How the Net Works

Downes, Stephen

National Research Council Canada
Conseil national de recherches Canada

Institute for Information Technology
Institut de technologie de l'information

NRC · CNRC

How the Net Works *

Downes, S.
October 2007

* published in CEGSA RAMpage Magazine. October 2007. NRC 49873.

Canada

How the Net Works

Stephen Downes
July 27, 2007

The purpose of this paper is to describe how network learning works and to show how an understanding of network learning can inform the design and evaluation of online learning applications.

1. Models

The title of this paper does not refer to the Internet or Internet technologies specifically, but rather, at the use of networks and network theory generally in support of teaching and learning.

The network approach to learning is perhaps best contrasted with what might be called the transmission model of learning. According to the translmission model, teaching consists essentially of the transfer of educational content from experts to learners. This creates a distance that must be bridge by pedagogical practices. Such a model informs, for example, Moore's transactional distance theory. (see http://www.ed.psu.edu/acsde/deos/deosnews/deosnews1_25.asp)

Most educators conform to the transmission model. In a startling study, Melissa Engleman and Mary Schmidt found that 85 percent of teachers surveyed fall into the 'SJ' category of the Myers-Briggs Temperament Indicator. While there is certainly room to question both the measure and the measurement, it is nonetheless illustrative that almost all teachers would select responses that indicate a preference for learning through identifying and memorizing facts and procedures, step-by-step presentation of material, and consistent, clearly defined procedures, order and structure. (see http://jolt.merlot.org/vol3no2/engleman.htm)

It is the transmission model that has informed much development of learning technologies to date. As Norm Friesen illustrates, the existing paradigm is to assemble units of learning, called 'Sharable Content Objects' (and later, 'Learning Objects') in a learning management system into sequences of learning. These would then be broadcast by various means into students' minds. "The end result of this approach," writes Friesen, "is to understand training and the technologies that support it as a means of 'engineering' and maximizing the performance of the human components of a larger system." (See http://www.learningspaces.org/n/papers/objections.html)

But learning is not accomplished merely by transferring information from sender to receiver. Learning is not merely the remembering of information. We can see this clearly by reflecting on cases where something has been remembered, but not learned:

- in language, for example, people can remember nonsense terms (such as a line from a Lewis Carroll poem, "Twas brillig..."), and people can remember (and attempt to use) words without knowing what they mean.

- in mathematics, for example, people can learn how to add and multiply, and yet fail to appreciate quantities; consequently, the retail industry has developed a skill, 'counting change', to prevent simple mathematical errors.

Rather than being a process of acquiring something, as commonly depicted, learning is in fact a process of becoming something. Learners do not 'receive' information which they then 'store', they gain experiences which, over time, result in the formation of neural structures. To learn is to instantiate patterns of connectivity in the brain. These connections form as a result practice and experience. They are not constructed; a student does not 'make meaning' or 'construct meaning', as

sometimes depicted in the literature. Connections are grown, not created; meaning is, therefore, grown, not constructed. (See http://www.ag.ndsu.edu/pubs/yf/famsci/fs609w.htm for some quick examples; I also recommend Joseph LeDoux, The Synaptic Self, for a detailed discussion of this point)

Knowing how we learn is important because it tells us a lot about what we learn. And this, again, gives us evidence showing that learning is not merely the acquisition of knowledge and information. It is not, because there isn't anything that can stand on its own as an instance of 'knowledge' or 'information' to begin with. We sometimes think of knowledge as structured, ordered, and sentential. 'Paris is the capital of France,' for example, might be an instance of knowledge. But this is not in fact what we learn. We may use the same sentence to communicate, but what was in your mind and what is in my mind is very different.

Specifically:

- a great deal of knowledge - possible most of what we know - is 'tacit'. That means it is 'ineffible'. It cannot be expressed in words at all. As Michael Polanyi describes in Personal Knowledge, our knowing how to ride a bicycle cannot be expressed in words. (See http://www.infed.org/thinkers/polanyi.htm)

- knowledge is also irredibly personal. What something means depends on the context in which it is understood. Context infuses all levels of language and communication, from the meaning of a given word to scientific explanations and attributions of cause. What something means depends crucially on what else it *could* be, and this is not a matter of fact, but rather of one's beliefs and opinions. A good way to see this is to think of the 'meaning' of a painting. The meaning of words works in a similar way. (See http://www.artandperception.com/2006/11/the-hijacking-of-meaning.html)

2. Learning

To understand what learning is, it is necessary first to understand what knowledge is. As stated above, knowledge is *not* the accumulation of a set of propositions. Rather, it is the development of a pattern of connectivity in the brain. These patterns of connectivity correspond to the skills, abilities, intuitions and habits that we develop over time. A good example - and a good way to understand how knowledge characteristically works - is the process of *recognition*. When we see something, we say we 'know what it is' when we recognize it. What has happened is that a phenomenon in front of us, a tiger, say, has stimulated an appropriate pattern of connectivity in the brain - a different pattern for each person, depending on what their previous experiences of tigers (and things related to tigers) has been.

Learning, on this model, is perception. It is the having of the experiences that lead to the formation of a certain pattern of connections in the mind. It is the growing of new patterns of connectivity through repeated exposure to certain phenomena or the repeated performance of certain activities. Learning is thus very similar to exercise. At first it's awkward and you don't know it very well. But with repeated use and practice, it becomes instinct. Habitual. Expert, as described by the Dreyfus model. (See., eg., Dreyfus, H. (2001) On the Internet) and elsewhere.

The 'knowledge' we have is, in essence, the patterns of connectivity we have in our mind. Or, we might say, the knowledge *is* the network. What does this mean? It means that what we think of as 'knowledge' has changed:

- we used to think of knowledge as governed by rules, principles and universals - statements like 'all ducks are animals' or 'rain is caused by evaporation'

- but knowledge actually consists of - and should be described in terms of - patterns and similarities. Knowledge consists of being able to recognize ducks, for example, or to be able to recognize when it is likely to rain. (To really get this,

compare section 4.1.2, 'The semantics of similarity', in http://www.bbsonline.org/Preprints/OldArchive/bbs.tenenbaum.html with Tarski semantics, as described in http://plato.stanford.edu/entries/tarski-truth/)

When we think of knowledge as 'recognition', we can think of numerous cases where we've seen it in operation before. 'Knowing' is like 'snapping to attention'. Like when you find 'Waldo' in "Where's Waldo", for example. (See http://en.wikipedia.org/wiki/Where's_Waldo) or when you recognize a duck-rabbit image as a duck or a rabbit (again, notice how context and personal variability plays a role here). (See http://www.mindfake.com/illusion_25.html) Or any of the numerous 'out of the blue' experiences described by Tom Haskins. (See http://growchangelearn.blogspot.com/2007/02/emergent-learning.html)

The way networks learn is the way people learn. Network learning is the same thing as personal learning.

3. Personal Learning

By 'personal learning' we mean learning conducted by oneself, for oneself, what Jay Cross means by 'informal learning'. (See http://informl.com/) Probably the best indicator of what works in informal e-learning is what works on the web in general. After all, this is where much informal learning is already taking place. And the web is a medium that supports informal, random-access on-the-job training.

Looking at successful websites in general (and looking at usability, information architecture, and other design documents) we can identify three major criteria: interaction, usability and relevance.

By 'interaction' what we mean is the capacity to communicate with other people interested in the same topic or using the same online resource. In a learning environment, interaction means the capacity to speak with your fellow students or your instructor. Of course, online, such roles are not so distinct - your student at one moment may be your instructor the next, depending on the subject.

Interaction is important for two major reasons. First, it helps us understand that there are people out there, that we aren't merely communicating with a machine - what Terry Anderson would call 'presence'. (See http://cade.athabascau.ca/vol14.2/rourke_et_al.html) We need presence to held devlop cognitive skills and to feel the supportive environment that supports growth. (See http://cde.athabascau.ca/online_book/pdf/TPOL_chp11.pdf)As any user of one of those automatic telephone answering services can attest, when you want to be heard there is little else more frustrating that speaking to a device that cannot understand you.

But more than the human contact, interaction fosters the development of human content. A bundled training program can give a learner a lay of the land. But even the best designers cannot create lessons for every contingency (and even the best learners are unlikely to sit through them all). This is why stories are so important in learning (See http://reviewing.co.uk/stories.htm) - and so frequently found on internet bulletin boards.

By 'usability' we mean the ease with which desired objectives may be satisfied using an application or appliance. For example, is a site is a search site, 'usability' refers to how easy it is to successfully locate a desired search result. Probably the most usable websites on the internet are Google and Yahoo. And between the two sites, designers have hit on what are probably the two essential elements of usability: consistency and simplicity.

Simplicity is the feature that strikes the user first. Many of us probably recall Google's debut on the web. At that time, it was little more than a text form and a submit button. Results listings were unadorned and easy to follow. Simplicity has long been the path to online success. Amazon made buying books online simple. eBay

made hosting an online auction simple. Blogger made authoring your own website simple. Bloglines made reading RSS simple. The web itself is actually the simplification of earlier, more arcane technologies like Gopher, Archie and Veronica.

Consistency is less well understood but we can get an idea by looking at the links on both Yahoo!'s and Google's cureent sites. What you won't find are things like dropdown menus, fancy icons, image maps and the other arcania of the typical website. Links on both Yahoo! and Google are not only simple, they are consistent: they are the same colour and the same type throughout the site, for the most part unadorned. They use the ultimate standard of consistency: words - a system of reference with which readers are already familiar.

By 'relevance' we mean the principle that learners should get what they want, when they want it, and where they want it. What learners want is typically the answer to a current problem or enquiry. This is what drives the use of search engines forward, as web users attempt to specify and work through results lists in an effort to state precisely what it is they are looking for. This is what drives the users of community and hobby groups on Yahoo! Groups and other discussion boards to pose increasingly detailed statements of exactly what it is they are trying to learn.

Placing relevant content in to exactly the right context at the right time is an art. It involves both aspects of effective content design and aspects of dynamic search and placement. Information needs are not static - they will change with both the situation and the changing capacities of the learner. Placement depends on the precise nature of the request sent by a piece of software or tool, and the ability of a piece of content to respond to that success. Game designers understand this - the game presents different information to users at different points of the game where it will be useful - and usable - by the player.

4. Network Learning

By 'network learning' we mean the principles that inform the development of new connections in a network - in other words, how networks learn. These principles are informed partially through the study of neuroscience and partially throught he development of networks in computer science, an approach called 'connectionism'. (See http://plato.stanford.edu/entries/connectionism/)

Though there are various ways networks can form sets of connections among entities, there are three major types of network learning that are informative in this discussion:

- Simple (or 'Hebbian') associationism - this is the principle that if two nodes in a network are activated at the same time, a connection will form between those nodes. Thus, for example, we recognize similar things (like tigers) by seeing them over and over again. (See http://psychclassics.yorku.ca/Hebb/)

- Backpropagation - this is the principle that allows the output of a network to be corrected by the sending of a signal back through the network instucting it to either strengthen or weaken the connections that produced the output. For example, a person might receive feedback - positive or negative - on their performance. (See http://www.seattlerobotics.org/encoder/nov98/neural.html)

- Boltzmann - this is a principle that allows connections to strengthen or weaken by 'settling' into thermodynamically stable configurations (much the way water will settle to a level surface in a pond), and a mechanism (called 'annealing') that disrupts the network of connections, to prompt the settlement into the most stable configuration possible. (See Hinton, G. E. and Sejnowski, T. J. Learning and relearning in Boltzmann machines. In Rumelhart, D. E. and McClelland, J. L., editors, Parallel Distributed Processing: Explorations in the Microstructure of Cognition.)

Most people don't think of themselves as associating, back propagating of settling. But the theory of learning described by these mechanisms is in fact relatively commonplace, and can be described (in slogan form) as follows:

To *teach* is to model and demonstrate, and to "learn* is to practice and reflect. To teach is, essentially, to provide or to make possible the having of experiences by students. These models and demonstrations, by virtue of their structural similarities with other models and demonstrations, allow students to form relevant networks of connections. Students then actively begin to learn by practicing - first by imitating, then later by creating something novel. The point of practize is to improve performance by receiving feedback. They then reflect on what they have experience and practiced - this is (somewhat) analagous to the Boltzmann mechanism.

5. Reliability

Both personal learning and network learning are characterized by dynamic patterns of interactivity in a networked environment. The same principles are at work in each case. But can this process be trusted? Is it reliable?

Networks can be trusted, as James Surowiecki shows in The Wisdom of Crowds (see http://www.randomhouse.com/features/wisdomofcrowds/). "Many cognitive, coordination and cooperation problems are best solved by canvassing groups (the larger the better) of reasonably informed, unbiased, engaged people. The group's answer is almost invariably much better than any individual expert's answer, even better than the best answer of the experts in the group." (See http://blogs.salon.com/0002007/2004/11/15.html) It is this wisdom we see not only in the sudience picking the right answer in "Who Wants to Be a Millionaire" but also in picking stocks in the stock market and picking governments in elections.

However, not just any network can be trusted. Networks can sometimes run away with themselves - for example, if one person in a community catches a fatal virus, it can spread to every other member, and kill the entire community. Such phenomena are known as cascade phenomena. In the realm of information networks (such as the brain, or a community) these are known as informational cascades. (See http://en.wikipedia.org/wiki/Informational_cascade) They are like 'jumping to a comclusion' or 'groupthink'.

Networks avoid informational cascades - and hence, are reliable - only if they satisfy the following four criteria (known collectively as 'the semantic condition'):

- Diversity - Did the process involve the widest possible spectrum of points of view? Did people who interpret the matter one way, and from one set of background assumptions, interact with people who approach the matter from a different perspective?

- Autonomy - Were the individual knowers contributing to the interaction of their own accord, according to their own knowledge, values and decisions, or were they acting at the behest of some external agency seeking to magnify a certain point of view through quantity rather than reason and reflection?

- Openness - Is there a mechanism that allows a given perspective to be entered into the system, to be heard and interacted with by others?

- Connectivity - Is the knowledge being produced the product of an interaction between the members, or is it a (mere) aggregation of the members' perspectives? A different type of knowledge is produced one way as opposed to the other. Just as the human mind does not determine what is seen in front of it by merely counting pixels, nor either does a process intended to create public knowledge.

6. Examples

How does the discussion above help us understand about and design learning technologies? They show us not only what to design but also help us understand what would be a better (or worse) design.

We begin with the principle, 'To *teach* is to model and demonstrate, and to "learn* is to practice and reflect.' This gives us a set of four types of things to create:

- Things that model - such as the wiki, concept maps, diagram tools such as gliffy, video / 2L 3D representation, and the like

- Things that demonstrate - such as code libraries, image samples, articles describing thought processes (see http://www.jerrypournelle.com/), case studies and stories

- Things that help us practice - such as games, sandboxes, job aides, simulations and ebnvrionments

- Things that help us reflect - such as presentations and seminars, blogs, wikis, discussion groups, and other ways of sharing and communicating

For any given application in each of the four categories, we can apply the remaining principles to provide an assessment of it likely effectiveness.

For example, consider the wiki. Does it support network learning? Yes - it provides examples to follow, allows correction and criticism, and rethinking and rewriting. Does it support personal learning? Yes, it engages interaction. It supports a genuine voice, experiences, opinions. It is a simple and consistent interface. It is (mostly) accessible where and when I need it.

Is the wiki reliable? Do I have diversity of sources? Yes - but only if there is a threshold number of users. Are the sources autonomous? They can be. And wikis support connectedness with links, etc, and can be open to a large number of contributors. These considerations argues against closed or private wikis, but suggest that wikis can be useful for large groups.

As another example, consider image libraries. They provide examples to follow, but our study suggests that image libraries should have (like Flickr (see http://www.flickr.com) communictaion channels, ratings and reviews, and ways to link images, such as tags. And an image library will be 'reliable' if it allows contributions from numerous photographers. We also see that we want people to have individual identities on Flickr, rather than just contributing to a pool, to preserve autonomy and diversity.

As a third example, consider Second Life. We can see why people are attracted to it. It allows us to create examples to follow, corrections and criticisms. It engages interaction and supports a genuine voice. But we also see weaknesses. Is Second Life a good place for reflection? There are limits on reusing what other people have created. It is also semantically weak. There is only one world, not a large number of diverse worlds. Autonomy is limited - you can't even pick your own name - and there are questions about governance. There is connectedness, through slurls, but it is not clear that it is an open platform.

7. Concluding Remarks

The purpose of this paper was to describe how network learning works and to show how an understanding of network learning can inform the design and evaluation of online learning applications.

Admittedly, there is room for debate and discussion regarding the nature and precise statement of the principles. What remains, however, is that the model of learning as a personal and a network activity provides us with concrete insights into the sort

of learning environments that are most likely to be successful online.