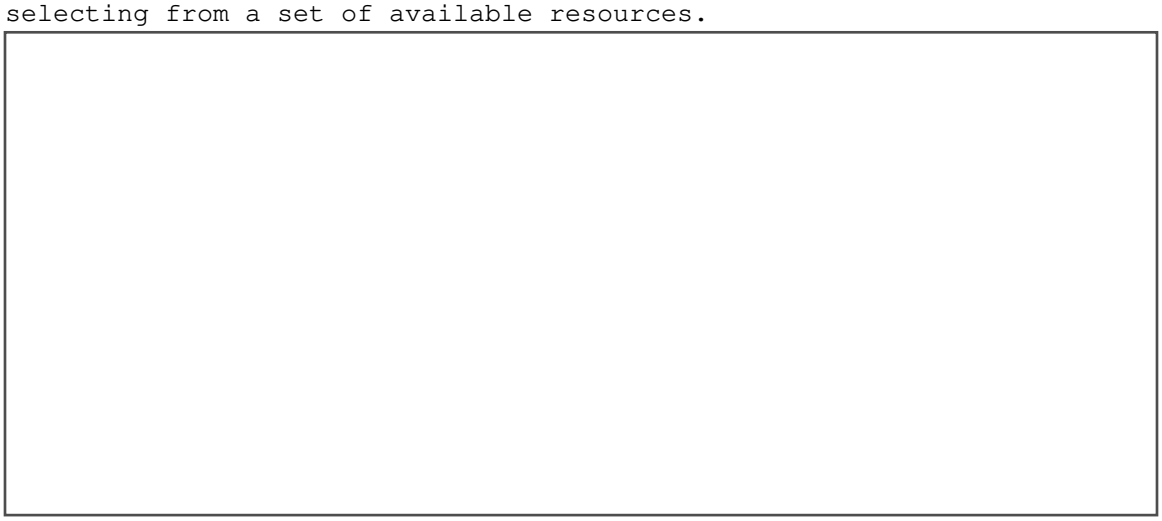

(Cisco Systems 2001)

Example: the Performance Support Portal. For example, instead of being given a week of orientation, a new employee may be given a short course introducing the performance support portal. "Here he can view performance support tools, read frequently asked questions, chat with others in his position and in the company, look up important information, and ask questions of cyber and human experts. His PSP also tracks his training progress. When his daily training time comes around, his PSP suggests skills he needs to learn. Based on his preferences, the short training bits from which he picks are in his preferred media - movies, animations or text. The trainings he sees are customized for his specific job, based on his knowledge and past training results." (Schatz, 2000)

14.5 D. The Learning Object Economy

This concluding section is intended to show how the development and distribution of learning objects by multiple institutions forms a larger system or educational resources worldwide

14.5.1 xxiv. Course Portals

Course Portals are a mechanism to help students select courses. A course portal is a website offered wither by a consortium of educational institutions or a private company working with educational partners that lists courses from a number of institutions. The purpose of a course portal is to enable a student to browse through or search course listings to simplify the student's selection of an online course.

TeleEducation. A New Brunswick, Canada, learning organization, TeleEducation NB hosts the TeleCampus Online Course Directory. Courses are submitted by institutions and screened to ensure that they are fully online. The database contains more than 50,000 courses, including about 3,000 free courses and 1,200 complete and fully online programs. TeleCampus provides a subject-based directory and search services. http://teleeducation.nb.ca/

UNext. Focussing on business education, UNext collaborates with major business schools such as the Columbia Business School, Stanford University and the London School of economics to provide courses in leadership and management, e-commerce, marketing, finance, accounting, and business communications through the private and for-profit institution, Cardean University. http://www.unext.com

Hungry Minds. Hungry minds offers more than 17,000 courses through its online campus, Hungry Minds University, from course providers such as the University of California at Berkeley, the University of California at Los Angeles and New York University. Hungry Minds also provides learning content through publishers such as For Dummies, CliffsNotes, and Frommer's. http://www.hungryminds.com/

Fathom. Created by Columbia University and including partners such as the University of Chicago, the London School of Economics and Political Science, Cambridge University Press, The British Library, The Smithsonian Institution's National Museum of Natural History, and The New York Public Library, Fathom is a centralized for-profit learning object repository. While Fathom provides lectures, interviews, articles, performances and exhibits, its major focus is an offering of online courses from member institutions. http://www.fathom.com

14.5.2 xxv. Course Packs

Course packs are packages of learning materials collected to support a course. Offered primarily by educational publishers, course packs are collections of learning materials offered to instructors for use in traditional or online courses. The course pack may be pre-defined or custom built by the instructor. The instructor is expected to supplement the course pack with additional content, educational activities, testing and other classroom activities.

Course packs are stand-alone or offered through a learning management system. Some course packs, such as those offered by XamEdu, are stand-alone. This means that the course pack is distributed as a separate product and purchased by the student directly through the college or university bookstore. Supplementary educational materials are offered by the instructor on his or her course website or are delivered in a classroom setting. Course packs delivered through a learning management system are more like 'default' online courses. The instructor then

customizes the course for delivery online.

WebCT Course Packs. The learning management system WebCT offers course packs consisting of a course structure and set of readings offered by publishers with a distribution agreement with WebCT. Course packs are purchased by the institution on a seat-license basis and are then customized by the instructor.

14.5.3 xxvi. Learning Object Repositories

Learning object repositories are collections of learning objects or metadata. Learning objects are stored in databases called learning object repositories. There are two major types of repositories: those containing both the learning objects and learning object metadata, and those containing metadata only. In the latter case, the learning objects themselves are located at a remote location and the repository is used as a tool to locate learning objects. In the former, the repository may be used to both locate and deliver the learning object.

Repositories are either stand-alone or included in another service. Most well known learning object repositories are stand-alone. These repositories function a lot like portals in that they contain a web-based user interface, a search mechanism, and a category listing. Another major class of learning object repositories functions more like a database attached to another product. An LCMS, for example, may contain a learning object repository intended for its exclusive use.

Repositories may be centralized or distributed. Two major models for learning object repositories exist. The most common form is a centralized form in which the learning object metadata is located on a single server or website (the learning objects themselves may be located somewhere else). An alternative model is the distributed learning object, in which the learning object metadata is contained in a number of connected servers or websites. Distributed learning object repositories typically employ a peer-to-peer architecture to allow any number of servers or websites to communicate with each other.

Learning object repositories are in a state of transition. Many learning object repositories, and especially stand-alone repositories, are former online course portals (see above). These repositories are in a state of transition, listing and offering both courses and learning objects. Because of the changing terminology and increasing importance of learning objects, course portals will sometimes represent themselves as learning object repositories.

Merlot. Described above, Merlot is probably the most well known learning object repository. Merlot is a centralized repository containing metadata only and pointing to objects located at remote locations. It is stand-alone, acting like a portal for learning objects. In addition to providing search and categorization services, Merlot provides a peer review service provided by communities of experts in different subject areas. http://www.merlot.org

Campus Alberta Repository of Educational Objects. CAREO is a centralized collection of learning objects intended for educators in Alberta, Canada. A stand-alone repository, CAREO contains metadata and provides access to learning objects located on remote web servers. http://www.careo.org

Portals for Online Objects in Learning. POOL is a distributed (peer-to-peer) repository system under development intended to create a pan-Canadian repository of learning objects. A primary objective of POOL is to develop and distribute tools for creating connected learning object repositories. http://www.newmic.com/pool/ (not currently functioning). See also

http://www.canarie.ca/funding/learning/1999backgrounders/pool.html See also http://www.siliconalleydaily.com/issues/sar08132001.html

National SMETE Distributed Library. In development for the (SMETE), NDSL is intended as a "federation" of learning object repositories, each library using different document formats, different systems of classification, and different database and repository management schemes. NDSL is intended to join these libraries using a common search engine called Emerge and a method for sharing resources called LOVE (Learning Object Virtual Exchange). (Chen, 2001)

14.5.4 xxvii. Certification and Review

Types quality: content, pedagogy, and compliance. The quality of learning object may be assessed across a number of dimensions. Of these, three major categories may be identified: quality of content, in which the information provided by a learning object complies which what is generally accepted to be true by experts in the field; pedagogical soundness, in which the learning object conforms to recognized principles of teaching and learning; and compliance, in which the learning object may be shown to comply with metadata and learning object standards.

Content quality assessment requires review by experts. The traditional form of content quality assessment in the publishing industry is accomplished by peer review. As mentioned previously, this is the assessment tool employed by at least one learning object repository, Merlot. Learning objects provided by publishers, such as those contained in XanEdu, also undergo a peer review process as part of their having been published in a recognized journal or magazine.

There is little in the way of pedagogical evaluation available. Most learning object repositories offer little or nothing in the way of evaluating of pedagogical standards. Indeed, most such standards are offered informally as a set of recommendations or guidelines for best practice. Probably the only organization that offers an evaluation of pedagogical standards as part of a formal certification program is the American Society for Training & Development (ASTD). http://www.astd.org/ecertification/

Standards bodies test for compliance with learning object standards. Compliance with learning object standards involves determining whether learning objects are described with correct metadata and whether the wrappers provided with learning objects operate correctly with learning management systems. As mentioned above, standards compliance is tested by the Aviation Industry CBT (Computer Based Training) Committee and Advanced Distributed Learning.

Evaluation has developed as a proprietary service. As the examples above illustrate, the review and evaluation of learning objects has developed as a proprietary service offered by particular organizations or services. In some cases, such as with ASTD certification or ADL Plugfests, there is a substantial fee involved. In other cases, such as Merlot’s peer review, use of a particular learning object repository is required. In all cases, evaluators are selected by the service or organization; there is no such thing as a genuinely independent learning object review or evaluation service.

14.5.5 xxviii. Publishers and Private Enterprise

The web creates a problem of free content. Many publishers and educational institutions are concerned about the potential Napsterization of education. Napster was a program intended to allow internet users to share music files on the internet. These files were copied from CDs and distributed for free through an

indexing service provided by Napster. Though Napster was eventually shut down, a host of similar peer-to-peer music sharing services continued this function. Music publishers therefore faced substantial losses as a consequence of lost sales.

Educational materials could also be shared for free. because students are able to access, and therefore copy, course materials, there is the potential that educational materials could be shared for free over the internet in much the same way music files were shared. The concept of a Learnster has been proposed by a number of organizations, including the University of New Brunswick's Electronic Text Centre and TeleEducation NB. The Open Learning Agency, based in British Columbia, Canada, owns the domain names learnster.org and learnster.com. http://www.unb.ca/happening/releases/B989alliance.html

Access is prevented using passwords and proxy servers. To prevent the unauthorized distribution of learning materials, including journal articles, academic journals have turned to dedicated databases for the electronic distribution of publications. Access to these databases is restricted by password protection, and in addition, viewers must enter the database via a recognized proxy server. Thus, only individuals belonging to organizations that have purchases subscriptions may retrieve and view these publications. The copying and distribution of articles is limited because the identity of readers is known and readers become accountable to their employers for the use of these materials.

Copying is prevented using Digital Rights Management (DRM) technology. To prevent the sharing of educational materials on the world wide web, publishers are investigating digital rights management (DRM) technology. While there are many approaches to DRM technology, they all have in common the function of preventing copying through the use of dedicated readers. One approach requires that students purchase dedicated electronic text readers (sometimes known as e-book readers). Another approach requires that students view text through specialized software programs.

Adobe Content Server. Adobe markets a combination of a document server and reader designed to prevent documents from being copied. Encrypted documents are delivered from the server and can be read only by a version of Adobe Acrobat. The reader will display the document on authorized computers only. The reader may restrict viewing to particular sections of an e-book or restrict viewing to a particular time, such as the duration of an online course. http://www.europedrm.com/presentations/clarke.pdf Mike Clarke, Digital Rights Management for Documents, Digital Rights Management Seminar, 20 November 2001

Digital Rights Management Systems are protected by law. While digital rights management software, such as Adobe's e-book reader, is a tempting target for hackers, the United States has instituted legislation making it illegal to distribute software intended to decrypt protected documents. this legislation, the Digital Millennium Copyright Act, has been emulated in a number of countries around the world. A number of individuals and organizations around the world have been prosecuted under the act for distributing prohibited software. http://www.educause.edu/issues/dmca.html

Growing resistance in the academic community. The use of digital rights management and restricted access databases is meeting increasing resistance by academics who argue that the free exchange of scholarly content is being jeopardized by the new legislation and technology. A number of free online scholarship initiatives have started up with the intent to promote boycotts of subscription-based journals, exchange academic content for free, and lobby for

changes to the existing legislation.

Example: The Scholarly Publishing and Academic Resources Coalition (SPARC). The purpose of SPARC is to provide an alternative mechanism for the exchange of academic papers and research. SPARC encourages journals to provide free or low-cost editions online and lobbies academics to publish their papers in SPACE-endorsed journals only. http://www.arl.org/sparc

14.5.6 xxix. The Learning Marketplace

Development of the marketplace model for the collection and distribution of learning objects (see "The Learning Marketplace", above, for a more detailed discussion)

The online learning economy: a closed marketplace? As the discussion above illustrates, vendors of educational material often link the provision of content to the purchase of a particular service or product. Content providers, in turn, must enter into distribution agreements with those service or product providers in order to reach the market for their content. The result is what might be called a closed market. Access to that market is controlled by what might be called 'gatekeepers', in this case, vendors of learning content management systems and learning object repositories.

The closed marketplace excludes most providers of educational content. With the exception of a limited number of name universities and educational institutions, content provided through the major learning object repositories and learning content management systems is provided by major publishers and private corporations. Other colleges and universities, as well as small-scale content providers, are required to partner with a publisher or private corporation in order to gain access to these markets. Such publishers or private corporations obtain the rights to the educational materials and receive the majority of the income.

An open learning marketplace is required. To obtain access to the educational content market, colleges, universities and small educational publishers require access to an open learning content marketplace. Such a marketplace would consist of a distributed (or peer-to-peer) learning object repository. It would also create mechanisms for non-aligned third parties, such as educational organizations, certification agencies and professional associations, to provide independent assessment and review of online learning materials.

An open learning marketplace supports multiple standards. In an open learning marketplace, educational providers may employ their own standards for course metadata. This allows them to select between, say, IMS, SCORM or CanCore. Providers create their own metadata (perhaps using a metadata authoring tool) which is in turn retrieved from the provider's server (or harvested) by the learning object repository.

Example: Edutella. Edutella is a prototype peer-to-peer network under development that will exchange of educational resources between German universities (including Hannover, Braunschweig and Karlsruhe), Swedish universities (including Stockholm and Uppsala), Stanford University and others. It builds in a structured query service to help users locate materials, an annotation service to allow users to comments on materials, and a mediation service to join data from different metadata sources. http://edutella.jxta.org/ See also http://www.cetis.ac.uk/content/20010927163232/viewArticle

References

Bray, Tim, 1999. Dave Hollander and Andrew Layman. Namespaces in XML. World Wide Web Consortium. http://www.w3,org/TR/REC-xml-names

Chen, Si-Shing. 2001. NBDL (National Biology Digital Library)

. Slide Presentation. http://www.cni.org/tfms/2001a.spring/handout/Progress-SChen2001Stf-NSF.ppt

Cisco Systems, 2000. Reusable Learning Object Strategy. Version 3.1 http://www.cisco.com/warp/public/10/wwtraining/elearning/implement/rlo_strategy_v3-1.pdf

Cisco Systems. 2001. E-Learning. Slide Show. http://www.cisco.com/warp/public/779/largeent/avvid/downloads/elearning.pdf

Department of the Interior (DOI) Online. 2002. Online Training Program Information. http://www.doi.gov/training/online_features.html

Downes, Stephen. 2001. Learning Objects: Resources For Distance Education Worldwide. International Review of Research in Open and Distance Learning 2:1. http://www.irrodl.org/content/v2.1/downes.html

Dublin Core Metadata Initiative. 2000. Memorandum of Understanding between the Dublin Core Metadata Initiative and the IEEE Learning Technology Standards Committee. http://dublincore.org/documents/2000/12/06/dcmi-ieee-mou/

Friesen, Norm. 2001. What are Educational Objects? Interactive Learning Environments, 9 (3) 219-230. http://www.careo.org/documents/objects.html

Heery, Rachel and Manjula Patel. 2000. Application profiles: mixing and matching metadata schemas. Ariadne Issue 25. http://www.ariadne.ac.uk/issue25/app-profiles/

Human Resources Development Canada. 2002. Knowledge Matters: Skills and learning for Canadians. http://www.hrdc-drhc.gc.ca/sp-ps/sl-ca/doc/toc.shtml

Institute of Electrical and Electronics Engineers, Inc. 2002. Draft Standard for Learning Object Metadata (IEEE P1484.12/D6.1). http://ltsc.ieee.org/doc/wg12/LOM_WD6-1_without_tracking.htm

Kirkpatrick, Donald. 1998. Evaluating Training Programs: The Four Levels. 2nd edition. San Francisco: Berrett-Koehler Publishers.

Kurtis, Ron. 2001. Return on Investment (ROI) from E-Learning, CBT and WBT. School for Champions. http://www.school-for-champions.com/elearning/roi.htm

Kruse, Kevin. 2002. Measuring the True Cost of E-Learning. E-Learning Guru. http://www.e-learningguru.com/articles/art5_2.htm

Lim Kim Chew. 2001. Overview of the Singapore eLearning Framework. Presentation slides. http://www.itsc.org.sg/doc/2001-10/KinChew-SELF-Overview.pdf

Longmire, Warren. 2000. A Primer on Learning Objects. Learning Circuits. http://www.learningcircuits.org/mar2000/primer.html

Major, Lee Elliott. 2002. OU Attracting More Young People. The Guardian. http://education.guardian.co.uk/universityaccess/story/0,10670,637574,00.html

Setaro, John L. 2001. Many Happy Returns: Calculating E-Learning ROI. Learning Circuits. http://www.learningcircuits.org/2001/jun2001/elearn.html

Schatz, Steven (2000) Paradigm Shifts and Challenges for Instructional Designers. IMS Project. http://www.imsproject.org/feature/kb/knowledgebits.html

Singh, Harvi. 2000. Achieving Interoperability in e-Learning. Learning Circuits. http://www.learningcircuits.org/mar2000/singh.html

Suthers, Daniel D. 2001. Evaluating the Learning Object Metadata for K-12 Educational Resources. Proceedings of the IEEE International Conference on Advanced Learning Technologies (ICALT 2001). http://lilt.ics.hawaii.edu/lilt/papers/2001/suthers-icalt-2001-lom.pdf

Vail, Kathleen. 2001. Online Learning Grows Up. Electronic School. http://www.electronic-school.com/2001/09/0901f1.html

Wason, Tom. 2000. Dr. Tom's Classification Guide. IMS Global Learning Consortium, Inc. http://www.imsproject.org/drtomclassification.html

Wiley, David A. 1999. The Post-Lego Learning Object. Unpublished Draft. http://works.opencontent.org/docs/post-lego/

Wiley, David A. 2000. Connecting learning objects to instructional design theory:

A definition, a metaphor, and a taxonomy. In D. A. Wiley, ed., The Instructional Use of Learning Objects: Online Version. http://reusability.org/read/chapters/wiley.doc

If The Learning Object Economy represents the cumulation of my preceeding though, this essay marks an important beginning in my present thought. The idea – to which I find myself returning constantly – is that multimedia objects are like works in a new language, and that children today are becoming fleunt in this new language even as teachers try to get them to learn the old. Moreover, importantly, what we can learn about the structure of the network that distributes these objects can be learned by the semantics of this new language. The logic of the new literacy is the logic of the network.

15 The New Literacy

Written October 4, 2002 Published in Learning Place, September, 2002. USDLA Journal, October, 2002. (If that seems impossible, keep in mind that I was several weeks late submitting my work to Learning Place).

Time and again we hear from academics bemoaning the loss of the cultivated and literate student in today's schools, the victim, they say, of a multi-media diet of McDonalds, music videos and post-modernist pablum. Such students fail, moan the critics, to engage in complex dialogue and complex thought. They are capable of understanding only simple and sanitized text, and even then only when it is accompanied with moving pictures and a soundtrack.

I have spent a large part of my working life in the company of the literati, listening to their seminars, attending their lectures, reading their journalistic contributions to the pool of public knowledge. For me, the greatest invention of recent years has been the introduction of wireless networking so I can have something to do while waiting through the interminable gaps in their reasoned arguments. Even while reading, I prefer to have the radio or television playing to occupy my mind as I wade my way through the text. I am not alone (136k), as one exasperated instructor after another struggles to keep online clat to a minimum during class time.

Scollon (et.al.) calls this polyfocal attention: "Perhaps the most striking thing about our students' attention is that it is polyfocal. That is, very rarely do they direct their attention in a focal, concentrated way to any single text or medium. When they watch television, they also listen to music and read or carry on conversations; traveling on the bus or Mass Transit Railway they read and listen to music–most commonly they 'read' while chatting, watching television and listening to music on CD." (Scollon, R., Bhatia, V., Li, D. and Yung, V. 1999. Blurred genres and fuzzy identities in Hong Kong public discourse: Foundational ethnographic issues in the study of reading. Applied Linguistics 20(1):22-43) (135k).

Why don't students pay attention to only one think. Scollon (et.al) suggest that new technology may allow new distractions, but that people have always been polyfocal – but had to content themselves with things like smoking cigarettes or eating hot dogs. I think it's more than that. It seems to me that for an information age student the most definiing characteristic of written text is that it is slow. Not quite as slow as listening to voice mail messages, but when compared to the rapid-fire pace of information transfer most of us are used to, it is achingly slow. The words struggle to pass from one to the next, a disappointingly linear presentation of what would more usefuly be a multi-streamed layering and threading of information, context and content. Today's students see no reason to wait. If there is a lull in the information stream coming from one direction, they quickly shift focus to another.

The problem with text is that it can only do one thing at a time. As I compose this article, for example, I would like to combine the multimedia version of Lawrence Lessig's free_culture with the recent study showing that there is a generational gap, a gap so wide as to even include how the different groups use their thumbs. With hyperlinking, I can at least fit these disparate thoughts into a single paragraph. With text only, it would be hopeless.

And yet it is important, in order to make the point, that these phenomena be seen side by side, acquired, ideally, in the same moment by the mind, so that the nuances of the one can be understood by the other. To see the depth of the generational gap I want readers to visualize the use of the thumb on keypads (as compared to the awkward way adults navigate the touch-tone with their index fingers) and to place that alongside the impact the spoken word adds to the slides in Lessig's show, to present all of these as a single thought.

What the critics of new media are missing is what may be called hyper-grammar. Textual language is bound by rules of syntax and semantics, with reference and meaning tightly constrainted by systems of representation. It is not a thought, in text, if it cannot be articulated without a subject and a predicate. It is not related to another thought, in text, if it cannot be logically conjoined. Waves of meaning are washed aside when the experience is rendered into words. That experience, so quaintly called "filling in the gaps with your imagination" by the literati, is lamented by the older generation when it is lost. And frustrating for the young, who would like to know what the author really meant with just that turn of a phrase.

Today's reader works with a much wider grammar. Even such simply typographic conventions, such as the use of italics, bold and capitals, can add new meaning to a text. The addition of symbols, such as smileys, convey emotion or sentiment. The breaking of linguistic rules – like this – can add urgency or clarity. The dropping of nouns, verbs or pronouns can express coreference (essentially, placing two separate thoughts into a single context). True, the haste with which people type online can result in a myriad of interesting typos and other errors – but then the error rate in a message also designates its degree of formality (conversely – to remove the errors reduces all text to the same sterile state of formality).

This is but one dimension of the new literacy. Here is another: go to any online chat room or IRC and observe the conversation. To the initiated, what emerges is a slew of seemingly unrelated comments. The participants roam back and forth from one topic to the next, sometimes within a single post. When I have hosted chats online among academic, participants complained that it was too complex, that they couldn't follow the conversation (and would I please ask people to stop posting messages). It would probably astonish such people that younger users may operate in several such chats simultaneously, each one in a separate window.

What should be understood is that these multiple threads layer into one another. It's not merely that attention is being shifted from one to another stream of information (though that does sometimes occur). Rather, the different topic streams are each facets of a multilayered presentation. The best analogy is in the explicit use of a soundtrack to add meaning to a dialogue (a technique used by the pop news shows so popular on television – as Homer Simpson says to his wife, "Oooo, he must be evil. Don't you hear the scary music?"). Words and images and text fuse into a single, complex message. Just as I can now no longer separate John Stuart Mill from the Devonian gardens (where I read On Liberty) or Quine's

discussion of rabbits from the Edmonton river valley (where I read "Word and Object"), these multiple media add nuance to the text that words alone cannot convey.

So let us now return to the original complaint: that students are unable to understand complex concepts. If it is true that students use hyper-grammer, that their attention is polyfocal, and that their interactions are multi-threaded, then it seems that even short exchanges are quite complex. The difference is in how that complexity is expressed. And it is arguable - and I would argue - that the sort of complexity sought after by the literati is an inferior complexity than that experienced by the information age student.

How so? In a famous passage Michael Polanyi in his book Personal Knowledge defined 'tacit knowledge' as being similar to knowing how to ride a bicycle. His point was that, no matter how much we read about the subject, it would be impossible to learn until we actually mounted the vehicle and took a ride for ourselves. Now in a certain sense, learning to ride a bicycle through practice is much simpler than the corresponding textual description. Indeed, it is likely that the person who has learned to ride the bicycle could not even understand the textual account of the same process (particularly if he mathematics of balance and motion are included). And yet, the person riding the bicycle has the very same knowledge as the person who has grasped the text - more, even, according to Polanyi.

What information technology brings us is the capacity to substitute experience for description. At the most basic level, we immerse ourselves in the darkness of a movie theatre and see and feel for ourselves what it must have been like to be on board the sinking Titanic. But add to this the possibility of multiple channels of communication, immersive simulation, multi-threaded interaction - a veritable medly of sight, sound and text - and we are able to move ourselves much closer to the experience, and thus to acquire a complex (though non-textual) comprehension of the event.

Moreover, the teen-age student may be in no better a position to describe this knowledge than a six year old who can ride a bicycle. Perhaps the only textual account he can give is a half-gutteral "whoa." But this does not mean that the information has not been acquired. It merely means that the information has not been abstracted from its experiential surround, abstracted, stripped of emotion and rendered in neat little syntactically correct packages. Such a student would fail utterly in contemporary evaluations of learning (literary criticism being a foreign art form, an earlier and drier version of Siskel and Ebert). But this is more a criticism of the testing instrument: an evaluation of what the student really learned would be found in practice (does he avoid icebergs?) and creativity (can he emulate and improve upon the representation of ships being struck by icebergs?).

It may be years before people cease to lament the decline of the literate student (after all, people today still bemoan the fact that students no longer learn Latin and Greek). But lament it we should not, because by avoiding the need to codify knowledge into sentences and seminars students today are acquiring not only different modes of learning, but much more efficient and effective modes of memory and recall. The new literacy may not be an even greater grasp of the fine points of language, but rather, a capacity to move beyond the limits of text and to manipulate experience directly.

When I joined the National Research Council in the fall of 2001, one of my first projects was to ensure New Brunswick participation in the project that eventually became eduSource, an assignment I accepted enthusiastically because it dovetailed with my work in learning object distribution networks. As I began to advocate for a system of content distribution as described above, I found myself in need of a simple description of RSS in order to explain what I meant. This paper was thus written in advance of an eduSource Atlantic meeting for that very purpose.

16 An Introduction to RSS for Educational Designers

Written November 2, 2002. Unpublished.

RSS stands for "Rich Site Summary" and is a type of XML document used to share news headlines and other types of web content. Originally designed by Netscape to create content "channels" for its My Netscape pages, RSS has been adopted by news syndication services, weblogs, and other online information services.

Because it is one of the simplest uses of XML, RSS has become widely distributed. Content developers use RSS to create an XML description of their web site. The RSS file can include a logo, a site link, an input box, and multiple news items. Each news item consists of a URL, a title, and a summary.

Content developers make their RSS files available by placing them on their web server. In this way, RSS "aggregators" are able to read the RSS files and therefore to collect data about the website. These aggregators place the site information into a larger database and use this database to allow for structured searches of a large number of content providers.

Because the data is in XML, and not a display language like HTML, RSS information can be flowed into a large number of devices. In addition to being used to create news summary web pages, RSS can be fed into stand-alone news browsers or headline viewers, PDAs, cell phones, email ticklers and even voice updates.

The strength of RSS is its simplicity. It is exceptionally easy to syndicate website content using RSS. It is also very easy to use RSS headline feeds, either by viewing a news summary web page or by downloading one of many free headline viewers. Though most RSS feeds list web based resources, several feeds link to audio files, video files and other multimedia.

16.1 Why RSS is Important for Educational Designers

RSS is the first working example of an XML data network. As such, and in this world of learning objects and metadata files, RSS is the first working example of what such a network will look like for educational designers. Just as news resources are indexed and distributed in the RSS network, so also educational resources can be indexed and distributed in a similar learning object network.

The model provided by RSS is very different from the model provided today by learning content management systems (LCMSs). In the world of the LCMS, everything is contained in one very large software application. Insofar as content is distributed at all, it is distributed in bundled content libraries. This means that educational institutions must make a major investment in software and expertise in order to access learning content.

RSS, by contrast, is not centralized. It is distributed. Content is not distributed in bundles, it is distributed one item at a time. There is no central store, repository or library of RSS content; it is all over the internet. To access and use RSS content in a viewer or in a web page, you do not need a large software application. A simple RSS reader will do the trick.

For this reason, the distribution of educational content over the internet will look a lot more like an RSS network than it will an enterprise content management system. Many more people will use a distributed learning object network not only because it’s easier and cheaper, but because they can access much more content for much less money.

As a result, the concept of syndicated educational content can really come into play. While there will always be a need for reusable learning objects (RLOs), anything that can have an educational application – including images, videos, journal articles, even news items – can be distributed through a learning object syndication network.

16.2 The RSS Network Architecture

An RSS network consists of three major components:

A (large) number of content providers, each providing news articles, and each providing their own RSS files describing these news articles.

A (smaller) number of RSS aggregators that read these RSS files from multiple sources, collect them into an index, and provide customized "feeds" of topic-specific news headlines from this index.

A (large) number of news viewing applications that, based on user input, connect to an RSS aggregator, access a news feed, and display it to the reader. On viewing the news feed, the reader can then select a news item (by clicking on the headline) and read the article directly from the content provider.

The RSS network architecture looks like this:

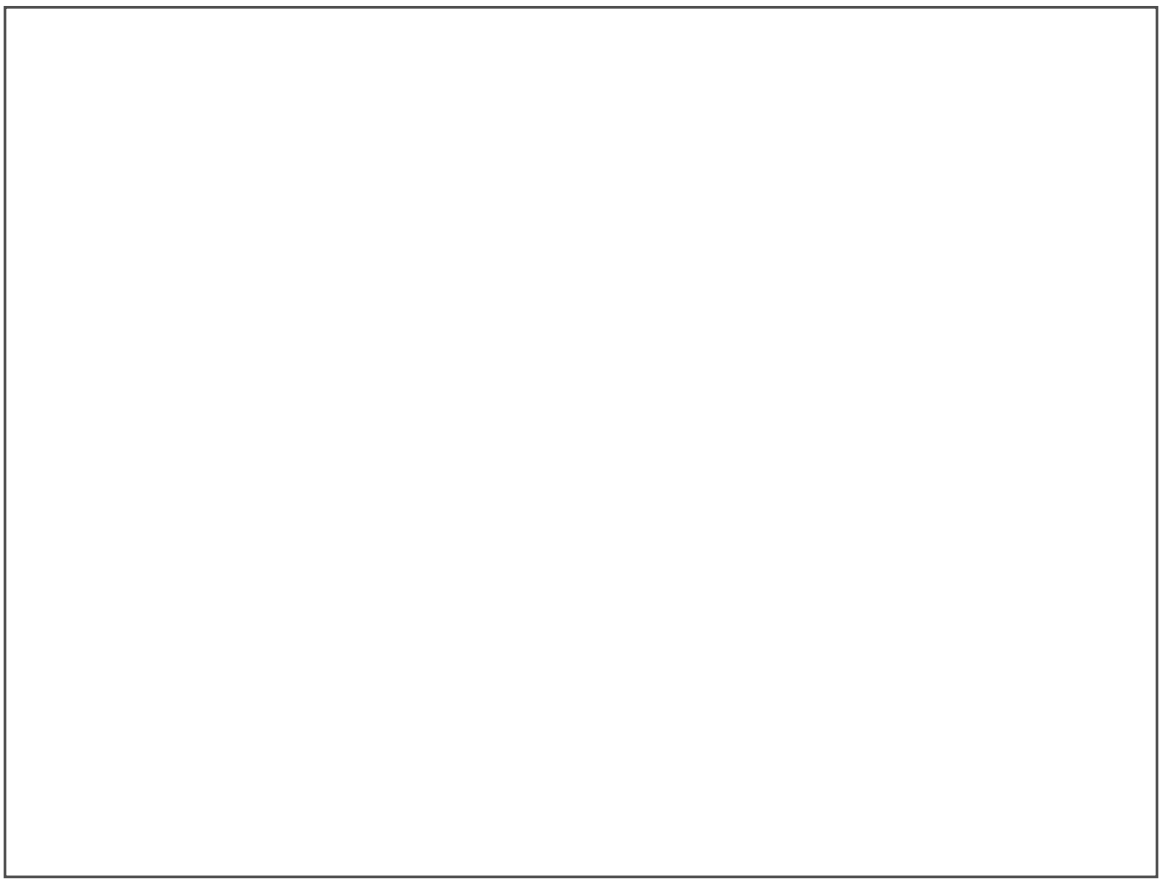

16.3 RSS Channels

A single RSS file is typically called an RSS channel. This is a lot like a television channel or a radio channel: it contains news items from a single source. For example, to the right is an HTML view of an RSS channel from the online magazine First Monday.

An RSS channel consists of two major sets of elements:

Channel Properties – the name of the channel (in this case, First Monday), a home URL for the channel, and an image for the channel.

Item Properties – the separate news items listed in the channel. In this case, there are ten news items listed. Each item has a headline and a URL. In some cases, an item will also contain a short summary, a publication date, author information, and more.

In order to define a channel like the one on the right, the channel properties and the item properties are defined in an XML file (or to be more precise, an RSS file), as follows:

At the top of the text box is a declaration of the type of XML file being used. Next we see an XML field describing the RSS channel. Within this field is the channel name, link and description. Finally, we see a list of the items available in the channel (I have only listed two items here). Each item is described with a title, and a URL.

16.4 Creating an RSS Channel

Because an RSS channel is an XML file, it can be created using a plain text editor – the same sort of editor that you might use to create an HTML page. It is usually easier to start with a template (such as the RSS file displayed on the previous page) and to insert your own values for each tag.

Typically, though, RSS files are created automatically. This is possible because an RSS file has a standard format. Thus, if you have a database of articles, then you can easily create an RSS channel from that database by extracting table data into XML data.

Another popular means of creating an RSS file is by means of scraping an HTML file. To scrape an HTML file is to extract link titles and URLs from an ordinary web page. This is done by analyzing the HTML tags and for the link title and URL. A script such as this in Perl

will generate a list of the URLs and titles in almost any HTML page. Thus it is very easy to write a script that will generate an RSS file from any web page.

There are online services, such as Moreover, that specialize in HTML scraping. Moreover scans the web pages of major newspapers from around the world and generates RSS channels for them. Moreover also provides a series of specialized RSS feeds.

16.5 Weblogs

Weblogs, or as they are sometimes called, blogs, have a unique role in the world of RSS. A weblog is, in the first instance, a web page that is updated on a regular basis. Thus a weblog resembles a diary or a journal; entries are dated and each day the weblog web page contains something new.

What distinguishes a weblog from a personal web page, though, is that the weblog consists of a series of entries associated with links to other resources on the web. Thus the typical weblog consists of a list of sites, descriptions of those sites, and some discussion.

My daily newsletter, OLDaily, pictured at right, is a typical example of a weblog.

OLDaily has channel elements, such as the newsletter title and home page URL.

The difference is in the items. I am not listing my own articles. I am listing articles published by someone else. The description, however, is mine. I am providing my description and interpretation of someone else’s material.

Also worth noting is that I did not obtain my items from a single source. As you can see by looking at the items, I have listed different articles by different authors working for different publications.

So a channel need not be produced by a content producer. A channel can be created by anybody with something to say about the items being described.

The RSS for OLDaily, though, looks exactly like the RSS created for First Monday.

If you were to look at the RSS for OLDaily, though, you would find several more tags, and specifically, tags to denote the author, publisher and publication date of the article, along with the URL and the title.

16.6 Aggregators

An RSS aggregator is a type of software that periodically reads sets of RSS files and indexes them for display or syndication. There are two major types of aggregator: centralized and personal.

A centralized aggregator is intended for use by a number of people. RSS files are read by the centralized aggregator and are then used to create a topic-specific web page or customized RSS feeds (as in the diagram above).

The Moreover aggregator, for example, culls RSS from a variety of sources (including HTML pages, which it scrapes). It then provides RSS feeds devoted to specific topics – such as Microsoft, as illustrated – that can be used on web pages.

Another well known centralized aggregator is a web site called News Is Free. At latest report, the aggregator collects headlines from 3744 sources and allows readers to browse the headlines or to search for the latest news. The site also offers headline syndication and web services.

A personal aggregator is an application that runs on a user's desktop. It can access a centralized aggregator (in which case it functions as a headline viewer) or, more frequently, it can access an RSS channel directly. This is called subscribing to the RSS channel.

Radio Userland, for example, accesses a list of channels from a centralized aggregator. The user selects a set of channels from this list and subscribes to them. Radio then updates item listings from the selected channels once an hour. Using the data supplied from the RSS files, it also facilitates the creation of a personalized weblog (which can in turn be published as another RSS channel).

Another popular personal aggregator is called Amphetadesk. Like Radio Userland, users can select from a list of channels supplied by a centralized aggregator. Amphetadesk users can also subscribe to a channel directly if the channel provider has provided a specialized subscription script.

Aaron Swartz has written a novel aggregator that converts RSS channels into email messages. Metadata Harvesting Generally

RSS aggregators are members of a wider class of software called harvesters. The purpose of a harvester is to retrieve and parse metadata located on remote servers, providing the information in a usable form for various applications.

In educational design, metadata harvesting involves the aggregation of metadata records associated with education and training resources. Aggregation provides greater exposure of those resources to the wider community. Aggregation also promotes the reuse of resources and encourages the development of interoperable resources.

The most well known metadata harvesting initiative in the education community is called the Open Archives Initiative (OAI). The purpose of the open archives initiative is to provide access to academic papers and other resources over the internet.

In the terminology employed by OAI, the content provider is called the data provider. The aggregator is called the service provider. And a harvester is a

client application used by the service provider in order to access metadata from the data provider.

A number of initiatives have emerged based on the OAI harvesting protocols. The Public Knowledge Project, for example, is an initiative based at the University of British Columbia intended to develop an Open Journal System (OJS). The OJS assists with every stage of the refereed publishing process, from submissions through to online publication and indexing.

Another project is the Illinois OAI Protocol Metadata Harvesting Project. The public face of this project resembles a centralized aggregator in that it provides a search window for academic articles. It then displays the full metadata record for the selected article.

The National SMETE Distributed Library (NSDL) is another organization looking at metadata harvesting. The model described by NSDL mirrors almost exactly the model used by the RSS community. The NSDL project is attempting to collect metadata not only from OAI compliant archives but also from a wider variety of metadata sources. This, reports the NSDL, does not cause a problem not in the collection process but does cause a problem in service delivery.

16.7 Headline Viewers

The purpose of a headline viewer is to provide a list of headlines obtained from an aggregator. When a user selects from this list of options (by clicking on a headline), the headline viewer retrieves the article from the source site and displays it for reading.

Many headline viewers exist. One of the earliest and best known is Carmen's Headline Viewer. This program runs as a stand-alone application on the desktop and taps into major centralized repositories such as My Userland, XMLTree, News Is Free and Grok Soup.

The major difference between a headline viewer and a personal aggregator (described above) is in the display of the results. Carmen's Headline Viewer, as can be seen from the screen shot at right, displays headlines sorted by topic. Thus the reader is not restricted to a small set of channel subscriptions; instead, they obtain topic-specific feeds.

Other headline viewers, such as Novobot, create web pages for viewing headlines. This has the advantage of being a familiar interface for most users. However, web pages listing many resources can take a while to load.

16.8 What RSS Does Not Provide

RSS is a powerful tool for content syndication. However, it lacks some important features needed to develop into a robust solution for the location and organization of educational content.

One of the major deficiencies (identified in the NSDL paper) is the lack of digital rights management for both original articles and metadata feeds. RSS assumes that all articles and metadata are published on the open web and are therefore freely available for viewing by anyone. This means that resources for which there is a fee cannot be accessed through the RSS network.

Another major problem, again identified in the NDSL report, is RSS's inability to deal with mixed metadata. Over the years various types of RSS have developed (RSS 0.9, RSS 1.0, RSS 2.0) and the tools have adapted on a case by case basis. RSS aggregators, however, still cannot access alternative forms of metadata, much less display resources from a wide array of both news and non-news sources.

A third problem for RSS is in the manner it handles weblogs. As described above, weblogs are commentaries on original resources. Yet they are displayed in the same format, and in the same place, as original articles. This can result in duplicate listings when the same resource is described in several weblogs. In addition, there is no means to display the comments from many weblogs side by side.

More Reading

Amphetadesk. http://www.disobey.com/amphetadesk/

Arms, William Y. et.al. October, 2002. A Case Study in MetadataHarvesting: the NSDL. http://www.cs.cornell.edu/lagoze/papers/Arms-et-al-LibraryHiTech.pdf

Carmen's Headline Viewer. http://www.headlineviewer.com/

Education Network Commonwealth of Australia. 2002. About Metadata Harvesting. http://www.edna.edu.au/harvesting/module2.html

Illinois OAI Protocol Metadata Harvesting Project. http://oai.grainger.uiuc.edu/search

Jackson, Dean. September 19, 2002. Aaron Swartz's RSS to Email Aggregator. http://www.w3.org/2002/09/rss2email/

Moreover. http://www.moreover.com

Novobot. http://www.proggle.com/novobot/pdef.shtml

News Is Free. http://www.newsisfree.com

Public Knowledge Project. University of British Columbia. http://www.pkp.ubc.ca/ojs/

Radio Userland. http://radio.userland.com/

Van de Sompel, Herbert, and Lagoze, Carl . June 14, 2002. The Open Archives Initiative Protocol for Metadata Harvesting. Version 2.0. Open Archives Initiative. http://www.openarchives.org/OAI/openarchivesprotocol.html

Winer, Dave. October 8, 2002. What is a News Aggregator. http://davenet.userland.com/2002/10/08/whatIsANewsAggregator

Though this paper was written before the next, it was written after (and in defense of) several of the pieces that comprise the next, leaving me with a chronological paradox. It is presented here in order to underscore some of the underlying thinking behind the next paper. The next paper is a detailed set of principles and architecture. This paper is the theory.

17 Componentization

Written on November 4, 2002. Published in Learning Place, December, 2002.

Stephen Lanahas challenges my conception of a component-based learning object repository network. He writes,

"In a sense the terms "component" and distributed are holdovers from client-server architectures, they don't properly describe what is being done on the internet (even though the internet is effect the ultimate client server structure, it has more or less transcended itself). The closest things to describing what we're looking for is "web services" and "semantic web" but these seem inadequate as well (or else why we would the IT community be compelled to go back and drag out the componentization paradigm again).

"We are in a state of transition wherein the previous vision of how this would work, didn't work and the new visions have not really got off the ground. On a more practical note, one of the greatest failings with the "distributed" approach was its over emphasis on application issues, leading to extreme complexity.

"These factors have been conquered before, but only in strictly controlled environments (where performance was guaranteed, regulations in place etc.)

"The e-learning repository network will have to be a lower cost proposition by nature in order to succeed. The reason that peer to peer music succeeded was its architectural simplicity - simpler by a factor of 10x or more than what we're looking at.

"It's a very tough call, I still think we are a step or two away from a viable conceptual architecture for distributed repositories, I don't think those steps are big - just haven't recognized the breakthoughs yet..."

I think he raises serious issues. But they are issues with the potential for a good response.

It would be useful to obtain an analysis of why previous efforts at componentization (is that a word?) failed, and indeed why (and in what way) the internet is growing beyond itself.

From my perspective, the internet - and web - is itself the clearest example of the success of componentization.

Part of my thinking is derived from my imagining what the web would have looked liked had we taken the enterprise system approach. On such a model, in order to have web access, you would have to buy a large content management system, download 'the web' (or those components of the web to which you have licensed access), store it on your own system, and access it in-house.

Obviously, the web could not have been successful without the development of some very simple components: a way of separating the browsing functionality from the web server functionality.

I agree that the cornerstone of componentization is simplicity. It is not clear to me that peer-to-peer as it has evolved has achieved that level of simplicity. The mess in instant messaging, for example, illustrates what can go wrong. That is why

I advocate standards tolerance, and that is why I advocate a decentralized system – so you cannot have an AOL, say, denying access to the network by competing technologies.

I believe that simplicity is obtained by taking a lot of what is (was previously) believed to be network (or enterprise system, on the other view) properties and incorporating them into the objects themselves. Build complexity into learning objects, not the system that transports them.

For example, most (all?) LMSs and LCMSs have built in discussion board (or chat) tools. There is no reason for this. Access to a discussion list is a service, the functionality (or front-end) for which can be encoded into a learning object (perhaps not on the current definition of 'learning object' but I refuse to be dragged into such digressions.

Building too much functionality into the network itself – rather than into the objects transported into the network – is what leads to the creation of alternative competing networks, and a fragmentation of the system. It would be like building streaming media into the definition of the HTTP protocol – then we would have a web that is compatable for MS media Player users and a completely different web for Real Networks users.

But by building functionality into a subsystem that can be transported over the network (aka plug-ins) this sort of fragmentation can be avoided. Oh sure, you still have compatibility issues (I have *never* made Quicktime work properly, for example), but at least the network as a whole isn't broken.

Semantic web, web service – yes. These are names for it, implementations of the concept. They haven't really taken off because we lack the simple components that allow people to use them productively. There is no 'web services browser' – nor would one make sense, under the current implementation. There is no 'semantic web' browser (Amaya notwithstanding).

That's why we need a simple tool for e-learning access, a simple API or protocol that can be easily adapted by developers, that provides people with a view of the entire learning object sphere. Until such a thing exists, the whole field of learning objects dies stillborn, a great idea that nobody could use.

But you don't get a simple browser without componentization. Which, since I think learning objects really are a good idea, is why we need componentization.

My own personal view (and of course there remains every possibility that I may be wrong) is that the breakthroughts we need are in this description:

multiple instances of essential components, to avoid bottlenecks

standards tolerance, to ensure interoperability

smart learning objects, to ensure simplicity

browsing tools, to ensure usability

The rest of it, the layers of third party metadata, of DRM, of personal information systems, etc., are to enable increased functionality. They are layered onto the basic system, much in the way images and later streaming media were layed onto the web.

Anyhow, that's my view. That's how I see the infrastructure unfolding. And I would be very surprised to see it unfold a substantially different way, at least over the long term.

Through the summer of 2002 I wrote a series of short papers (many of which ended up in Learning Place) in an effort to sway the eduSource project to my way of thinking. Then, in the fall of 2002, I was invited to speak in Milan, and so in desperate need of something to talk about I strung the papers together, added some bridging material, and produced what was probably my most important work of the year. It was not enthusiastically adopted by eduSource, as I had hoped, but it did receive some readership and acclaim. Oh, and it was very well received in Milan. The second part of the title only appeared in its published form; I would never have added the questiuon mark.

18 Design and Reusability of Learning Objects in an Academic Context: A New Economy of Education?

Written October 9, 2002. Published in USDLA Journal, January 7, 2003.

18.1 Introduction

The purpose of this paper is not to discuss the creation and use of learning objects per se but rather to look at systems for locating and distributing learning objects. What will be argued is that this system is currently poorly constructed, based essentially on what may be called a silo model of distribution. A series of problems and issues related to this model will be discussed. In place of the silo model, a distributed model of learning object repositories is proposed. This model is based on a set of principles intended to create an open and accessible marketplace for learning objects, in essence, a learning object economy. To conclude, a model for a distributed learning object repository network is proposed.

For readers unfamiliar with the concept of learning objects, the generally accepted definition is that learning objects "any entity, digital or non-digital, which can be used, re-used or referenced during technology supported learning." (IEEE, 2002) Wiley (2000) defines a learning object as "any digital resource that can be reused to support learning." Even so, as Wiley comments, "the definition is broad enough to include the estimated 15 terabytes of information available on the publicly accessible Internet." In this paper, a functional definition of learning objects is employed: a learning object is anything that is exchanged in what may be called the learning object economy.

18.2 The State of the Art

18.2.1 Overview

In this section common methods for locating and retrieving learning objects will be discussed. In particular, three major systems will be described: course portals, course packs, and learning object repositories. In addition, systems for collecting and organizing learning objects, learning management content systems, will also be described.

18.2.2 Course Portals

A course portal is a website offered wither by a consortium of educational institutions or a private company working with educational partners that lists courses from a number of institutions. The purpose of a course portal is to enable a student to browse through or search course listings to simplify the student’s selection of an online course. The following are examples of course portals.

TeleEducation. A New Brunswick, Canada, learning organization, TeleEducation NB hosts the TeleCampus Online Course Directory. Courses are submitted by institutions and screened to ensure that they are fully online. The database

contains more than 50,000 courses, including about 3,000 free courses and 1,200 complete and fully online programs. TeleCampus provides a subject-based directory and search services. http://teleeducation.nb.ca/

UNext. Focusing on business education, UNext collaborates with major business schools such as the Columbia Business School, Stanford University and the London School of economics to provide courses in leadership and management, e-commerce, marketing, finance, accounting, and business communications through the private and for-profit institution, Cardean University. http://www.unext.com

Hungry Minds. Hungry minds offers more than 17,000 courses through its online campus, Hungry Minds University, from course providers such as the University of California at Berkeley, the University of California at Los Angeles and New York University. Hungry Minds also provides learning content through publishers such as For Dummies, CliffsNotes, and Frommer's. http://www.hungryminds.com/

Fathom. Created by Columbia University and including partners such as the University of Chicago, the London School of Economics and Political Science, Cambridge University Press, The British Library, The Smithsonian Institution's National Museum of Natural History, and The New York Public Library, Fathom is a centralized for-profit learning object repository. While Fathom provides lectures, interviews, articles, performances and exhibits, its major focus is an offering of online courses from member institutions. http://www.fathom.com (You, 2001)

18.2.3 Course Packs

Course packs are packages of learning materials collected to support a course. Offered primarily by educational publishers, course packs are collections of learning materials offered to instructors for use in traditional or online courses. The course pack may be pre-defined or custom built by the instructor. The instructor is expected to supplement the course pack with additional content, educational activities, testing and other classroom activities.

Some course packs, such as those offered by XamEdu, are stand-alone. This means that the course pack is distributed as a separate product and purchased by the student directly through the college or university bookstore. Supplementary educational materials are offered by the instructor on his or her course website or are delivered in a classroom setting. Other course packs are available for use only in a learning management system (LMS). Course packs delivered through a learning management system are more like 'default' online courses. Using tools provided in the LMS, the instructor selects the course and customizes it for delivery online.

The following are examples of course pack providers:

WebCT Course Packs. The learning management system WebCT offers course packs consisting of a course structure and set of readings offered by publishers with a distribution agreement with WebCT. Course packs are purchased by the institution on a seat-license basis and are then customized by the instructor.

Canada's SchoolNet. In Canada, the leading learning resources portal is probably Canada's SchoolNet. A list of resources is displayed, each with a short description and a link to an external website. SchoolNet also provides information about each site and provides an "advanced search" using metadata. Each resource in the "curriculum" area is approved by a professional "pagemaster". For the most part, however, SchoolNet links to institutional home pages, and not to learning resources per se. Teachers using the SchoolNet service must still search through these sites in order to locate suitable materials. http://www.schoolnet.ca/

MarcoPolo. MarcoPolo is a compilation of teaching resources from six educational institutions which provide free internet content for K–12 education. What the six partners have in common, and what makes this an important and interesting development in online learning, is an adherence to national curriculum and evaluation standards in the subject areas. Material is categorized by grade level and individual items are matched to individual learning topics. Despite its strengths, however, MarcoPolo is a closed project; only the six member institutions contribute content. There is no centralized search facility and no metadata listings for the resources. http://marcopolo.worldcom.com/

XanEdu. Xanedu is a learning resource site that collects articles from journals, magazines and other resource providers. Instructors may compile 'course packs' consisting of collections of these materials; students who subscribe to XanEdu may access these course packs. The materials are sorted by category and may also be located using a search mechanism. Like MarcoPolo, however, XanEdu is a closed project. It draws materials only from selected publishers. And while it allows subscribed students to browse through its materials, the vast bulk of resources available on the internet cannot be found through XanEdu. http://www.xanedu.com/

18.2.4 Learning Object Repositories

Learning objects are stored in databases called learning object repositories. There are two major types of repositories: those containing both the learning objects and learning object metadata, and those containing metadata only. In the latter case, the learning objects themselves are located at a remote location and the repository is used as a tool to locate learning objects. In the former, the repository may be used to both locate and deliver the learning object.

Most learning object repositories are stand-alone. That is, these repositories function a lot like portals in that they contain a web-based user interface, a search mechanism, and a category listing. Another major class of learning object repositories functions more like a database attached to another product. An LCMS, for example, may contain a learning object repository intended for its exclusive use.

Two major models for learning object repositories exist. The most common form is a centralized form in which the learning object metadata is located on a single server or website (the learning objects themselves may be located somewhere else). An alternative model is the distributed learning object, in which the learning object metadata is contained in a number of connected servers or websites. Distributed learning object repositories typically employ a peer-to-peer architecture to allow any number of servers or websites to communicate with each other.

The following are examples of some learning object repositories:

Merlot. Described above, Merlot is probably the most well known learning object repository. Merlot is a centralized repository containing metadata only and pointing to objects located at remote locations. It is stand-alone, acting like a portal for learning objects. In addition to providing search and categorization, Merlot provides a peer review service provided by communities of experts in different subject areas. http://www.merlot.org

Campus Alberta Repository of Educational Objects. CAREO is a centralized collection of learning objects intended for educators in Alberta, Canada. A stand-alone repository, CAREO contains metadata and provides access to learning objects located on remote web servers. http://www.careo.org

Portals for Online Objects in Learning. POOL is a distributed (peer-to-peer) repository system under development intended to create a pan-Canadian repository of learning objects. A primary objective of POOL is to develop and distribute tools for creating connected learning object repositories. http://www.newmic.com/pool/ (not currently functioning). See also http://www.canarie.ca/funding/learning/1999backgrounders/pool.html

National SMETE Distributed Library. In development for the (SMETE), NDSL is intended as a "federation" of learning object repositories, each library using different document formats, different systems of classification, and different database and repository management schemes. NDSL is intended to join these libraries using a common search engine called Emerge and a method for sharing resources called LOVE (Learning Object Virtual Exchange). (Chen, 2001)

18.2.5 Learning Content Management Systems

Learning objects are typically small, consisting of no more than the equivalent of an hour or two of instructional time (there is some debate as to how small a learning object may be and whether educational content must contain pedagogical features, such as a statement of learning objectives, in order to qualify as a learning object). Most educational institutions deliver larger chunks of instruction, called courses. To create a course, therefore, a set of learning objects must be assembled into a package.

Packages organize learning objects sequentially. In order to create a course out of, say, a dozen lessons, where each lesson is a separate learning object, a course author arranges these lessons into a sequence. In some cases, where the learning objects are smaller units, course designers may need to create lessons composed of a sequence of individual modules, then the course as a whole out of the sequence of lessons. However created, the sequence of objects is used to define course-specific entities as the course outline or table of contents.

Packages are created using a Learning Content Management System. While a course author could locate and assemble learning objects by hand, it would be tedious and unproductive to do so. Courses created using learning objects are typically created using a development environment called a Learning Content Management System (LCMS). The LCMS performs two major functions: it provides authors with a means of locating learning objects, and it assembles them into standards compliant learning packages (or courses). (Ellis, 2001)

Though many types of LCMS are available, the typical LCMS will contain four essential features: an authoring application similar to the computer assisted software environment (CASE) described above, a collection of learning objects (called a repository), a means of sending the completed course to a delivery system (called a delivery interface), and administration tools.

http://www.internettime.com/itimegroup/lcms/IDCLCMSWhitePaper.pdf

http://www.internettime.com/itimegroup/lcms/

Using an LCMS, a course author defines major features of the course: its topic area, say, or its grade level. The author then instructs the LCMS to search through the learning object repository for relevant resources (because the data is in XML, the search can be very precise). From the search results, the author may review a learning object or select it for inclusion in the course. The LCMS retrieves the object metadata from the repository and inserts it into the course package. The LCMS automatically adds institution-specific formatting and prepares the package for delivery.

18.3 Problems and Issues

18.3.1 Overview

In general the issues surrounding the location, distribution and reuse of learning resources online have to do with system architecture and resource based on what I call the "silo model." On the silo model, resources are not designed or intended for wide distribution. Rather, they are located in a particular location, or a particular format, are intended for one sort of use only.

The silo model is dysfunctional because it prevents, in some essential way, the location and sharing of learning resources. In an important sense, such resources or architectures are broken because they require some additional step, usually involving manual labour, in order for developers or learners to make use of the material. The requirement of such a step adds significantly to the cost of a learning resource and in some case may prohibit its use altogether. In fairness, this cost or prohibition may be imposed by design. But from the point of view of a learning object economy, the resource or architecture is unusable.

There are numerous ways a learning resource or architecture may follow the silo model. In this section, a number of these are listed. Few products embody all of these problems. But most contain instances of at least one of these problems. And even a single instance of the silo model is enough to prevent a learning resource or architecture from being used as part of a network.

18.3.2 Proprietary Standards

A standard is proprietary when it is secret or when patents, copyrights or other restrictions prohibit its use. The standard is created by a commercial entity and specifies "equipment, practices, or operations unique to that commercial entity." (National Communications System, 1996) With the advent of the internet, proprietary standards are much less of an issue than in years past. Nonetheless, proprietary standards continue to abound, especially in the realm of multimedia formats.

The use of a proprietary standard divides a distribution network into those people or systems able to use the standard, and those people or systems unable to use the standard. For example, a document created using DXF for Autocad may not display properly in Cadkey, which uses CADL, or ACIS, which uses SAT. Another example is XrML, a digital rights management language developed by ContentGuard. Developers have been reluctant to use the standard because of Microsoft’s control over the standard. (DRM Watch, 2002)

Proprietary standards pose numerous risks to developers. One risk is that the standard will cease to be supported in new software. Documents encoded in older MS Word formats, for example, need to be converted before they can be used. There is the risk that licensing terms may change, and as a consequence, require that user pay unexpected licensing fees. If the standard is not widely shared or distributed, as is the case, for example, with Microsoft Windows, it is difficult to develop new applications, and the holder of the standard enjoys an advantage over competing products. Additionally, the choice of viewing software may be limited. Because of these risks, it is difficult to encourage wide adoption of proprietary standards.

Several of the systems listed in the previous section depend in whole or in part on proprietary standards. Course packs designed for Web CT, for example, cannot easily be used in competing learning management systems. It is necessary to use a content migration utility (some versions of which are no longer supported) to

obtain interoperability. http://www.webct.com/IMS

18.3.3 Overly Strict Standards

Even when a standard is non-proprietary, it may be the case that the standard is too limiting for widespread use. If, for example, a standard requires that only a limited type of data will be transported by a data transmission system, then novel applications using different types of data will be impossible to develop.

Much of the criticism around the Sharable Content Object Reference Model (SCORM) was focused on this sort of objection. SCORM was developed to support self-study modules designed for use by the U.S. Military. Learning objects defined using SCORM are mutually independent, meaning that only the most basic sort of