New Kinds of Learning Objects

Steve Bennett
Senior Lecturer in Computer Science
Blended Learning Unit Teacher
S.J.Bennett@herts.ac.uk, ext 4309

Summary
This article describes two aspects of thinking about learning objects, their desired "standalone" nature and certain best-practice recommendations regarding their production. The article disputes that a learning object can ever truly be "standalone" and argues that the best way to get them produced is by the lone academic herself rather than the traditional managed production teams. To support this it cites the example of open-source software production and online collaborative enterprises such as Wikipedia which achieve their aims without any traditional management layer, but rather function on a volunteer basis. The author concludes with a case study where he attempted to put these methods into practice in producing materials for a Master's multimedia course.

In recent years a number of software engineering terms have been applied to the processes of learning materials development, of which the two most common are 'reusability' and 'learning objects'. In laymen's terms, 'learning objects' are standalone units of learning material created to allow their easy aggregation into larger units and courses. These range from small interactions or diagrams, to chapters of books, web pages, but potentially as large as lesson series or whole courses! These developments are a response to the problem of past resistance to the transferability of electronic learning materials between institutions, particularly the dreaded *not invented here* syndrome. In this new 'learning object economy', it is hoped that people and institutions will generate and share their learning materials, and the development costs of artefacts thus contrived will be amortised by a cheery cross-institutional barter.

This is an attractive proposition, but there are two potential weaknesses in it. The first concerns whether the software engineering/instructional design analogy actually fits: can a lesson really be a *component* that can be *slotted* into another course in the same way that a *widget* can be *slotted* into a word processor? Isn't the *nature* of learning objects completely different from that of software components? Can courses be assembled like Ikea flat-packs?

The second issue relates to which form of software engineering the analogy is based on and how this impacts on the production process. Recently we have seen the emergence of large software projects, such as the Linux operating system, and related endeavours like the huge online encyclopaedia, Wikipedia, driven by a loosely co-ordinated set of people all over the globe on an almost volunteer basis. This is known as the *open source* model and is quite different from the highly *managed* methods of software production in the past – and yet, it is precisely this old monolithic method that seems the preferred analogy for learning technology enthusiasts.

My question, therefore, is have we got the fundamental nature of learning objects wrong, but have we also got the production process required to generate such objects cock-eyed?

The standalone learning object
To examine the first issue I propose to look at one of the most eloquent expressions of the learning object paradigm, namely Tom Boyle's *'Design principles for authoring dynamic, reusable learning objects'.*

This paper essentially describes principles for authoring learning objects, such that objects may be re-used in other courses, or similarly, that they may be optionally invoked by the student doing a course, depending on their level of proficiency. In order to facilitate re-use, says Boyle, the learning object should be 'standalone'.

> "This principle states that the unit (software module/learning object) should have minimal bindings to other units. Thus the content of one learning object should not refer to and use material in another learning object in such a way as to create necessary dependencies. One object then cannot be used independently of the other (Sommerville, 2000; Pressman & Ince, 2000). This principle is crucial in design for reuse. The learning object should, as far as possible, be free standing" (Boyle, 2003).

To most teachers, this proposition might be taken as a bit puzzling. While 'minimal binding' is likely to be the most important feature for the

programmer in developing class libraries, in order that the developer who wants to re-use them is not encumbered with built-in premises and dependencies, for the teacher, the notion of a 'minimally bound' piece of content, something that doesn't refer to anything that came before it, or anything that comes after it, might be regarded as absurd. The analogy might be a lecture where you are forbidden to cite previous lectures, and are no longer allowed to say, "you remember last week when we covered…"

 If other content is invoked in a URL, says Boyle, it should not be from the learning object, which, like a royal bride, needs to preserve its purity, but rather from some kind of structure around it.

> "The primary design feature is that the URLs must not be mixed in with content. They must be kept and managed on a distinct area of the screen. This produces minimal explicit bindings between the main content and the URL links. The URLs can be added to, subtracted from, or modified without affecting the core object structure" (Boyle, 2003).

The fundamental problem with this, however, is its unnaturalness. It is the most normal thing in the world for web pages to link to other web pages. Also, the great advantage of the online link is its rich contextualisation of the material that is being referenced. If I write in a page, "a wonderful explanation about xyz can be found at…" – I am not only pointing to a resource, but I am also affirming its status and validity. The mere putting of links to other learning objects without comment, however, seems to affirm an aridly goal-driven concept of learning, where the link becomes an instrumental expedient to achieve the 'learning outcome'.

One of Boyle's major subheadings is 'Towards a synthesis of software engineering and pedagogical principles' – but perhaps this is precisely the problem. The way in which things are re-used in the engineering as opposed to the cultural sphere is completely different. Re-use in the cultural sphere takes place all the time: in music we have sampling, in journalism we have quotations. However, is it true that the core property that encourages re-use in this sphere is context independence or in Boyle's words 'minimal bindings'?

I would suggest not. Moreover, I would argue that human beings are perfectly capable of dealing with contextually situated meaning, where certain of its particulars are 'bound' to some precursory context, but also capable of abstracting it from that context in order to gather some value from it. Metaphor is based on this capability. Why was *Pride and Prejudice* recently turned into a feature film when many aspects of its context and 'bindings' (the mores of the early 19th century gentry) are very distant from us? I would argue, because we ourselves find a way to re-use them, to transplant them from their originating contexts and make them serviceable to our own preoccupations. In fact, one way out of this bind would be to replace the concept of 'reusability' with a concept of 'iterability'.

This idea was developed by Jacques Derrida to describe the inherent reusability of all language, the ability of all communication to break beyond its originating context and to continue to signify in contexts unforeseen by the author (Derrida, 1977). This helps explain why *re-usability* is not something that truly needs to be worked at, but occurs naturally whenever human beings communicate. But the real question is what makes other people *want* to re-use things? In an essay *Against George Lukács,* Bertolt Brecht writes about aesthetic value in a way that I would suggest is also valid

for pedagogical value:

> "Anything that was worn out, trivial, or so commonplace that it no longer made one think, they did not like at all ('You get nothing out of it'). If one needed an aesthetic, one could find it here" (Bertolt Brecht, quoted in Eagleton (1976)).

Therefore, let us return to Boyle's point about dependencies:

> "Thus the content of one learning object should not refer to and use material in another learning object in such a way as to create necessary dependencies. One object then cannot be used independently of the other" (Boyle, 2003).

I disagree. It is possible for one learning object to refer to another and still be used independently in another context. All the time we read articles where some of the particulars and references are opaque to us, but do we give up whenever something like that occurs? No, we live with it. It is for the reader to decide whether the bindings are a prerequisite or not for comprehension.

Models of production

My first exposure to learning technology came as a developer in the TLTP project GeographyCal – a project for producing some core first-year geography materials. In that project, academics wrote 'storyboards', which developers like me would then turn into Toolbook multimedia presentations. These were then checked by the original authors before finally being distributed on a series of floppy disks.

We might call this process of development the ERICC model (Elicit, Repurpose, Inspect, Confection and Circulate). And yet while the technology has moved on to such an extent that a number of the assumptions implicit in that model are no longer operative (that academics could only supply *raw material,* not the finished artefact, that the presentations had to be distributed on physical media), it seems that sophisticated variants of that way of thinking are still very current.

A detailed example of that ERICC-style thinking can be found in Leeder and Morales' 'Universities' Collaboration in eLearning (UCeL): Post-Fordism in action.' Here they describe a model of production of 'learning objects', which happens in 8 stages:

1. Content submission by academics.
2. Rewriting by editors.
3. Peer review.
4. Addition of media by editors.
5. Despatch to major developer who combines everything together.
6. Checking for functionality.
7. Second peer review.
8. Delivery and indexing on the web.

While it is certain that such a model has the potential to establish a very effective proofreading system, I fear that it may do so at the cost of diminishing the joy to be derived in creating such materials.

There are three distinct authorial inputs: the academic, the editorial and the multimedial. At a number of points along the way, these agencies are also supplemented by peer review, both at academic and usability level. While in functional teams this is likely to have positive feedback loops, equally there is the danger of academics feeling they lack ownership of their materials.

From my own experience, I know that two consequences of this style of development, where academic *content* and multimedia *form* arrive from two separate agencies, are 1) a kind of lazy *interactivisation* and 2) a manic

beautification. By lazy interactivisation is meant those kinds of pages where concepts are expressed in buttons which, when clicked on, subsequently reward the viewer with crisp definitions popping up. This creates a kind of pseudo-interactivity, but is actually less efficient than simply printing on paper. Manic beautification is usually expressed in graphs where the histograms animate up one by one or line graphs where the curve gradually reveals itself.

The author's own experience of such models, in the TLTP GeographyCal project and the EU Euromet project, has shown the ERICC model can certainly work, but does so usually when the developers themselves have keen subject knowledge. However, in subsequent smaller projects with which I have been involved, significant problems have been found relating to the sense of academics' ownership over the materials to which they were contributing.

While Leeder and Morales talk about 'post-Fordism', they are nonetheless sketching a system of production that is heavily *managed*. In addition, in their requirement for indexing of the learning material, they are also describing a product that is heavily *fixed*, i.e. not something subject to variation, editing or refactoring, but true for all time. However, the most interesting phenomena in the area of cultural production these days is the anarchic and most definitely unmanaged and unfixed world of blogging and wikis, as in the enormously influential online encyclopaedia, Wikipedia.

Emergent authoring
Blogs are essentially web diaries where the author posts his/her thoughts on an almost daily basis. Wikis are collective websites authored by groups of people, where articles posted are editable by other members of the group, and where new articles may be created by simply putting in a new link from the current

one. Wikipedia is the best-known example of this, and it is an online encyclopaedia where authors worldwide post definitions and illustrations of various terms.

These have been tremendous successes, but what characterises them is the sheer immediacy of their productions, and their completely unmanaged nature. Opposite to the ERICC model above, these place the act of publishing as the *first* step in the materials development process and not the last. It is the basis upon which quality control begins, not the act that seals it. Moreover, the quality control in the process comes through the queries, refutations and, in the case of Wikipedia, edits and corrections that are spontaneously made by other interested users. There is no sense that the words or images posted are true for all time: they are contingent, impermanent, self-consciously aware that they are reflections of the state of knowledge at any one time, always revisable.

Wikis, largely, are the application in the cultural domain of techniques of production that first became common in the software development domain, of which some of the most extraordinary results are the operating system Linux and databases like mySQL. A brilliant description of such self-organising complexity has been written by Eric Raymond in his seminal essay 'The Cathedral and the Bazaar'.

The cathedral and the bazaar
Raymond's basic question is how can a system as complex and robust as Linux emerge from the unco-ordinated (or only lightly co-ordinated) efforts of a loosely coupled network of volunteer programmers.

"The Linux community seemed to resemble a great babbling bazaar of differing agendas and approaches (aptly symbolized by the Linux archive sites,

who'd take submissions from *anyone*) out of which a coherent and stable system could seemingly emerge only by a succession of miracles…The fact that this bazaar style seemed to work, and work well, came as a distinct shock" (Raymond, 2000).

To describe what makes this work, Raymond sets out a series of 20 'lessons', which explain why stable and coherent programs can emerge from a chaotic maelstrom. Some of them are very specific to the process of software production, but certain others would have great relevance to the production of learning material, and it is these I will seek to address.

> "Good programmers know what to write. Great ones know what to rewrite (and reuse)" (Raymond, 2000).

Academics do this too: either by borrowing the notes of colleagues within their institutions, or those from other institutions who have published to the web! However, equally, most academics have to make these materials *their own*, just as programmers need to make the routines they incorporate into their programs *their own*, insert them in a particular order in the presentation and unite them with other examples, perhaps not those that the original author used. They need to metabolise what they re-use.

Raymond's most famous lesson is number 7:

> "Release early. Release often. And listen to your customers" (Raymond, 2000).

Unlike the 'Cathedral model' of software development, where releases take place only after extensive testing, in the 'Bazaar model', new builds are produced on an almost nightly basis. What he is trying to say, of course, is that it does not matter if you release unstable versions of software (most Linux-style software has the option of latest builds by day, but also 'latest stable release'). The important thing becomes the speed to market. Publication does not occur after a long sequence of bug checking, but self-confessedly in an imperfect state, in order that the imperfections can be found, through 'listening to your customers'.

There is a huge emphasis in Raymond on the development of community. Describing his efforts at community building during his project Fetchmail he writes:

> "In order to build a development community, you need to attract people, interest them in what you are doing and keep them happy about the amount of work they're doing … The personality you project matters too" (Raymond, 2000).

To a certain extent, fellow academics are our community, but the community we have to create is ultimately that with the students. Just as Raymond polled his beta-testers about design decisions, we need to poll our students about the materials we give them: we need their evaluations, and for them to point out our deficiencies and for us to be honest with them.

Applying Raymond to learning objects
So how do we apply these Raymond principles to producing learning objects? Might an equivalent of Wikipedia be created in which are contained lectures and quizzes instead of definitions? Yes, but only if the speed of creation becomes as fast in writing learning materials as it is in Wiki contribution or coding in open-source software. There are three principles necessary for generating fast recombinable learning objects, which I will call the three 'I's: Individualism, Immediacy and *In media res.*

Individualism

What makes people want to participate in these big online social enterprises is ultimately a desire for personal achievement. Raymond writes about the motivations of software hackers:

> "The 'utility function' Linux hackers are maximizing is not classically economic, but is the intangible of their own ego satisfaction and reputation among other hackers. (One may call their motivation 'altruistic', but this ignores the fact that altruism is itself a form of ego satisfaction for the altruist)" (Raymond, 2000).

The Leeder and Morales *post-Fordism* denies this gratification: while the degree of pre-publication editorial validation might inhibit the expression of falsehoods or infelicities, it equally might inhibit expression *itself*. Since there is not a financial market for lectures typically produced by your average academic, the motivation for publication has to come from elsewhere, and, if there is no money involved, the peer-approbation that Raymond describes seems the only likely source. Is there a greater motivation to produce these things than the fact that your own students appreciate them, and your peers want to use them?

Immediacy

By this I mean both a sense of 'speed' and a sense of being 'unmediated'. If academics are to freely produce their material, it must be something that fits into their own ways of working, which is aligned to their daily practice and not something that requires wholly new technical skills, nor appreciation of the finer points of graphic design. Neither should it have to go through committees of reviewers and subsequent editors. Ultimately, the location where academics generate the majority of their learning material is in the lecture theatre itself.

Therefore, the kinds of things we need are tools to capture lectures in such a way as to require as little post-editing as possible.

This, however, does not mean just the PowerPoint file itself, which is often little more than gargantuan cue cards when projected in a lecture theatre. Nor is it just a recorded mp3 of the whole thing, which would be an object of such ungranular dimensions as to render it difficult to be incorporated into anything beyond its original context. It has to be a *via media* between the minimalism of the PowerPoint file and the unbearable loquacity of the unbroken mp3 file. In the case study below, I offer my own proposal as to what that *via media* could be.

In media res

This is the famous method that the epic poet Horace advised to the writers of epic poems. Quoting from Wikipedia itself:

> "*In media res* (Latin for 'into the middle of things') is a literary technique where the narrative starts in the middle of the story instead of from its beginning (ab initio). The characters, setting, and conflict are often introduced through a series of flashbacks." (Wikipedia: http://en.wikipedia.org/wiki/In_media_res)

Ultimately, all *learning objects* have to be like this. They presume certain knowledge; they rely on certain past lectures or activities. And they are quite entitled to *invoke* these things, in the middle of their discourses, wherever they want. The aspiration for totally *standalone* learning objects as described by Boyle is a chimerical one. Moreover, the fact that previous material is cited does not hinder the learning process. Whenever we read we are always plunging *into the middle of things,* trying to figure out the context from which it came and to where its references point.

Nevertheless, the absence of those references does not stop us learning, except when crucial information is being withheld. Human beings can live with not having the total context completely at hand.

A small case study

What follows is a brief sketch of how I tried to follow these principles in a Masters course run in Semester A of the 2005/06 academic year.

Essentially, for all my lectures, I recorded them as I gave them, and then broke them down into a series of mini-lectures afterwards. For most of the mini-lectures I wrote a small multiple-choice quiz. Towards the end of the semester I wrote a contextualising paragraph for each of the mini-lectures, which also contained hyperlinks to other relevant material regarding topics upon which the mini-lecture touched. When ready it was released to the students to help them with their revision.

The technique used to record the lectures was the tool Smirk (a multimedia authoring tool) in 'live' mode. Using this, a PowerPoint file was imported and used for the lecture, and voice was recorded through a radio lapel microphone. Since Smirk also records pen strokes made during the recording, what actually was generated was a rich medley of sound, slides and pen markings, but also of very large storage dimensions (approx 260mb for the recording of a two-hour lecture). After the lecture, I would transfer all the data from the laptop to my office PC using a pen drive, and then do the tricky work of breaking the lecture down into logical subunits, things that could exist by themselves.

Along with this, the sound was examined for imperfections and periods of silence eliminated. Finally, these sub-lectures were placed into a categorisation scheme of topics and subtopics, the former being the title of the original whole lecture, and the subtopics, the titles of the sub-lectures into which that lecture was segmented.

One element of quality checking I undertook before release was to make sure that for every sub-lecture there was a contextualising paragraph, and that for most of them there was a quiz. By the end, the number of sub-lectures had reached 58 and this meant there was quite a task of managing the material. In order to do this a small utility was coded to establish the topic/subtopic hierarchy, and check for the presence of contextualising paragraphs and quizzes. This meant a very systematic way of putting in the contextual material could be achieved.

At the end of the whole process I calculated that the time taken to produce the materials (in the sense of record/break down/transfer/add quizzes and contextualising text compared to the time taken to deliver them) was something like 4:1. While this might seem like a lot, assuredly it is much less than would have been required if they were all turned into full text notes, or the old TLTP-style multimedia packages. Moreover, this ratio may be artificially high, since this is the first time that this method for producing material has been attempted. Hopefully, in the future, when tools have matured, a ratio of 3:1 might be possible.

This method followed the three 'I's mentioned above:

1. It was individualistic. I did the whole thing myself.
2. It was immediate. I just recorded myself and then broke up the presentations, wrote contextualising paragraphs and quizzes.
3. All the sub-lectures came up *'in media res'*; the sub-lectures broke immediately into the midst of their topic, making references forward to upcoming lectures and backwards to

previously heard ones.

Initial student feedback

After taking a test, the students were asked to answer a simple questionnaire. In summary, the major suggestions for improvement were:

- Having quizzes on every presentation (not all the presentations had them).
- Improving the sound on certain presentations (this was known to the author since the quality of the microphone deteriorated during the course).

However, most of the feedback was overwhelmingly positive. In response to the question: "What is your overall opinion of the web lectures and what would you recommend to improve them?":

"It was a great help in revising the course. The lectures I had missed, I could listen to them and make notes. It felt as if I am really attending the lecture in the class."

"I think it's great that we have web lectures as it has helped me with my revision by miles! I think the quizzes were a good way to revise and learn so the only recommendation I have is that all the quizzes are there!"

"Great job on the Smirkboard site, definitely the most organized and best presented notes that I have come across. I found the use of audio really helped the revision process, breaking each lecture down into sub-topics made the process even easier. The option to take a quick quiz at the end of each section really made me confident coming into the exam."

What I take from this feedback is the importance of the quizzes: for many students it seems that they are an essential tool for them to self-evaluate their degree of comprehension. Secondly, as the final quotation makes clear, the importance and utility of granulising the course into smallish units. The value of this ultimately is that it gives a structure to the material of the course. Stefano Penge writing of the experience of the online learner writes:

"In the digital era, the author is not (only) a creator of original content, but also and above all an organiser of contexts in which content can serve a precise end…While the traditional student feels always 'protected' by the walls of the lecture theatre where she sits in the company of the class, the distance learning student is only in front of a cold monitor screen, and has obviously a much greater need of points of reference, stimuli, and signs to help her orient herself" (Penge, 2004, p72 – my translation).

It is this facility of self-orientation that is the biggest benefit from structuring and granularising the material.

Conclusion

As they are understood right now, learning objects in terms of their own phenomenology or the conditions of their causation, demonstrate an *insufficiently digital sensibility*. The kind of mob-handed authoring advocated by Leeder and Morales, or the virginal purity desired in the *standalone* learning object of Boyle belong to a pre-Internet, unconnected, world. The lessons of the evolution of the Linux operating system and the Wikipedia project show us how the speed of communications introduced by the Internet can create self-regulating volunteer communities where evaluation or quality control no longer needs to be planned preparatory to publication, but

happens as the consequence of publication. This being so, anything thus published is always in a process of amendment and evolution and can never truly be regarded as 'fixed'.

Therefore, rather than seeking to produce learning objects good for all time through complicated procedures of peer-evaluation, which then get glazed with meta-data description, we should seek to produce ones which are good enough for now, and which will be improved by exposure to the world and its feedback, and carry on evolving in the future.

The reusability beloved of the software developer is not that of the academic author. The academic instead seeks to produce a more gregarious artefact, one that glories in its implication within a network of other sources of knowledge. It is also an artefact that is continuously aware of its own impermanence, its 'editability', the fact that it might say something different on one day to what it said the day before. It is not something that can be tied down forever to the ball and chain of its own meta-data. That is a particularly Gutenbergian and undigital way of relating to a digital resource. Stefano Penge in his superb article 'Being an Author Online' writes:

"A distance learning system is, again, a middle way: it allows, even requires, the author to plan beforehand the whole course, to publish it but also to go and change it or integrate it. And while we shouldn't think of a course as a definitive text, equally neither is it an unfinished rough draft. We are not used to this infinite provisionality of the digital in any sector. Many of us look at the newspapers online, or train times, but then we save the page and even print it 'to have it always at hand'. While television was born as a flux in time, we first got to know digital objects as static forms (floppy disks and cd-roms), and only afterwards as dynamic forms (bulletin boards and internet): and thus we continue to treat them as always the same unto themselves" (Penge, 2004, p71 – my translation).

The precise quality of digital communication is its evanescence, its impermanence, the fact that a URL does not indicate a fixed resource but one open to continual mutation. The action of keeping those resources relevant, of constantly updating them in response to the feedback of the community – that is the true vocation of the teacher and academic.

Notes
For the TLTP GeographyCal visit:-
http://www.geog.le.ac.uk/cti/Tltp/

For the EU Euromet project visit:-
http://www.eumetcal.org/euromet/english/navig/beginn.htm

Smirkboard can be seen at:
http://smirkboard.herts.ac.uk/multimedia/index.htm

References
Boyle, T. (2003) 'Design principles for authoring dynamic, reusable learning objects', *Australian Journal of Educational Technology,* **19**(1), 46-58.

Derrida, J. (1977), 'Signature Event Context', *Glyph*, Vol. I (1977), pp. 172-97 and *Limited Inc.* (2nd ed., Evanston, Ill.: Northwestern University Press, (1988)

Eagleton, T. (1976) *Criticism and Ideology.* London: New Left Books.

Healey, M., Robinson, G. & Castleford, J. (1996) 'Innovation in Geography Teaching in Higher Education: Developing the Potential of Computer-Assisted Learning' in van der Zipp, T; van der Schee, & Trimp, H (compilers) *Proceedings: Commission on Geographical Education, 28th Congress of the International Geographical Union,* The Hague, 4-10 August 1996, pp.199-203

Leeder, D. & Morales, R. (2004) 'Universities' Collaboration in eLearning (UCeL): Post-Fordism in action'. In Kommers, P. & Richards, G. (Eds.), *Proceedings of World Conference on Educational Multimedia, Hypermedia and Telecommunications 2004* (pp.5353-5358). Chesapeake, VA: AACE.

Penge, S. (2004) 'Essere autore on line' in Maragliano, R. (Ed.) *Pedagogie dell' e-learning,* Editori Laterza, Bari

Raymond, E. (2000) 'The Cathedral and the Bazaar', http://www.catb.org/~esr/writings/cathedral-bazaar/cathedral-bazaar/

Biographical notes
Steve Bennett has been teaching Computer Science at the University of Hertfordshire for the last six years and was Learning Technology development manager there before that. He has been involved in two projects which have won prizes at the European Academic Software Award and was runner up in the *Times Higher Education Supplement* E-tutor of the year in 2004. He is also the principal software architect of the Smirk authoring tool.