Learning networks in practice
Downes, Stephen

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version. / La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version acceptée du manuscrit ou la version de l'éditeur.

Publisher's version / Version de l'éditeur:

Emerging Technologies for Learning, 2007

National Research Council Canada   Conseil national de recherches Canada

Canada

National Research
Council Canada

Conseil national
de recherches Canada

Institute for
Information Technology

Institut de technologie
de l'information

NRC·CNRC

*Learning Networks in Practice**

Downes, S.
2007

Canada

Learning Networks in Practice

While the learning management system succeeded in emulating the classroom online, a second wave of applications and approaches, drawing on what has come to be characterized as Web 2,0, is redefining the concept of online learning. This second wave is characterized by the 'personal learning environment' (PLE). The values that underlie the PLE and Web 2.0 are the same: the fostering of social networks and communities, the emphasis on creation rather than consumption, and the decentralization of content and control. But why should we think that these values improve learning. This paper argues that the personal learning environment and Web 2.0 are instances of a more fundamental concept, the learning network, and that networks with identifiable properties, such as the fostering of diversity and autonomy, are more reliable producers of learning and knowledge.

The Personal Learning Environment

Beginning in 2005 and continuing through 2006 discussion at the forefront of the educational technology community centered not around instructional design and the learning management system, but rather on approaches that dramatically shift the centre of e-learning; things like social networking applications such as ELGG, things like informal learning and e-portfolios, and most of all, things like the personal learning environment (PLE). These in turn are centred around, and draw from, a concept in the world of online computing called Web 2.0.

The use of Web 2.0 technologies in education came to be called e-learning 2.0. "Like the web itself, the early promise of e-learning - that of empowerment - has not been fully realized," writes Stephen O'Hear in the Guardian. "The experience of e-learning for many has been no more than a hand-out published online, coupled with a simple multiple-choice quiz. Hardly inspiring, let alone empowering. But by using these new web services, e-learning has the potential to become far more personal, social and flexible." These technologies, in other words, would *empower* students in a way previous technologies didn't.

As O'Hear writes, "The traditional approach to e-learning… tends to be structured around courses, timetables, and testing. That is an approach that is too often driven by the needs of the institution rather than the individual learner. In contrast, e-learning 2.0 takes a 'small pieces, loosely joined' approach that combines the use of discrete but complementary tools and web services - such as blogs, wikis, and other social software - to support the creation of ad-hoc learning communities."

Through 2005 and 2006 in thinking about e-learning 2.0 the concept of the Personal Learning Environment (PLE) began to slowly take form in the educational technology community, coalescing with a 'Future VLE' diagram released by CETIS's Scott Wilson. As described by Colin Milligan, PLEs "would give the learner greater control over their learning experience (managing their resources, the work they have produced, the activities they participate in) and would constitute their own personal learning environment, which they could use to interact with institutional systems to access content, assessment, libraries and the like."

The idea behind the personal learning environment is that the management of learning migrates from the institution to the learner. The PLE connects to a number of remote services, some that specialize in learning and some that do not. Access to learning becomes access to the resources and services offered by these remote services. The PLE allows the learner not only to consume learning resources, but to produce them as well. Learning therefore evolves from being a transfer of content and knowledge to the production of content and knowledge.

As Mark van Harmelen writes, PLEs are motivated by the need for "a standard interface to different institutions' e-learning systems" as well as "pedagogic approaches which require that e-learners; e-learning systems be under the control of the learners themselves." Such a system is needed, additionally, to support mobile learning or offline learning "in a wireless-free hospital, or on a remote mountainside."

The PLE is a recognition that the 'one size fits all' approach characteristic of the LMS will not be sufficient to

meet the varied needs of students. It is, indeed, not even an application *per se*, but is rather a characterization of an approach to e-learning. "The PLE is not a software application as such," writes Graham Attwell, "but rather a 'mash up' of different applications and services although of course, it is possible to develop applications such as ELGG which bring together much of this functionality and allow ease of access to different services."

As such, the key to understanding the PLE consists not in understanding a particular type of technology so much as in understanding the thinking that underlies the concept, and in turn, the responses to that thinking as found in Web 2.0 and, as Attwell notes, "the changing ways in which people are using technologies to communicate and to learn and the accompanying social effect of such use."

The PLE, then, consists in effect of a set of related concepts, each associated with the technologies and applications of Web 2.0, and each describing a shift in emphasis away from that which would characterize learning using the traditional LMS.

Learning in Communities

Frequently mentioned from Wenger forward is the occurrence of learning in what have come to be called communities of practice. According to Wenger, "Communities of practice are groups of people who share a concern or a passion for something they do and learn how to do it better as they interact regularly."

In essence, in this theory, to learn is to immerse oneself in the network. It is to expose oneself to *actual* instances of the discipline being performed, where the practitioners of that discipline are (hopefully with some awareness) *modeling* good practice in that discipline, or as Thomas Kuhn would say, knowing how to solve the problems at the end of the chapter. The student then, through a process of interaction with the practitioners, will begin to *practice* by replicating what has been modelled, with a process of *reflection* (the techies would say: back propagation) providing guidance and correction.

Learning, in other words, occurs in communities, where the practice of learning is the participation in the community. A learning activity is, in essence, a *conversation* undertaken between the learner and other members of the community. This conversation, in the web 2.0 era, consists not only of words but of images, video, multimedia and more. This conversation forms a rich tapestry of resources, dynamic and interconnected, created not only by experts, but by all members of the community, including learners.

Hence in the first instance the tools that characterize Web 2.0 are communication tools. Communication tools support direct interaction between individuals. They provide an individual with a means of communicating with one or more members of a network, and hence, support social networking. Members typically have a unique identity in such systems and communicate with a collection of other people organized either by membership in a group or forum or by belonging to a list of 'friends' or 'buddies' created by the individual.

Instant messaging (IM) has been identified as the predominant form of communication among younger net users. Each of the major IM tools – ICQ, AIM, YIM and MSN – allows a user to create a list of contacts (known as 'friends' or 'buddies'). A similar functionality, SMS, operates on mobile phones. IM is an advance over email because it promotes diversity and decentralization. Each person's list of contacts is unique. Conversations are typically person-to-person (and hence, these are called peer-to-peer (P2P) networks) though in some cases multi-party conferences are created. P2P file sharing networks, such as Gnutella or Kaaza, work along similar principles, though the creation and maintenance of contact lists is handled automatically by the software.

Instant messaging and conferencing tools have expanded from text into audio and video. Skype, for example, is an application that allows free online audio conversations. Each Skype user has a unique identity and Skype users maintain a contact list of other Skype users. Videoconferencing, meanwhile, is already supported by several of the commercial IM products, such as AIM, as well as (more recently) by Skype.

Probably the greatest misapplication of online community in online learning lies in the idea that a community is an adjunct to, or follows from, an online course. This is perhaps most clearly exemplified by the existence in itself of online class discussions. It is common to see the discussion community created with the first class and disbanded with the last. The community owes its existence to the course, and ends when the course does. We see this in the evolution of community on the web as well. Early online communities followed the model proffered

by [Hegel and Armstrong](#), where the community was centered around a certain website, which in turn, would monetize that activity. In both cases, the depiction is community as a group centered on some location or activity.

Community on the web evolved differently, however. While individuals did from time to time cluster around a certain website or service, they did not confine their communications to a single mode or channel. An online community might be a much looser set of associations, what social network theorists such as [Mark Granovetter](#) would call 'weak ties'. A community in this sense could best be described as a cluster of common associations, where these associations are represented as membership in buddy lists, connections in peer-to-peer networks, and other sorts of contact lists. Weak ties are necessary in order to allow the spread of knowledge, and in order for weak ties to be created, "there must be several distinct ways or contexts in which people may form them."

So learning occurs in communities, but communities cannot be based on the group, but rather, the network, where connections cut across existing boundaries, via weak ties, to form layers of association. The implication is that the course content (if any) ought to be subservient to the discussion, that the community is the primary unit of learning, and that the instruction and the learning resources are secondary, arising out of, and only because of, the community. And, in the Web 2.0 world, it was only a matter of time before they were created *by* the community.

Creation, not Consumption

Even LMS-based learning recognizes that learning is best accomplished through some sort of activity, rather than through rote content consumption and memorization. That said, the history of online learning is remarkable for its emphasis on content consumption, as evidenced by the activity surrounding course creation and learning object design. "As learners move beyond content consumption and into stages of critical thinking, collaboration, and content creation," [writes George Siemens](#), "LMS weaknesses become apparent." Content creation tools enable the creation of content. What distinguishes the current set of content creation tools is that the content creation occurs, or is largely supported, online, and hence converts the act of creating content into a social and connected act.

Learning management systems, insofar as they support content creation at all, support online community, or 'group', tools that have their origins in the early days of the web. Their influence has been widespread. Both Yahoo ([Yahoo Groups](#)) and Google ([Google Groups](#)) support massive mailing list and bulletin board services. Large communities have also formed around some specialty sites, such as Slashdot, Metafilter and Digg, each of which displays a series of selected posts, around which a discussion occurs. Smaller communities have also developed using popular content management systems such as Drupal, Plone, PostNuke and Scoop. Some learning management software, such as Moodle, can be used in this way, as for example by EdNA Groups.

These sites all have in common, however, their focus on the group or institution, rather than the individual. Typically, such sites will be managed by one or two people, and other people contribute subject to the benevolence of the owner. Autonomy, therefore, is minimal, and in some cases, diversity is either tacitly or explicitly discouraged. A common complaint found on such sites is the plea to '"stay on topic" or "keep the discussion off-list". Many such groups require registration and identification before posting is allowed, maintain strict acceptable use policies, and often prohibit non-members from viewing the discussions.

Consequently, recent years have seen the rise of personal content authoring and delivery services. The prototypical personal publishing system is the weblog. These greatly simplify personal publishing, allowing writers an autonomous voice, and thus have greatly diversified the content available online. Some blog services are hosted, that is, they are located on a remote server and accessed using a web browser. Early hosted services included Blogger and LiveJournal. Additionally, blogging software allows a user to host a blog on their own server. Moveable Type was an early commercial application, while WordPress is the most popular open source blogging application.

Related to blogging applications is a set of tools known as social networking applications. These services essentially combine the 'buddy lists' of IM with the content creation capacities of blogs. Arguably, LiveJournal

was one of the first social networking applications. Other such systems include Friendster, Tribe, Orkut and Yahoo 360. These sites stressed social interaction. Social networking sites combining personal content creation and interaction, however, took the lead. In 2005, the social networking site MySpace, a music fansite, emerged as a phenomenon, becoming the most popular site on the web. MySpace allowed people to upload photos, music and video. Sites similar to MySpace include Bebo and Facebook, both of which are marketed directly to students.

Content creation sites have formed the vanguard of 'Web 2.0'. This movement is based on the idea that the web is a place where people can create and communicate – in other words, to network. Flickr allows people to store their photos online – and to share them with a network of contacts and friends. Podcasting, a phenomenon that began in 2003, involved the creation of MP3 audio files edited using (free) software such as Audacity, then distributed to the world via sites such as Audioblogs, Odeo or iPodder. Some communication tools became content creation tools. Skype, for example, became a popular way to record online interviews and conversations. In 2006, user-created video took the centre stage, with YouTube, a video hosting service, taking the top spot from MySpace. Hundreds more services, allowing users to create all manner of content, were launched, some of the more popular being Jotspot (wiki), Writely (word processing), Gliffy (diagrams) and Jumpcut (online video editing).

What we have seen, in essence, is a convergence between the characteristics that have redefined online community and those that have characterized online content creation. In order to express themselves, web users have moved away from the group sites. The constraints of creating content within a limited environment have been overcome through the use of a network of separate services, each with its own particular capacity, joined together with social networks. The result is that people, including students, have a much greater capacity to create, and therefore, insofar as a capacity to create supports learning, a much greater capacity to learn.

The 'pedagogy' behind the PLE – if it could be still called that – is that it offers a portal to the world, through which learners can explore and create, according to their own interests and directions, interacting at all times with their friends and community. "New forms of learning are based on trying things and action, rather than on more abstract knowledge. 'Learning becomes as much social as cognitive, as much concrete as abstract, and becomes intertwined with judgment and exploration.'" (Graham Attwell)

And – crucially – *teaching* becomes the same thing as well. As I wrote in 2002, "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 forms of learning in much the same way, for much the same reasons."

Context, Not Class

When learning becomes the creation of content in the context of a community of practice, then learning becomes something that is characterized not by instruction in a classroom, but rather by dialogue and communication within a given context. Jay Cross is talking about a similar thing when he talks about informal learning. He writes, " For sixty years, we've thought of learning as residing in the formal models exemplified by schools, universities, and training programs. Common to these top-down formats is a curriculum that rests on the beliefs and worldview of the authorities in charge. Informal learning is more democratic. It's responsive to learners and often ad hoc."

What needs to be understood is that learning environments are multi-disciplinary. That is, environments are not constructed in order to teach geometry or to teach philosophy. A learning environment is similar to some 'real world' application or discipline: managing a city, building a house, flying an airplane, setting a budget, solving a crime, for example. In the process of undertaking any of these activities, learning from a large number of disciplines is required. Indeed, as in the case of electronic performance support systems, these environment may *be* some real world application.

These environments cut across disciplines. Students will not study algebra beginning with the first principles and progressing through the functions. They will learn the principles of algebra and other fundamental subjects as needed, progressing more deeply into the subject as the need for new knowledge is provoked by the demands of the simulation. Learning opportunities - either in the form of interaction with others, in the form of online learning resources (formerly known as learning objects), or in the form of interaction with mentors or instructors -

will be embedded in the learning environment, sometimes presenting themselves spontaneously, sometimes presenting themselves on request.

The idea of context-sensitive learning is not new. It is already supported to a large degree in existing software; Microsoft's help system, for example, would be an example of this were the help pages designed to facilitate learning and understanding. In a similar manner, learners interacting with each other through a learning environment will access 'help' not only with the software but also with the subject matter they are dealing with. Learning will be available not so much in learning institutions but in any given environment in which they find themselves.

The Personal Learning Environment (PLE) ought to be seen in this light. It is tempting to think of it as a content management device or as a file manager. But the heart of the concept of the PLE is that it is a tool that allows a learner (or anyone) to *engage* in a distributed environment consisting of a network of people, services and resources. It is not *just* Web 2.0, but it is certainly Web 2.0 in the sense that it is (in the broadest sense possible) a *read-write* application.

What makes this possible, and what distinguishes the current crop of applications from those that are *merely* content creation tools, is RSS. Originally designed to list indices of newspaper and magazine articles, RSS worked well for personal publishing, and especially serialized content as is found in blogs. RSS allows individual web users to create custom subscription pages for themselves using applications called News Readers. Early RSS readers were stand-alone applications, such as Carmen's Headline Reader and Amphetadesk. Today, news readers have also become online applications, with services like BlogLines and Google Reader being popular choices. Both the Internet Explorer and Firefox web browsers have built-in news readers, while another application allows you so subscribe to blogs by email.

Some services have emerged in an attempt to aggregate all RSS or blog content. Early listings of popular sites included Blogdex, DayPop, PodDex and PubSub. The current leader in this field is Technorati, which indexes some 50 million blogs. Technorati also introduced to the environment the concept of 'tagging', a system whereby, instead of classifying articles according to a pre-determined taxonomy, readers simply picked whatever words they felt appropriate, hence 'tagging' the articles with a vocabulary of their own choosing. Tagging quickly spead to other social networking applications, most notably, Flickr.

What RSS does is to transform a piece of content created by a student or instructor from something that is a static and stand-alone object into something that resembles a stream or a flow. Contents syndicated in RSS become *part of* other contents, and this interaction occurs seamlessly, with no conscious intervention on the part of the creator needed to make this happen. A learning environment that contains RSS feeds becomes dynamic; the contents of those feeds are what makes it dynamic.

The system of linking and metadata employed by blogs using RSS created an open network with a very low threshold for joining. This approach is being emulated in other areas, from the simple and easy web services model, REST, to the grass-roots personal profile metadata format, FOAF. Each step in content organization has tended toward increased diversity and increased autonomy on the part of readers. Additionally, content creation and aggregation applications have become increasingly transparent as RSS and similar formats allow people to extract content, while APIs (such as the Blogger API) allow people to submit content.

Support Tools

In addition to the standard network infrastructure, such as the web browser, probably the most important support tool for Web 2.0 applications will be an identity manager. Numerous attempts have been made, and the web has seen a large number of centralized (or Federated) approaches – from Microsoft Passport to Liberty to Shibboleth. None of these has caught on widely, and while Google and Yahoo have added their own (proprietary) single-signon systems, no user-centered system yet exists. As of this writing, there is hope in the form of some initiatives. Two major commercial distributed identity systems, LID and SxIP, have been proposed. The developers of LiveJournal have proposed an open and non-commercial OpenID system. Various developers have attempted to collaborate, forming a (now quiet) initiative called YADIS (Yet Another Distributed Identity System).

Another major issue surrounds digital rights management. As content is created, reused, repurposed and fed

forward around the web, it becomes both more important (especially from a commercial perspective) and more difficult to assert ownership, much less enforce use conditions. A variety of digital rights management schemes have been proposed, but users have stayed away from such systems (as one person commented, nobody is demanding to be able to do *less* with their stuff), favouring open protocols such as MP3 and (more recently) Flash video. In addition, distributed and lightweight rights expression models, such as Creative Commons and ODRL, have been widely adopted. By *expressing*, rather than *enforcing*, digital rights, these systems enable, rather than restrict, the free flow of information.

The semantic principle postulated by learning networks is a theoretical principle. But an examination of the trends exhibited by Web 2.0 software illustrate this principle in practical use. Online applications in Web 2.0 are supporting greater user autonomy, from greater content creation capacities to better ways to personalize their information sources. They support diversity, allowing not hundreds but millions of different voices to be heard, and to be heard not only in text but in all manners of multimedia. The applications support openness. They tend to support simply and widely usable protocols, open standards, open source applications, and even open identification and open digital rights. And they support interactivity, supporting communication at all levels.

Learning Networks

Why this, rather than that? Why the PLE and learning networks, rather than the LMS, the lecture and the class?

Taken together, the ideas that underlie the PLE – learning in communities, creation over consumption, and context over class – constitute an instance of a more general approach that may be characterized as 'learning networks'. A network is a collection of connected entities, where a connection is something that allows one entity to send a signal to another entity. The internet is a network; it connects computers together and allows their operators to send messages to each other. And as we have seen, the users of Web 2.0 applications organize themselves into a network as well.

When networks are properly designed, they reliably facilitate learning. This is because, when properly designed, the network will itself learn. Through the process of interaction and communications, the entities that constitute the network will form a mesh of connections. Knowledge is embedded in this mesh of connections, and therefore, through interaction with the network, the learner can acquire the knowledge. Foresters learn about trees by working with foresters; lawyers learn about the law by working with lawyers.

It is the *organization* of the network that supports learning, and that if the network is designed appropriately, it will *organize itself* – just as we see happening in Web 2.0 communities – in order to best support learning. Thus, when we talk about 'learning networks' we are talking about networks in two distinct ways: first, we are talking about *the use of networks to support* learning, and second, we are talking about *networks that learn*. Though these may seem to be very distinct, the central thesis of 'learning networks' as a theory is that these two things are one and the same.

The theory, though, does not describe the particular type of organization that best facilitates learning, partially because there is no one way that fits that description, but also because any such organization is so complex it defies description (it would be akin to attempting to describe the knowledge that 'Paris is the capital of France' by describing a particular set of neural connections). Hence, what is described are the *properties* of the network that are known to most reliably lead to network knowledge. And as seen from an examination of learning networks therefore depend on a 'semantic principle', consisting of four parts:

First, **diversity**. Entities in a network should be diverse. In a society, this means involving the widest possible spectrum of points of view. In a human mind, this means being exposed to a wide spectrum of experiences. Diversity allows us to have multiple perspectives, to see things from a different point of view. These views moderate each other, and prevent us from jumping to a conclusion. Diversity is supported through weak ties. The loose connections enabled through the use of social networking applications allows us to reach beyond our groups and to connect with, and learn from, a wide range of influences.

Second, and related, **autonomy.** Each entity operates independently of the others. This does not mean that it operates without input, but rather, it means that it operates according to an individual and internal set of principles and values. Autonomy is what allows diverse entities to respond and react in a diverse manner. Autonomy is enabled through personal software environment. In Web 2.0, it is enabled through the provision of content creation tools such as blogging software. In learning, it is enabled through a personal learning environment.

Third, interactivity, or **connectedness**. The knowledge produced by a network should be the product of an interaction between the members, not a mere aggregation of the members' perspectives? A *different* type of knowledge is produced one way as opposed to the other. Comparing two points of view, for example, allows us to see what they have *in common*, while merely counting or aggregating views forces us to pick one or the other. Web 2.0 software is about much more than listing connections or tallying memberships. It is about the conversation that happens between individuals. And so, too, the personal learning environment supports not just content consumption but interaction and communication.

Fourth, and again related, **openness**. Each entity in a network must be able to contribute to the network, and each entity needs to be able to receive from the network. Openness is what makes interactivity possible; barriers that make it difficult or impossible to communicate within the network limit the network's capacity to learn. Web 2.0 software freed users from the confines of mailing lists and discussion boards, environments owned by authorities where access was controlled and often restricted. Personal learning environments allow the learner to take their learning out of the classroom and to make it something they can share with the world, to make learning the *result* of sharing with the world.

All learning technology will be at least to some degree network technology, since it is designed to facilitate the interaction between public knowledge and personal knowledge. Thus though these principles may be theoretical in origin, they can be employed in practice as a metric for selecting and designing learning technology. Learning technology that promotes autonomy, encourages diversity, enables interaction and supports openness will, in the main, be more effective than technology that does not. And thus we will see learning technology evolve from the approach defined by the learning management system to the idea that is the personal learning environment.