


Access :: Future

Practical Advice on How to Learn and What to Learn



Stephen Downes
National Research Council of Canada

Access :: Future
Practical Advice on How to Learn and What to Learn

Stephen Downes
Copyright (c) 2011

This work is published under a Creative Commons License



Attribution-NonCommercial-ShareAlike
CC BY-NC-SA

This license lets others remix, tweak, and build upon your work non-commercially, as long as they credit you and license their new creations under the identical terms.

[View License Deed](#) | [View Legal Code](#)

Learning Today

Introduction	5
The Purpose of Learning	7
Things You Really Need To Learn	9
The Mark of Wisdom	19
Critical Thinking in the Classroom.....	23
Necessary and Sufficient Conditions.....	36
Not All... ..	42
Educational Blogging.....	43
Your Career	74
Managing Your Blog Entry: 11 Better Tips.....	75
Blogs in Education.....	79
How To Write Articles Quickly and Expertly	82
Principles for Evaluating Websites.....	91
Applying Critical Reasoning.....	101
How Do You Know?.....	103
Having Reasons	114
How Memory Works.....	118
How The Net Works.....	126
An Operating System For The Mind.....	134
Personal Knowledge: Transmission or Induction?	144
Virtues Education.....	150
Free Learning and Control Learning.....	155
The Science of Learning	169
E-Learning 2.0	177
To The School or Classroom 2.0 Advocates	185
The Issues in Front of Us	189
The Form of Informal	195
Uniqueness and Conformity	201
New Technology Supporting Informal Learning	203
How I Would Organize A Conference	217

What I Learned in High School..... 219

My Personal Passion Trajectory..... 223

My Academic Upbringing..... 228

Social Media and Me..... 230

Seven Habits of Highly Connected People..... 234

The Reality of Virtual Learning..... 238

Nine Rules for Good Technology 257

What Not To Build..... 262

Ten Futures 272

Introduction

In the preparation of a volume such as this there is always for me the sense that I am giving up on the content. The flaws in the material are there for all to see, and while no doubt a few years of additional research and writing would immeasurably improve the work, other matters and other issues press me for time, until, inevitably, there is nothing left to do but to tie the collection together and send it forth.

Some of my most popular work is in this collection. From the outset, what I sought to do was to combine some of my most useful work for learners – my guides to learning and blogging, for example – with my less formal works on education. The effort here is to present a case for a type of learning by producing a volume that is also an instance of that learning, and that has as its subject the content of that learning.

It is hard – or should be hard – to tell where the material addressed to learners ends and where the material addressed to teachers begins. The subject of this work is, first and foremost, thinking and learning, and it is this subject that ought to interest most both teachers and learners. If we can see that learning and reasoning are part and parcel of the same activity, and that it *is* an activity, then, i think, the most elusive aspect of education has been overcome, and learning can begin.

This is not an academic work. Goodness knows the world needs one, if only to address these issues to academics, who might not read them otherwise, but I am not the person to write such a work. For me, writing – far from the finely honed craft of the technician and the artisan – is most akin to breathing, something I do in the course of every day, without too much regard to whether I get it right, the purpose being primarily to interact with the world rather than erect some edifice to seek homage from it. This work is fast, it is loose, it works at and harries themes rather than pronouncing on them, it is colloquial, it is rooted in practice and pragmatism far more than theory, and it has no doubt all been written before by someone who should receive far more credit than I.

The most praising and damning thing that can be said of this work is that it is composed of blog posts. That is not strictly true – there are some transcripts from some of my talks as well – but it's close enough to true to be worth mentioning. It can be argued, on the one hand, that the blog post lacks the careful reflection and precision of an academic work. On the other hand, from my perspective, such reflection and precision are fictions in and of themselves. No less care is taken in the writing of these posts, but it is a care of a different sort, a care that attends much more to genuine voice and authentic experience than to methodological rigor and (pseudo-)scientific objectivity.

Taken together, these posts represent a treatise, an approach to thinking and learning, and if they lack a careful definition statement, summarization of canonical literature, and body of evidence in test subjects and student surveys, that is part of the thesis, that a rough-and-ready critical stance and practical application in a sea of conflicting data and changing technologies *is* the best, most appropriate, and only route to an understanding of this (or any other) science.

These are not carefully prepared treatises authored by me *about* teaching and learning, they are examples of me, working with the best tools I have, *engaged* in teaching and learning. I am not arguing or explaining, I am not describing, I am *demonstrating*. You are not intended to read these articles with the idea that they are objective statements of fact, but as an explorer, being told of someone else's experiences.

The Purpose of Learning

For me, an education was not a given. Yes, I was born and raised in one of the richest nations in the world, a country where schooling is not merely available, but required, and yet my education was still not a given. I rebelled early, and then had to scratch and claw my way through four high schools and three colleges before finally getting a degree.

Would I do it again? Absolutely. The same way? I'm not sure - it would depend on what options are available. I didn't have many choices. I took what I could get. I paid for it with promises and IOUs. And I never did quite finish my PhD. Almost everything the education system stood for, I opposed. And in many ways, I made my own education, spending at least as much time learning outside the formal system as within.

It's ironic that what ultimately led me away from my studies was the experience of standing in front of students, mostly adult and disadvantaged, teaching in northern Alberta. This wasn't planned; I didn't set out to 'do community work' or any such thing; that's a luxury allowed people who had more financial freedom than I. But it was immensely rewarding, not the least because I could see my face among the students in those classes, and I knew exactly what I was trying to provide for them.

It's hard to state what that is without becoming a bit hackneyed, but there's truth in every cliché. With the basic tools of literacy, critical literacy and reasoning skills, combined with a whole dose of self-confidence, these students had at least a chance to make something of their lives, to shape their own futures, to be something more than flotsam in the currents of social change and disorder.

It's no magic pill, and it's altogether too little for both those people who have to struggle out of poverty just to get their foot in the door, and those born out of affluence who have no comprehension of the work required to become a person of strong character and self-determination. Yet in the right meter, and in combination with the right experiences, an education is sufficient to lift a person into a life of self-awareness and reflection. It is the great liberator, and even should an educated person never rise out of poverty, that person will never again be poor.

John Stuart Mill said that the principle of liberty is the right of each person to pursue their own good, in their own way. But he never intended this right to be given only to a nation of sheep, and he understood that the highest principle of liberty was in fact both the right and capacity to actually define one's good, to freely choose one's ambition and purpose in life, and to enact the means and mechanisms to carry it out. Freedom is not merely the absence of restraint, but the right to live meaningfully.

An educated population is probably the least governable, the most likely to rebel, the most stubborn and the most critical. But it is a population capable of the most extraordinary things, because each person strides purposefully forward, and of their own volition, together, they seek a common destiny.

Submitted as my contribution to [Purpos/ed](#)

Moncton, February 2, 2011

Things You Really Need To Learn

Guy Kawasaki last week [wrote](#) an item describing 'ten things you should learn this school year' in which readers were advised to learn how to write five sentence emails, create powerpoint slides, and survive boring meetings. It was, to my view, advice on how to be a business toady. My view is that people are worth more than that, that pleasing your boss should be the least of your concerns, and that genuine learning means something more than how to succeed in a business environment.

But what should you learn? Your school will try to teach you facts, which you'll need to pass the test but which are otherwise useless. In passing you may learn some useful skills, like literacy, which you should cultivate. But Guy Kawasaki is right in at least this: schools won't teach you the things you *really* need to learn in order to be successful, either in business (whether or not you choose to live life as a toady) or in life.

Here, then, is my list. This is, in my view, what you need to learn in order to be successful. Moreover, it is something you can start to learn *this year*, no matter what grade you're in, no matter how old you are. I could obviously write much more on each of these topics. But take this as a starting point, follow the suggestions, and learn the rest for yourself. And to educators, I ask, if you are not teaching these things in your classes, why are you not?

1. How to predict consequences

The most common utterance at the scene of a disaster is, "I never thought..." The fact is, most people are very bad at predicting consequences, and schools never seem to think to teach them how to improve.

The prediction of consequences is part science, part mathematics, and part visualization. It is essentially the ability to create a mental model imaging the sequence of events that would follow, "what would likely happen if...?"

The danger in such situations is focusing on what you want to happen rather than what might happen instead. When preparing to jump across a gap, for example, you may visualize yourself landing on the other side. This is good; it leads to successful jumping. But you need also to visualize *not* landing on the other side. What would happen then? Have you even *contemplated* the likely outcome of a 40 meter fall?

This is where the math and science come in. You need to compare the current situation with your past experience and calculate the probabilities of different outcomes. If, for example, you are looking at a 5 meter gap, you should be asking, "How many times have I successfully jumped 5 meters? How many

times have I failed?" If you don't know, you should know enough to attempt a test jump over level ground.

People don't think ahead. But while you are in school, you should always be taking the opportunity to ask yourself, "what will happen next?" Watch situations and interactions unfold in the environment around you and try to predict the outcome. Write down or blog your predictions. With practice, you will become expert at predicting consequences.

Even more interestingly, over time, you will begin to observe patterns and generalities, things that make consequences even easier to predict. Things fall, for example. Glass breaks. People get mad when you insult them. Hot things will be dropped. Dogs sometimes bite. The bus (or train) is sometimes late. These sorts of generalizations - often known as 'common sense' - will help you avoid unexpected, and sometimes damaging, consequences.

2. How to read

Oddly, by this I do not mean 'literacy' in the traditional sense, but rather, how to look at some text and to *understand*, in a deep way, what is being asserted (this also applies to audio and video, but grounding yourself in text will transfer relatively easily, if incompletely, to other domains).

The four major types of writing are: description, argument, explanation and definition. [I have written about these elsewhere](#). You should learn to recognize these different types of writing by learning to watch for *indicators* or keywords.

Then, you should learn how sentences are joined together to form these types of writing. For example, an argument will have two major parts, a premise and a conclusion. The conclusion is the point the author is trying to make, and it should be identified with an indicator (such as the words 'therefore', 'so', or 'consequently', for example).

A lot of writing is fill - wasted words intended to make the author look good, to distract your attention, or to simply fill more space. Being able to cut through the crap and get straight to what is actually being said, without being distracted, is an important skill.

Though your school will never teach you this, find a basic book on informal logic (it will have a title like '[critical thinking](#)' or something like that). Look in the book for argument forms and indicator words (most of these books don't cover the other three types of writing) and practice spotting these words in text and in what the teacher says in class. Every day, focus on a specific indicator word and watch how it is used in practice.

3. How to distinguish truth from fiction

I have written extensively on this elsewhere, nonetheless, this remains an area schools to a large degree ignore. Sometimes I suspect it is because teachers feel their students must absorb knowledge uncritically; if they are questioning everything the teacher says they'll never learn!

The first thing to learn is to actually question what you are told, what you read, and what you see on television. Do not simply accept what you are told. Always ask, how can you know that this is true? What evidence would lead you to believe that it is false?

I have written several things to help you with this, including my [Guide to the Logical Fallacies](#), and my article on [How to Evaluate Websites](#). These principles are more widely applicable. For example, when your boss says something to you, apply the same test. You may be surprised at how much your boss says to you that is simply not true!

Every day, subject at least one piece of information (a newspaper column, a blog post, a classroom lecture) to thorough scrutiny. Analyze each sentence, analyze every word, and ask yourself what you are expected to believe and how you are expected to feel. Then ask whether you have sufficient reason to believe and feel this way, or whether you are being manipulated.

4. How to empathize

Most people live in their own world, and for the most part, that's OK. But it is important to at least recognize that there are other people, and that they live in their own world as well. This will save you from the error of assuming that everyone else is like you. And even more importantly, this will allow other people to become a surprising source of new knowledge and insight.

Part of this process involves seeing things through someone else's eyes. A person may be, quite literally, in a different place. They might not see what you see, and may have seen things you didn't see. Being able to understand how this change in perspective may change what they believe is important.

But even more significantly, you need to be able to imagine how other people feel. This means that you have to create a mental model of the other person's thoughts and feelings in your own mind, and to place yourself in that model. This is best done by imagining that you *are* the other person, and then placing yourself into a situation.

Probably the best way to learn how to do this is to study drama (by that I don't mean studying Shakespeare, I mean learning how to act in plays). Sadly, schools don't include this as part of the core curriculum. So instead, you will need to study subjects like religion and psychology. Schools don't really include these either. So make sure you spend at least some time in different role-playing games (RPGs) every day and practice being someone else, with different beliefs and motivations.

When you are empathetic you will begin to seek out and understand ways that help bridge the gap

between you and other people. Being polite and considerate, for example, will become more important to you. You will be able to *feel* someone's hurt if you are rude to them. In the same way, it will become more important to be honest, because you will begin to see how transparent your lies are, and how offensive it feels to be thought of as someone who is that easily fooled.

Empathy isn't some sort of bargain. It isn't the application of the [Golden Rule](#). It is a *genuine* feeling in yourself that operates in synch with the other person, a way of accessing their inner mental states through the sympathetic operation of your own mental states. You are polite because *you* feel bad when you are rude; you are honest because *you* feel offended when you lie.

You need to *learn* how to have this feeling, but once you have it, you will understand how empty your life was before you had it.

5. How to be creative

Everybody can be creative, and if you look at your own life you will discover that you are already creative in numerous ways. Humans have a natural capacity to be creative - that's how our minds work - and with practice can become very good at it.

The trick is to understand how creativity works. Sometimes people think that creative ideas spring out of nothing (like the proverbial 'blank page' staring back at the writer) but creativity is in fact the result of using and manipulating your knowledge in certain ways.

Genuine creativity is almost always a response to something. This article, for example, was written in response to an article on the same subject that I thought was not well thought out. Creativity also arises in response to a specific problem: how to rescue a cat, how to cross a gap, how to hang laundry. So, in order to be creative, the first thing to do is to learn to look for problems to solve, things that merit a response, needs that need to be filled. This takes practice (try writing it down, or blogging it, every time you see a problem or need).

In addition, creativity involves a transfer of knowledge from one domain to another domain, and sometimes a manipulation of that knowledge. When you see a gap in real life, how did you cross a similar gap in an online game? Or, if you need to clean up battery acid, how did you get rid of excess acid in your stomach?

Creativity, in other words, often operates by *metaphor*, which means you need to learn how to *find things in common* between the current situation and other things you know. This is what is typically meant by 'thinking outside the box' - you want to go to *outside the domain* of the current problem. And the particular skill involved is *pattern recognition*. This skill is hard to learn, and requires a lot of practice, which is why creativity is hard.

But pattern recognition *can* be learned - it's what you are doing when you say one song is similar to another, or when you are taking photographs of, say, flowers or fishing boats. The arts very often involve finding patterns in things, which is why, this year, you should devote some time every day to an art - music, photography, video, drawing, painting or poetry.

6. How to communicate clearly

Communicating clearly is most of all a matter of knowing what you want to say, and then employing some simple tools in order to say it. Probably the hardest part of this is knowing what you want to say. But it is better to spend time being sure you understand what you mean than to write a bunch of stuff trying to make it more or less clear.

Knowing what to say is often a matter of structure. Professional writers employ a small set of fairly standard structures. For example, some writers prefer articles (or even whole books!) consisting of a list of points, like this article. Another structure, often called 'pyramid style', is employed by journalists - the entire story is told in the first paragraph, and each paragraph thereafter offers less and less important details.

Inside this overall structure, writers provide arguments, explanations, descriptions or definitions, sometimes in combination. Each of these has a distinctive structure. An argument, for example, will have a conclusion, a point the writer wants you to believe. The conclusion will be supported by a set of premises. Linking the premises and the conclusion will be a set of *indicators*. The word 'therefore', for example, points to the conclusion.

Learning to write clearly is a matter of learning about the tools, and then practice in their application. Probably the best way to learn how to structure your writing is to learn how to give speeches without notes. This will force you to employ a clear structure (one you can *remember!*) and to keep it straightforward. I have [written more on this](#), and also, check out Keith Spicer's book, [Winging It](#).

Additionally, master the tools the professionals use. Learn the structure of arguments, explanations, descriptions and definitions. Learn the indicator words used to help readers navigate those structures. Master basic grammar, so your sentences are unambiguous. Information on all of these can be found online.

Then practice your writing every day. A good way to practice is to join a student or volunteer newspaper - writing with a team, for an audience, against a deadline. It will force you to work more quickly, which is useful, because it is faster to write clearly than to write poorly. If no newspaper exists, create one, or start up a news blog.

7. How to Learn

Your brain [consists of billions of neural cells](#) that are connected to each other. To learn is essentially to form sets of those connections. Your brain is always learning, whether you are studying mathematics or staring at the sky, because these connections are always forming. The difference in what you learn lies in how you learn.

When you learn, you are trying to create *patterns of connectivity* in your brain. You are trying to connect neurons together, and to strengthen that connection. This is accomplished by repeating sets of behaviours or experiences. Learning is a matter of practice and repetition.

Thus, when learning anything - from '2+2=4' to the principles of quantum mechanics - you need to repeat it over and over, in order to grow this neural connection. Sometimes people learn by repeating the words aloud - this form of rote learning was popular not so long ago. Taking notes when someone talks is also good, because you hear it once, and then repeat it when you write it down.

Think about learning how to throw a baseball. Someone can explain everything about it, and you can understand all of that, but you still have to throw the ball several thousand times before you get good at it. You have to grow your neural connections in just the same way you grow your muscles.

Some people think of learning as remembering sets of facts. It can be that, sometimes, but learning is more like *recognition* than remembering. Because you are trying to build networks of neural cells, it is better to learn a *connected whole* rather than unconnected parts, where the connected whole you are learning in one domain has the *same pattern* as a connected whole you already know in another domain. Learning in one domain, then, becomes a matter of recognizing that pattern.

Sometimes the patterns we use are very artificial, as in '[every good boy deserves fudge](#)' (the sentence helps us remember musical notes). In other cases, and more usefully, the pattern is related to the laws of nature, logical or mathematical principles, the flow of history, how something works as a whole, or something like that. Drawing pictures often helps people find patterns (which is why mind-maps and [concept maps](#) are popular).

Indeed, you should view the study of mathematics, history, science and mechanics as the study of [archetypes](#), basic patterns that you will recognize over and over. But this means that, when you study these disciplines, you should be asking, "what is the pattern" (and not merely "what are the facts"). And asking this question will actually make these disciplines easier to learn.

Learning to learn is the same as learning anything else. It takes practice. You should try to learn something every day - a random word in the dictionary, or a random Wikipedia entry. When learning this item, do not simply learn it in isolation, but look for patterns - does it fit into a pattern you already know? Is it a type of thing you have seen before? Embed this word or concept into your existing knowledge by using it in some way - write a blog post containing it, or draw a picture explaining it.

Think, always, about *how* you are learning and *what* you are learning at any given moment. Remember, you are *always* learning - which means you need to ask, what are you learning when you are watching television, going shopping, driving the car, playing baseball? What sorts of patterns are being created? What sorts of patterns are being reinforced? How can you *take control* of this process?

8. How to stay healthy

As a matter of practical consideration, the maintenance of your health involves two major components: minimizing exposure to disease or toxins, and maintenance of the physical body.

Minimizing exposure to disease and toxins is mostly a matter of cleanliness and order. Simple things - like keeping the wood alcohol in the garage, and not the kitchen cupboard - minimize the risk of accidental poisoning. Cleaning cooking surfaces and cooking food completely reduces the risk of bacterial contamination. Washing your hands regularly prevents transmission of human borne viruses and diseases.

In a similar manner, some of the hot-button issues in education today are essentially issues about how to warn against exposure to diseases and toxins. In a nutshell: if you have physical intercourse with another person you are facilitating the transmission of disease, so wear protection. Activities such as drinking, eating fatty foods, smoking, and taking drugs are essentially the introduction of toxins into your system, so do it in moderation, and where the toxins are significant, don't do it at all.

Personal maintenance is probably even more important, as the major threats to health are generally those related to physical deterioration. The subjects of proper nutrition and proper exercise should be learned and practiced. Even if you do not become a health freak (and who does?) it is nonetheless useful to know what foods and types of actions are beneficial, and to create a habit of eating good foods and practicing beneficial actions.

Every day, seek to be active in some way - cycle to work or school, walk a mile, play a sport, or exercise. In addition, every day, seek to eat at least one meal that is 'good for you', that consists of protein and minerals (like meat and vegetables, or soy and fruit). If your school is not facilitating proper exercise and nutrition, demand them! You can't learn anything if you're sick and hungry! Otherwise, seek to establish an alternative program of your own, to be employed at noontime.

Finally, remember: you never have to justify protecting your own life and health. If you do not want to do something because you think it is unsafe, then it is your absolute right to refuse to do it. The consequences - *any* consequences - are better than giving in on this.

9. How to value yourself

It is perhaps cynical to say that society is a giant conspiracy to get you to feel badly about yourself, but it

wouldn't be completely inaccurate either. Advertisers make you feel badly so you'll buy their product, politicians make you feel incapable so you'll depend on their policies and programs, even your friends and acquaintances may seek to make you doubt yourself in order to seek an edge in a competition.

You can have all the knowledge and skills in the world, but they are meaningless if you do not feel personally empowered to use them; it's like owning a [Lamborghini](#) and not having a driver's license. It looks shiny in the driveway, but you're not really getting any value out of it unless you take it out for a spin.

Valuing yourself is partially a matter of personal development, and partially a matter of choice. In order to value yourself, you need to feel you are worth valuing. In fact, you *are* worth valuing, but it often helps to prove it to yourself by attaining some objective, learning some skill, or earning some distinction. And in order to value yourself, you have to say "I am valuable."

This is an important point. How we think about ourselves is as much a matter of learning as anything else. If somebody tells you that you are worthless over and over, and if you do nothing to counteract that, then you will come to believe you are worthless, because *that's* how your neural connections will form. But if you repeat, and believe, and behave in such a way as to say to yourself over and over, *I am valuable*, then *that's* what you will come to believe.

What is it to value yourself? It's actually many things. For example, it's the belief that you are *good enough* to have an opinion, have a voice, and have a say, that your contributions *do* matter. It's the belief that you are *capable*, that you can learn to do new things and to be creative. It is your ability to be *independent*, and to not rely on some particular person or institution for personal well-being, and *autonomous*, capable of making your own decisions and living your life in your own way.

All of these things are yours by right. But they will never be given to you. You have to *take* them, by actually believing in yourself (no matter what anyone says) and by actually being autonomous.

Your school doesn't have a class in this (and may even be actively trying to undermine your autonomy and self-esteem; watch out for this). So you have to take charge of your own sense of self-worth.

Do it every day. Tell yourself that you are smart, you are cool, you are strong, you are good, and whatever else you want to be. Say it out loud, in the morning - hidden in the noise of the shower, if need be, but *say* it. Then, *practice* these attributes. *Be* smart by (say) solving a crossword puzzle. *Be* cool by making your own fashion statement. *Be* strong by doing something you said to yourself you were going to do. *Be* good by doing a good deed. And every time you do it, remind yourself that you have, in fact, done it.

10. How to live meaningfully

This is probably the hardest thing of all to learn, and the least taught.

Living meaningfully is actually a combination of several things. It is, in one sense, your dedication to some purpose or goal. But it is also your sense of appreciation and dedication to the here and now. And finally, it is the realization that your place in the world, your meaningfulness, is something you must create for yourself.

Too many people live for no reason at all. They seek to make more and more money, or they seek to make themselves famous, or to become powerful, and whether or not they attain these objectives, they find their lives empty and meaningless. This is because they have confused means and ends - money, fame and power are things people seek in order to do what is *worth doing*.

What is worth doing? That is up to you to decide. I have chosen to dedicate my life to helping people obtain an education. Others seek to cure diseases, to explore space, to worship God, to raise a family, to design cars, or to attain enlightenment.

If you don't decide what is worth doing, someone will decide for you, and at some point in your life you will realize that you haven't done what is worth doing at all. So spend some time, today, thinking about what is worth doing. You can change your mind tomorrow. But begin, at least, to guide yourself *somewhere*.

The second thing is sometimes thought of as 'living in the moment'. It is essentially an understanding that *you control your thoughts*. Your thoughts have no power over you; the only thing that matters at all is this present moment. If you think about something - some hope, some failure, some fear - that thought cannot hurt you, and *you* choose how much or how little to trust that thought.

Another aspect of this is the following: what you are doing *right now* is the thing that you *most want to do*. Now you may be thinking, "No way! I'd rather be on Malibu Beach!" But if you *really* wanted to be on Malibu Beach, you'd be there. The reason you are not is because you have *chosen* other priorities in your life - to your family, to your job, to your country.

When you realize you have the power to choose what you are doing, you realize you have the power to choose the consequences. Which means that consequences - even bad consequences - are for the most part a matter of *choice*.

That said, this understanding is very liberating. Think about it, as a reader - what it means is that what I *most wanted* to do with my time right now is to write this article so that you - yes, *you* - would read it. And even more amazingly, I know, as a writer, that the thing you most want to do right now, even more than you want to be in Malibu, is to *read my words*. It makes me want to write something meaningful - and it gives me a way to put meaning into my life.

Moncton, August 30, 2006

The Mark of Wisdom

Posted to ITForum

I think that wisdom is the property of the wise, something a wise person develops for him or her self over time, and therefore not the subject of an external standard. What makes a person wise is not the having of a certain opinion, or even the taking of a certain perspective or point of view. To be wise is to have mastered and to have used successfully for a certain period of time the practices that lead to wisdom. The question of whether someone has become wise (and not, note, 'acquired wisdom') is indeterminate in its resolution. Certainly, there is no acid test that identifies the wise (though we can all think of numerous examples we think would demonstrate that someone is unwise).

As Clark Quinn [mentioned in his paper](#), I have listed the practices that I think lead to wisdom in a [column on my blog](#). It should be noted on reading that each of the ten items describes a skill. It describes 'how' to do something. This is deliberate. The outcome of the application of any such skill must be left to the learner. For otherwise, acquisition of the outcome, rather than practice of the skill, becomes definitive of wisdom, and therefore the tendency will be to focus on the outcome at the expense of the skill, at the cost of never learning the skill.

For example, one of the skills I describe is 'How to predict consequences'. The statement of this principle could be questioned - Jonathon Richter quite rightly argues that the skill needs to apply not only to simple Newtonian systems with determinate outcomes but also to complex and chaotic environments in which the prediction is of only a probability or a range of outcomes. But it would be a mistake to say that the acquisition of this skill entails believing, say, that good intentions lead to good actions, or believing, say, that continued inaction on global warming will lead to catastrophe.

To make the test of the skill the test of a believe is to measure, not the having of the skill, but rather, the adherence to orthodoxy. Not everybody predicts the same effect from the same causes: that this can be the case, and that each of the two predictions can be substantiated is evidence, not of predictive failure, but rather, of the chaotic world Richter alludes to.

In a similar manner, some of the writers to this list have characterized 'humility' as one of the attributes of wisdom. Jan Visser characterizes humility as "the awareness of who we are, of our place in the universe, of what lies beyond our own ephemeral existence." Thus characterized, however, humility is characterized not as a practice or a skill but rather as an awareness that certain things are the case. But this test, then, is of what someone believes rather than how they came to believe it, and hence, is not a test of wisdom but of orthodoxy.

It is possible to be wise without humility. To be honest, I do not know my own place in the universe (and routinely disregard the admonishment "that's not your place"). I do not know what lies outside my own

existence, if anything. I consider the following two statements to be equally likely: "I am in the universe" and "the universe is in me". From which it follows that there is a great deal I do not know about myself. Being wise may lie in accepting any of these dilemmas, or in opting, on the basis of faith or intuition, one statement or the other.

Clark Quinn's own argument takes the perspective that wisdom results from practice. "Wisdom is making decisions on a systemic basis..." he writes. And that "it's very much a journey, not a destination." But in so much of what he writes the outcome, rather than the process, does the heavy lifting. The sentence quoted above reads, in full, "wisdom is making decisions on a systemic basis that are in line with our [interests] in the long-term as well as the immediate moment, and in line with our values for not only ourselves but others and society and the world as a whole." I have had to insert the word 'interests' as it is missing in the original, though it could read 'needs' or even 'fashions' at still make the same point.

And the point is: wisdom is characterized by having a certain type of perspective, of using a certain metric. But it seems to me that this should be an outcome of wisdom, and not definitive of it. Is being wise tantamount to enlightened self-interest? This is what the sentence seems to imply. But being wise may equally entail disregard for one's own self-interest, in rising above one's own self-interest. Certainly some Buddhists are wise, but central to the philosophy of Buddhism is what one might characterize as "the cessation of craving" or the cessation of clinging - taking one's own interest and seeing it as the cause of pain, of [Dukkha](#).

As with humility and self-interest, the matter of values is also one of content over process. Quinn writes that the wise person makes decisions, as noted above, "in line with our values for not only ourselves but others and society and the world as a whole." One could hardly find a more compelling example of a call for orthodoxy! He adds, "We'd need to discuss values and deliberately choose a value system to embody. Whichever one we choose (and this is difficult subject all on its own), we'll want to make it explicit."

Certainly there may be arguments for and against a values-based education, just as (say) there may be arguments for or against values-based government, or values-based religion, or values-based economic systems. But to make basing one's reasoning in values definitive of wisdom seems very much to be over-shooting the mark.

Why would I say this? For after all, people familiar with my own work will be well aware of a set of values that permeates it through and through. Yet deeper study will show that I well regard my own work as deeply situated within a certain society, a certain context, and that the values I espouse recognize this, and hence, do not presume to pass judgment on how the other person elects to live his or her own life. And it is a part of this belief to recognize that another person may not live his or her life according to values at all, and yet may nonetheless be wise.

Consider what it is to live one's life or to make decisions according to values. Quinn helpfully quotes Gladwell: "One of humankind's biggest problems in decision-making is assigning the wrong weights to the variables.... If I have an ethical system, I have a way of assigning those weights" - Malcolm Gladwell. Having values is depicted as - and indeed, arguably is - nothing more than a system of assigning weights to variables, to possibilities, to options. The simplest of value systems is to assign to all things a binary value: right or wrong, good or bad. More sophisticated approaches assign weightings of good and bad, and allow for an interplay between the entities so judged.

How completely the mercantile philosophy has permeated our way of life, that we cannot even imagine the concept of wisdom without also imagining some way of applying weights and measures to its application! Are the philosophies of epicurianism, of hedonism, wrong not because they are bad values but because they deal in quality of experience, quality of life, rather than with the calculation and measurement of what is right and wrong? Lao Tzu says that the attachments of values to things are nothing more than labels, nothing more than signs, and that wisdom consists in recognizing that there is always another way of representing the same reality. Friedrich Nietzsche describes, in *Beyond Good and Evil*, the transvaluation of value. J.L. Mackie is explicit, in *Ethics Without God*, that we are "inventing right and wrong".

The appeal of values, as are the appeal of natural laws and the idea of essences, is the appeal of the universal, the idea that there is some means though which we may abbreviate the complexities of our lives and our universe through appeal to some sort of underlying principle, whether that principle represents the Nature of the world, or whether it represents nothing more than an abbreviation of our experience and beliefs. But this is only one way of representing the world or of determining our course of action. Numerous alternatives exist.

Consider, for example, what you would do were you confronted by a strange creature, one that you had never seen before, consisting of an animal with brightly coloured fur, with sharp teeth and savage claws, and a roar that sent chills down your spine. Would wisdom at this point consist in the application of some set of values? Is there some principle of measurement to which you would appeal in this case? Almost certainly not. We most certainly have not formulated any rules describing what to do in such a case - certainly we do know what to do (run!) but any drafting of a rule would happen after the fact. Our actions would be determined by recognition and similarity. It looks like a tiger, best treat it as though it were a tiger.

If we consider our own actions, even for a moment, it becomes clear how many of them are not determined by any sort of value at all (and which are, nonetheless, wise). We breathe, our hearts beat, our blood circulates, not because we will it to be the case, not because it is good or right that it be the case, but merely because we are the sort of creatures that live and breathe. And surely part of being wise lies in being what you are, rather than what you are not? This is why we thought the characters of 'Flatliners' were unwise.

Others of our decisions are governed by what I would call 'network phenomena'. One of the principles of a successful network is 'diversity' - diverse networks are more reliable than those that are not. Fostering diversity, however, means fostering instances of actions and entities that precisely do not adhere to orthodox values. We've all heard the expression, "The exception proves the rule." We recognize that 'generosity is good' (say) by observing (or telling tales) of Scrooge-like people who are not generous, and observing their fate. There may not have been a King Midas turned to gold, but we are all aware from observation the poverty of greed.

What is important here is not that the value was discovered, nor the specific content of the value, but rather, that a process was followed that would allow this learning to take place. That is, we have the capacity to, should we so desire, learn what values there are, and how to apply them, should this be the sort of life we desire. Other people, equally wise, may choose to live their lives value-free, and will report a very different set of experiences. And it has yet to be shown, without prior appeal to values, that their experiences are in some way 'bad' or 'wrong'. As Nietzsche would say, the Superman makes his own rules - and who are we to say he is wrong?

To the extent that there is wisdom in society, it is the result of certain practices, that make people wise, rather than adherence to certain outcomes, that people (today) say are the mark of the wise. Learning this distinction, perhaps, is the first sign of wisdom.

Moncton, January 21, 2007

Critical Thinking in the Classroom

Introduction

Critical thinking is the use of reason in reading and writing. It enables the reader to evaluate the material being read, to recognize argument patterns and to detect inappropriate reasoning. And it allows the writer to present his or her points in a logical and reasonable manner.

As such, critical thinking is not reserved for the domain of logic and philosophy classes alone. It is a skill which has application throughout all disciplines. Indeed, expertise in any discipline is impossible without knowledge and application of critical thinking.

The purpose of this essay is to introduce the instructor to critical thinking and to suggest means of applying it in the classroom. As such, it is not a teaching document; it does not pause and repeat nor stimulate learning with examples and exercises. Rather, its purpose is to provide an overview of the field and to suggest a common terminology. A list of references is provided for those desiring more detailed study.

This essay will not attempt to persuade the reader of the merits of teaching critical thinking in the classroom. That is assumed. Rather, it focuses on what critical reasoning is and how to apply it. This essay proceeds in three major sections. First, the three major types of reasoning are described. Second, errors of reasoning in these three major types of reasoning are described. Finally, third, methods of application to the classroom are suggested.

Some notes are necessary about the approach taken. First, the methods of creating and criticizing arguments are presented as 'tools' for a student (or anyone) to use to achieve a desired outcome. Second, and related to this, it is taken that the use of a tool is flagged with 'indicator words'. That is, there are certain characteristic ways of telling the reader that you are trying to achieve a particular outcome. Hence, words themselves are regarded as tools for the expressions of an idea.

Types of Reasoning

(i) Deductive Reasoning

Deductive reasoning is the oldest and most venerable of the types of reasoning. Examples of deductive reasoning include mathematics, categorical reasoning, set theory, and computer programming. Deductive reasoning is by its very nature abstract; for this reason, students find it the most difficult to master.

A deductive argument is formed from one or more premises and a conclusion. The conclusion is the

opinion the author is attempting to prove is true. The premises are the reasons given in order to persuade the reader that the conclusion is true.

The premises and the conclusion of an argument are identified by indicator words. There are two types of indicator words: premise indicators, and conclusion indicators. Premise indicators always precede premises, while conclusions always precede conclusions. In general, the structure of deductive arguments is as follows:

(Using a premise indicator):

____ *because* ____ .

Since ____, ____ .

(Using a conclusion indicator):

____ *therefore* ____.

Notice the use of not only the indicator words ('because', 'since' and 'therefore') but also the use of punctuation and conjunctions to indicate the structure of the argument. Good writing follows a clear argument structure, and hence, good writing uses grammatical elements to show clear argument structure.

In the absence of an indicator word (some people are sloppy writers), the reader is reminded that the conclusion is an opinion. Hence, the conclusion is usually 'hedged' in some way. By that, what is meant is that the conclusion is not stated directly, but rather, is qualified with expressions like, 'I think that' or 'It must be that', or the like. Compare, for example, the difference between "The sky is blue" and "The sky must be blue". The latter is clearly hedged, hence, it must be an opinion and therefore probably the conclusion of an argument.

Not all arguments are deductive arguments. Deductive arguments may be recognized by their characteristic forms. The form of an argument can be recognized by identifying keywords. Because deductive arguments constitute a particular sort of reasoning, they entail the use of a particular set of words. In particular, there are three types of key words to watch for.

1. Mathematical keywords: plus, minus, equals
2. Categorical keywords: is, all, some, no, every, any, only
3. Propositional keywords: both...and, either, ... or, if .. then, unless

These keywords are not used only to recognize deductive arguments. Knowledge of the role of these

keywords also enables the writer to write clear, structured sentences. This will be discussed in more detail in section three.

(ii) Inductive Reasoning

The purpose of an inductive argument is to produce generalizations from matters of fact or experience. It is not as old as deductive argumentation, nor is it as well respected. Nonetheless, without inductive argumentation it would not be possible to live in the world at all.

Types of inductive reasoning include statistical generalizations, analogy, reasoning concerning cause and effect, and probability.

Like a deductive argument, an inductive argument is formed from one or more premises and a conclusion. And like a deductive argument, the purpose of an inductive argument is to persuade the reader that the conclusion is true, and the premises are given as reasons to believe that the conclusion is true. All that was said above of indicator words and hedging is also true of inductive arguments. Hence, the two can be distinguished only by their keywords.

Here are characteristic keywords of some inductive arguments:

1. Statistical keywords: most, many, five percent, usually, generally
2. Analogical keywords: is like, is similar to, like, as
3. Probabilistic keywords: the chances of, probably, likely
4. Causal keywords: causes, depends on, effect

Again, the use of these keywords tells the reader what sort of argument is being used. A reader can, for example, recognize an analogy much more clearly if the words 'like' or 'as' are used than if they are not.

(iii) Abductive Reasoning (Inference to the best explanation)

Abductive reasoning was recognized as such only in the late nineteenth century by Charles Sanders Peirce, though there are instances of it through antiquity. It is now the most common form of argument in the sciences, for it involves the postulation of theories which explain some event or regularity.

The form of an inference to the best explanation differs from that of deductive or inductive argument, though (confusingly) the same indicator words are used. In an abduction, the conclusion is some event or regularity which needs to be explained, while the premises are the theories or sets of conditions which do the explaining. That said, the word 'why' is used much more frequently in explanations, hence,

the word 'why' can be used to distinguish abductions from other forms of argument.

The most common form of an inference to the best explanation is:

The reason why ____ is because ____.

Note again that the conclusion should be some fact or regularity, while the premise is typically a theory. Very often the conclusion which is being explained is also the conclusion of an inductive argument. A writer will use induction to show that some generalization is true, and then use abduction to explain why it is true.

Abductive arguments do not have characteristic keywords (other than 'explains' and 'why'). The only way to distinguish between an inductive or deductive argument and an abduction is to determine whether the conclusion is a fact (in which case it's an abduction) or an opinion (in which case it's a deductive or inductive argument). It is important to watch for hedging words while making this distinction.

Errors of Reasoning

(i) Deductive Errors

There are two ways a deductive argument can fail: (i) the premises may be false, or (ii) the conclusion may not follow from the premises. Students often attempt a third method of evaluation: arguing directly against the conclusion. While this is allowed, it amounts to ignoring the argument in favour of the conclusion, and hence, is never decisive.

Whether or not the premises are true, if the conclusion follows from the premises, then the argument is valid. To say that an argument is valid is to say that the premises are appropriately related to the conclusion. The premises need not be true. To see this, consider the following argument: "If Mulroney is a Marxist, then he likes Castro, and he is a Marxist, hence, he likes Castro." As it happens, the premises are false. But suppose they were true. Then we can see clearly that the conclusion would have to be true as well; the premises support the conclusion.

In order to show that a deductive argument is invalid, it is necessary only to show that there is some way the premises could be true while the conclusion could be false. If this is possible, then we can see that the premises do not make the conclusion true. Consider the following example: "If the mill is polluting the river, then we can see dead fish, and we can see dead fish, therefore, the mill is polluting the river." Even if the premises actually are true, we can see that they do not support the conclusion, for it could be that something else is killing the fish, and that the mill is not polluting the river at all.

There are two major forms of invalid argument:

Denying the Antecedent. Any argument of the form "If A then B, and not A, therefore B" is invalid.

Affirming the Consequent. Any argument of the form: "If A then B, and B, therefore A" is invalid. The example of the mill (above) affirms the consequent.

The second way of criticizing a deductive argument is to show that the premises are false. Students are particularly hesitant to do this, however, it is often (all too often) accomplished with ease. Consider a categorical premise of the form "All A are B", for example, "All things which swim in the sea are fish." This is easily shown to be false by observing that there can be some A which is not B, for example, a dolphin swims in the sea, but is not a fish.

In general, premises are shown to be false by showing that their contradictories are true. Here are some common contradictions:

1. 'All A are B' contradicts 'Some A is not B'
2. 'No A are B' contradicts 'Some A is B'
3. 'If A then B' contradicts 'A and not B'
4. 'Either A or B' contradicts 'Not A and not B'
5. 'Both A and B' contradicts 'Not A'

An argument which is both valid and has true premises is called a sound argument. Sound arguments are also sometimes called cogent arguments.

(ii) Inductive Errors

All inductive arguments base their success on the similarity between the objects or events described in the premises and those described in the conclusion. This is most clear in the case of an analogy, and so we turn to the first error of inductive reasoning:

False Analogy. The two things being compared are not similar in a way which is relevant to the conclusion. For example, suppose someone argued, "An employee is like a nail. Just like a nail, an employee must be hit in the head in order to get him to work." This argument may be criticized by showing that employees are not like nails in that (i) incentives will not persuade a nail to work, but they will persuade employees to work, and (ii) a nail won't resent being hit, but an employee will.

Statistical generalizations are arguments which use some sort of sample to draw a conclusion about a population. For example, a pollster will collect a sample of opinions and draw conclusions about the population as a whole. In order for the sample to tell us anything useful about the population, the sample must be similar to the population. The two major inductive fallacies are cases where the sample may be dissimilar to the population:

Hasty Generalization. The sample is too small, and hence, we can't be sure that it is similar to the population.

Unrepresentative Sample. The sample can be shown to be in some way different from the population. For example, a survey taken in only one city is unrepresentative of the nation as a whole.

Unrepresentative samples are very common. Phone-in or write-in polls are classic examples of unrepresentative samples. So are testimonials. Many instructors value student opinions and observations in class. No doubt this makes the students feel good, but such information should not form the basis of instruction, for the individual experiences of one person constitute an unrepresentative sample.

There is a variety of things which can go wrong in causal reasoning. In order to say that A causes B, a minimum of two things must be true:

1. Generally, if A happens, then B happens
2. Generally, if A does not happen, then B does not happen

In addition, many theorists argue that there should be a third condition:

3. There must be a law of nature connecting A and B

The most common causal fallacy occurs when only the first condition is true and yet a causal relation is assumed to hold:

Post Hoc Ergo Propter Hoc (After this therefore because of this). This fallacy consists in assuming that because one thing follows the other, the one thing is caused by the other.

Good inductive arguments are called strong arguments. Bad inductive arguments are called weak inductive arguments.

(iii) Abductive Errors

There are two major ways an inference to the best explanation can go wrong: either (i) the fact to be explained is not a fact at all, or (ii) the theory which does the explaining is inadequate. Let us consider these in turn.

The fact to be explained may be false because of:

Non-support. For example, Jenny may wonder why John knows so much about physics. This 'fact' is false because of non-support if John knows nothing about physics.

Subverted Support. The argument which supports the 'fact' is not a good argument. For example, if a generalization such as "Edmontonians are cheap" was formed on the basis of one person's experience, then it is supported by an unrepresentative sample. Pointing out that this putative fact is not well supported is to subvert support.

There are also two ways a theory can be inadequate:

Untestability. Theories which cannot be tested are not good theories. Theories are tested by being used to make a prediction. If a theory cannot be used to make a prediction, then it is a poor theory. For example, if someone theorized that "Coffee keeps you awake because it has wakening properties" then this theory could be criticized because we cannot use it to predict what other things will keep us awake.

Better Alternative. If another theory can explain the same phenomenon and is a better theory, then the new theory can be used to criticize the old. There are two major criteria for the betterness of a theory: (1) the theory has a wider scope, that is, it applies to more things; and (2) the theory is simpler.

(iv) Informal Fallacies

There is also a range of error which can be committed in any type of argument. These are grouped under the heading of 'informal fallacies' ("fallacy" is a ten-dollar word for "error of reasoning").

The first grouping is Fallacies of Relevance. These are fallacies because they change the subject in some way. The following are major fallacies of relevance:

Attacking the Person. Authors commit this fallacy when they argue that because their opponent is a certain type of person, then their opponent is wrong. Students often argue that this form of argument is legitimate. For example, they argue that if a person has an interest in the outcome of an argument (say, a developer argues that some land should be rezoned) that a valid criticism may be made. This assumption is wrong.

Appeal to Force. In this fallacy, the reader is advised that some bad consequence will occur if the conclusion is not believed.

Appeal to Pity. In this fallacy, the reader is appealed to for support because the writer is in some bad state. For example, if a politician tells you how hard he worked on a piece of legislation, he is appealing to pity.

Prejudicial Language. A writer commits this fallacy when some moral value is attached to believing or not believing a conclusion. For example, "Clear thinkers agree that murder is bad" is a fallacy because it implies that people who disagree are not clear thinkers.

Appeal to Popularity. This fallacy is committed when it is argued that because most people believe a conclusion, then the conclusion is true. History is replete with examples where the majority was wrong.

The second grouping is Fallacies of Distraction. These are fallacies because while the premises in question appear to be true at first glance, closer examination shows them to be false.

False Dilemma. In this fallacy, the reader is presented with two options, and since one is unacceptable, we are forced to choose the second. The fallacy occurs when more than two options actually exist.

Argument from Ignorance. This fallacy is by far and away a student favourite. In this fallacy it is argued that because some proposition has not been proven to be true, it is therefore false.

Slippery Slope. The writer argues that if some proposition is believed, a chain of consequences will follow, leading to some unacceptable conclusion. The fallacy occurs when there is no reason to believe the consequences will actually occur.

Complex Question. This fallacy occurs when two separate points are presented as a single point. This fallacy is committed a lot on surveys, where a reader may be asked, for example, "Do you support reducing the deficit and cutting social programs?"

Begging the Question. Very often, this is the only way students know how to argue. Instead of offering support for a conclusion, the arguer instead restates the conclusion in a slightly different manner. Obviously, when the conclusion is simply restated, no support has been given for the conclusion.

The third grouping concerns Fallacies of Authority. Students tend to be very trusting of authority, even when the authority is inappropriate.

Unqualified Authority. This occurs when an authority is quoted outside his or her field of expertise. Celebrity endorsements fall within this category.

Disagreement. Even when an authority is an expert in the field, it may be that experts in the field disagree on the point in question. In such a case, an appeal to an authority is fallacious, since it is possible to quote an equally qualified authority who holds the opposite view.

Unnamed Authority. This fallacy is committed when an authority is implied but not named. This fallacy may be detected by the use of phrases such as "experts agree..." or "it is said that...". This is a fallacy because there is no way to know that the authority is an expert.

The fallacies listed in this section constitute only a partial list; they were chosen because they are committed the most frequently and because they are most often believed by students.

Applications in the Classroom

Critical reasoning has many more applications in the classroom than merely the correcting of faulty arguments. Critical thinking concerns the nature of argumentation itself, and all branches of knowledge involve some form of argument. This section will describe a number of applications of critical reasoning in the classroom.

(i) Writing

Knowledge of logical structures improves a student's writing in a direct and dramatic fashion. When logical structures are understood, the construction of a sentence is understood as an application of a particular logical structure. The following is a brief example of this process.

Simple sentences using categorical form. The structure of a categorical proposition, 'All A are B', mirrors the structure of a simple sentence. The 'A' in question is the subject of the sentence, while the 'B' is the predicate. This is useful because it helps correct problems with noun-verb agreement. Clearly identifying the subject and the predicate reminds the student that they work as a pair.

Another application of categorical form involves the use of subordinate clauses. The subject-predicate form clearly illustrates to the student the idea that subordinate clauses modify the subject (or predicate) they are attached to. Showing the student a sentence of the form:

"All men are mortal"

clarifies the form of:

"All men who are kings are mortal."

Complex sentences using logical operators. Complex sentences are formed out of simple sentences using logical operators. Consider, for example, how a complex sentence may be constructed from the simple sentences "All men are mortal" and "Socrates is a man".

If all men are mortal then Socrates is a man.

Either all men are mortal or Socrates is a man.

All men are mortal and Socrates is a man.

Even more complex sentences or paragraphs using indicator words. Using simple and complex sentences as described above, the structure of paragraphs can be detailed to students. We identify the premises and conclusion of an argument as a set of sentences. Then these sentences are assembled into a paragraph using indicator words.

If all men are mortal then Socrates is a man, and all men are mortal, therefore Socrates is a man.

All men are mortal, and Socrates is a man. Therefore, Socrates is mortal.

More complex paragraphs are constructed from more complex arguments. Consider the following:

All men are mortal and Socrates is a man. Thus, Socrates is mortal. All things which are mortal eventually die. Therefore, Socrates will eventually die.

(ii) Abstraction

Knowledge in many disciplines is abstract knowledge. This is most clearly the case in mathematics, where notation such as " $x+y=z$ " is abstract, but it is also true in many other cases. For example, in geography, students may be taught that a river meanders in a particular way. This is abstract because we are not talking about any particular river. Or in music, students are taught to read sheet music. This is abstract because sheet music is not generally written for a particular music.

Critical thinking forces a student to reason abstractly because sentences and arguments are thought of as abstract structures. The long paragraph just above should be recognized by the student as an instance of:

All A are B and S is A. Thus, S is M. All M are D. Therefore, S is D.

The benefits of abstract thought should be clear. Lessons learned in one domain are more easily applied in another domain when abstract features of the two domains are identified.

How might this be applied in a classroom? In essence, it involves imparting to the student not merely knowledge of particular matters of fact, but also the abstract form of whatever knowledge is being taught. For example, the proposition that "Rome fell because of a lack of morality" is an instance of the more general "Civilizations fall because of immorality". Students may be shown this, and also shown that the same pattern occurs in "Sodom and Gomorrah fell because of immorality" and "This civilization will fall because of immorality".

(iii) Reading

Students often misunderstand what they are reading. Often this is because they do not know what to look for in a piece of writing. This is understandable; there are many ways to go wrong when reading even a short paragraph.

For example, students often misunderstand a particular sentence. One common mistake occurs, for example, when a student interprets "Not all men are mortal" as meaning "No men are mortal". Knowing

that the contradictory of "All A are B" is "Some A are not B" would allow the student to understand that "Not all men are mortal" means "Some men are not mortal".

Students often believe that information contained in a subordinate clause is the main point of a sentence. Making the structure of categorical propositions clear corrects this error.

Students frequently miss the main point of a paragraph as a whole. Pointing to indicator words makes conclusions clear, and the conclusion is a main point of a paragraph. If a student learns to look for conclusions, misunderstandings of this sort can be reduced.

Students should be reminded on a regular basis how to extract information from a text. From time to time, it is useful to identify a key paragraph in a piece of writing and to provide an analysis of it, showing the student how to identify what each sentence says and showing the student how to identify the author's main point. Consider, for example, the following paragraph:

A country, after all, is not something you build as the pharaohs built the pyramids, and then leave standing there to defy eternity. A country is something that is built every day out of certain basic shared values. And so it is in the hands of every Canadian to determine how well and wisely we shall build the country in the future. (Pierre Trudeau, *Memoirs*, p. 366)

The use of the indicator word "so" clearly shows that the last sentence is the conclusion. There are no logical operators in the last sentence, hence, it is a simple sentence of the form "Every Canadian should determine...". The student should also note the use of an analogy in the first sentence. And notice the reasoning, in very abstract form: "A country cannot be left unattended, therefore, all people must attend to the country".

(iv) Critical Evaluation

This is the clearest application of critical thinking in the classroom. Essentially, it involves questioning the truth of premises and the validity of arguments, in other words, not taking the written (and spoken) word as Gospel. Students (and especially those coming straight from high school, where everything is Gospel) find this difficult to do.

A criticism of a point of view is, like everything else in academia, a form of argument. The conclusion is always that some argument has committed a logical error. The premises are the reasons for believing that the error occurred. The form of all critical evaluations is as follows:

The argument does such-and-such, and
Such-and-such is a fallacy,
Thus, the argument is a fallacy.

(Very often the second premise is left implicit.)

Students need to be shown that all sources, including their textbooks and their instructors (not to mention the media and their friends) can commit errors of reasoning. The best means to show them this is to critically evaluate any materials used for instruction. My own experience is that this can be very confusing for a student (one student commented, "I've never seen an instructor criticize the text before).

It is important, therefore, to state the criticism and the reason for the criticism clearly. It is also important to state the intent of posing such criticisms, specifically, that the student should not accept everything as being true, and that the student is expected to perform a similar sort of evaluation on any material. It is especially useful to encourage students to criticize the instructor, and to occasionally concede some points. Even when there is a response to be made, much more progress is made when a good criticism is acknowledged as such.

Finally, students should be required to stand the test of good reasoning. Comments in papers or in class which commit logical errors should be identified as errors in reasoning. This requires some tact. The approach should not be that the student is wrong, but rather, that the student's reasoning is flawed.

Suggested Readings

Barker, Stephen F. *The Elements of Logic*. McGraw-Hill, 1989.

Bergmann, Merrie, Moor, James, and Nelson, Jack. *The Logic Book*. Random House, 1980.

Cedarblom, Jerry, and Paulsen, David. *Critical Reasoning*. Third Edition. Wadsworth, 1991.

Copi, Irving M. and Cohen, Carl. *Introduction to Logic*. Second Edition. Macmillan, 1990.

Davis, Wayne. *An Introduction to Logic*. Prentice-Hall, 1986.

Huff, Darrell. *How to Lie With Statistics*. Norton, 1954.

Kelly, David. *The Art of Reasoning*. W.W. Norton, 1988.

Mayfield, Marlys. *Thinking For Yourself*. Wadsworth, 1987.

Pospesel, Howard. *Introduction to Logic: Propositional Logic*. Second Edition. Prentice-Hall, 1984.

Purtill, Richard. *Logic for Philosophers*. Harper and Row, 1971.

Quine, W.V.O. *Methods of Logic*. Fourth Edition. Harvard University Press, 1982.

Salmon, Merrilee. *Introduction to Logic and Critical Thinking*. Harcourt Brace Jovanovich, 1984.

Salmon, Wesley. *Logic*. Third Edition. Prentice-Hall, 1984.

Skyrms, Brian. *Choice and Chance: An Introduction to Inductive Logic*. University of Michigan, 1966.

Thomason, Richmond. *Symbolic Logic: An Introduction*. Collier-Macmillan, 1970.

Yanal, Robert. *Basic Logic*. West Publishing, 1988.

Edmonton, March 2, 1993

Necessary and Sufficient Conditions

The discussion of 'necessary and sufficient conditions' is well understood in philosophy, and as a result, I sometimes make the mistake of assuming it is commonly understood in the wider community. This post redresses this by sketching the concept and why it is important.

Conditions

To say that one thing is a *condition* for another is to say that the one thing is involved in making the second thing happen.

The most common example of a condition is a *cause*. For example, striking a billiard ball with a cue causes the ball to move. Thus, the striking of the ball is a 'condition' of the movement of the ball.

But conditions need not be causes. Giving *permission* is another type of condition. For example, a driver's license gives you permission to drive. This, having a driver's license is a 'condition' for driving.

Necessary and Sufficient Conditions

There are two ways to express conditions:

B if A (alternatively: if A then B)

B *only* if A

The first is called a *sufficient* condition. The second is a *necessary* condition.

The idea of a sufficient condition is that it is *enough* to make something happen. For example, in most cases, pushing on the gas is enough to make the car go forward. It's not the only thing that would do it; you could make the car go forward by pushing it, for example.

The idea of a necessary condition is that something will not happen *unless* the condition happens. For example, we might say that the car will not go forward unless we have turned off the parking brake. Turning off the brake is thus a necessary condition to the car going forward.

Necessary and sufficient conditions are typically used to *explain* why something happens. "Why did the car go forward?" we ask. The brake was turned off; that was necessary for the motion to happen. And then somebody pressed on the gas; that was sufficient to make it move forward.

The Logic of Conditions

The logical structures of necessary and sufficient conditions do a dance around each other.

The simplest statement of a sufficient condition is as follows:

If A then B

This is equivalent to:

If not B then not A, and

It is *not* equivalent to:

If B then A

Meanwhile, the simplest statement of a necessary condition is as follows:

If B then A

And we often use special words to indicate this special status:

B *only if* A

Not B *unless* A

This is also equivalent to:

If not A then not B

And it is *not* equivalent to:

If not B then not A

The Conditional Fallacy

Why is this important? Because it points to what is probably the most common fallacy involving conditions: *not sufficient means not necessary*.

For example, we often hear this kind of argument:

Studies show that simply spending money will not improve test scores in schools. So we should be looking at something else, like quality teachers.

What makes this a bit tricky is that the conclusion is often implicit. The conclusion, if spelled out, is that we should be doing something *instead* of throwing money at the problem.

Here's [an example](#) of the fallacy being committed. Ewan McIntosh writes, "In 2006 there was \$2 trillion spent on education by the world's governments. But money alone is not the reason we see improvement, not always." He then recommends "Getting the right people to become teachers, developing them into effective instructors (and) ensuring that the system is able to offer the best possible instruction for every child ." Presumably, *instead* of spending money on the problem - after all, Singapore didn't have to.

Here is Tom Hoffman [identifying the fallacy](#) in McIntosh's reasoning: "I don't have the slightest idea what school budgets look like in Scotland, so maybe over there it is appropriate to put across the message that more funding isn't necessary to improve education, but on this side of the pond, even this study makes it clear that improving American education requires spending more money."

The situation is represented thus: spending money is *necessary but not sufficient* to improve educational outcomes.

What this means is that *simply spending money* won't solve the problem. There are many ways to spend money that are not effective, as evidenced by many actual spendings of money that are not effective. Purchasing each mathematics class a Lear jet, for example, would certainly spend money. But it would not be very effective.

The response to this fallacy is to say, as Tom Hoffman did, that spending money is *necessary* in order to solve the problem. What this means is that, while the mere spending of money is no guarantee, nonetheless, the problem will not be solved *unless* money is spent. The supposition that the problem can be solved without spending money is a fallacy.

Causation

As you may imagine, with the logic of conditions being so entwined, it is very easy to get tangled in a mess of necessary and sufficient conditions. This is especially the case when attempting to state whether one thing will *cause* another to happen.

Many people mistake a cause as the *sufficient* condition for something to happen (sometimes thought of as the '[efficient cause](#)' or the 'causal agent'). But formally, we should think of a 'cause' as 'a necessary and sufficient condition for an effect'.

That is to say, the description of a cause needs to include, not only the sufficient conditions, but also the

necessary conditions, for an effect.

So if we sat that 'A' is a set of necessary and sufficient conditions, then when we say that 'A causes B' we mean that:

'If A then B' and 'If not A then not B'

You need both parts to ascribe a cause. You need to show that when A happens, that B also happens, but also, that it is *not a coincidence*, that is, when A does not happen, B does not happen either.

Some people at this point may argue that only a *correlation*, and not a cause, has been established. They argue that, in addition to a correlation, a causal argument must also appeal to a general principle or law of nature. This may be the case; if so, then we can simply say here that showing that 'If A then B' and 'If not A then not B' is *necessary*, but not sufficient, to show that A causes B.

Ceteris Paribus

The phrase *ceteris paribus* is Latin for 'all other things being equal' and is an important principle for understanding the concept of necessary and sufficient conditions.

Strictly speaking, the description of a cause for any event would be endless. For example, if I wanted to say that 'the car caused the accident' then I would need to say that the car exists and that the accident happened and that the earth exists and that the laws of nature are as we understand and that the accident was not a sub-temporal sentient being and that Merlin did not intervene and... well, you get the idea.

Usually, when we say that one thing is a cause for another, or that one thing is a condition for another, we assume a certain background state of affairs, which continues as it always has. This is especially important when talking about sufficient conditions, but will also come into play when talking about necessary conditions

When I said 'pressing on the gas was sufficient to move the car', I *assumed* that, as usual, the parking brake was not engaged. Because, after all, were the parking brake engaged, pressing on the gas might *not* be sufficient to move the car. Really, I should say, 'Releasing the parking brake and pressing on the gas is sufficient to move the car'. But since the parking brake is almost never engaged, it is not usually necessary to say this; I just assume it.

Similarly, when I said that 'releasing the parking brake is necessary to move the car', the presumption was that the parking brake was engaged. But most of the time, releasing the brake is *not necessary* because the brake was not engaged in the first place. I do not need to state the necessary condition.

This is why the concept of 'control' is so important in scientific experimentation. If you say 'all else being equal', then if you are measuring for results, then you need to know that all else was, in fact, equal.

Expectations

When you say 'all else being equal', you are assuming that a certain state of affairs holds, described in shorthand as 'all else'. But, of course, *something* changes, for otherwise causation would be impossible.

When you say 'all *else*' you mean 'everything not affected by the cause'. But this is essentially a statement of *expectations*. When you say 'A caused B' what you mean, in full, is that 'A caused B *instead of C*', where C denotes the alternative that would have been the case, all else being equal, had A not occurred.

[Bas van Fraassen](#) explains this at length. When you plant sunflower seeds beside the house and they grow to be six feet tall, someone may ask, "Why did the sunflowers grow here?" What they *mean* is, 'what caused them to grow (instead of to not grow)?' and not 'what caused sunflowers to grow instead of rutabagas?'.

When Tom Hoffman writes, sarcastically, "I don't get [Ewan's Scottish spin on this McKinsey \(i.e., American\) study](#) of educational systems around the globe," he is speaking of expectations. He is suggesting the production of a given effect involves spending *more* money in one context, where in the other the production of the same effect, it is implied, does not mean the spending of more money.

Tricks Involving Ceteris Paribus

This is where *ceteris paribus* gets tricky. Very often, the presumption 'all other things being equal' does not mean, strictly, '*all* other things', but rather, a subset of other things, and specifically (and importantly), the *set of necessary conditions* for the effect to happen.

Let us suppose that McIntosh said: "We can hire better teachers, but we do not need to spend money in order to do so." This is a bit of a caricature, but it is implied in the suggestion that the problem will not be solved by spending money.

Strictly speaking, this is impossible. It is not possible to hire teachers without spending money. What can only be meant is that it is not necessary to spend *more* money. He is stating, in other words, that enough money is *already being spent* to hire quality teachers. But, of course, this money is currently being spent on something else. So in *this* case, '*ceteris paribus*' means 'same amount of money spent' but does *not* mean 'spent on the same things'.

The unstated argument here is that the money being spent elsewhere should be reallocated to spending on quality teachers. But this very *necessary* condition remains unstated. This is a fallacy; the necessary condition is hidden in the *ceteris paribus* clause.

A similar fallacy exists elsewhere in the same argument. McIntosh writes, Less than 1% of African and Middle Eastern children perform at or above the Singaporean average - to be expected, you might believe, because those Singaporeans must hemorrhage cash into their education system. Wrong. Singapore spends less on Primary education than 27 of the 30 OECD countries. Fair enough. But is McIntosh recommending that funding for education in the UK be adjusted to match the funding provided to education in Singapore? Almost certainly not!

This is a case of *shifting ceteris paribus* clauses. In *Singapore*, 'all else being equal' means expenditures at Singapore's levels. But in Britain, this means something very different.

Why is this important? Because, if the expenses in Britain are *not* the same as those in Singapore, this means that there is something *very different* about Singapore which makes it possible to spend much less on education. But if Singapore is very different in precisely this way, then it is a poor analogy and cannot be used to define 'all else being equal', for, in this case, 'all else' is very *different*.

Summary

Arguments involving the use of conditions and causation are often deceptive because of the misuse of necessary and sufficient conditions.

When reading such arguments, you should not be swayed into believing that something is not necessary simply because it is not sufficient.

You should also be wary of hidden, and often shifting, assumptions about necessary conditions implicit in (frequently unstated) *ceteris paribus* clauses.

When evaluating such arguments, ask yourself simple questions. Like: if they did A, would the result be B? If they did not do A, would B not result?

Trust your intuitions. And keep in mind that if the appeal, by analogy, is to something that is unfamiliar to you - like Singapore, or like Estonia - the reason is most likely to hide some hidden difference that makes them a special case.

Moncton, January 07, 2008

Not All...

My series on homeschooling will continue. But for now, this interruption.

Dave Taylor, who is normally rational, [writes](#): "all cars should not be black." It's part of a presentation where he's trying to encourage people to "Allow experiment & change."

Fair enough. But does it have to be expressed in the form of a basic error of logic?

The sentence "all cars should not be black" *means* "No cars should be black." But this is *not* what he meant; this is as absolute as the behaviour he is trying to discourage.

What he *wants* to say here, of course, is "not all cars should be black." This allows that some cars can be black, and some cars can be other colours.

There really is no excuse for such a basic error in logic, and this particular error is far too common. Every time I read another case I wonder why basic illiteracy seems to be spreading through the educational community.

And this is not simply a matter of choice of expression, or of the changing nature of language. It is a matter of logic, not language, and logic, unlike language, does not vary with usage or over time.

For the uninitiated, the rules governing universals and negations are very simple:

All N are P = No N are not P

No N are P = All N are not P

Not All N are P = Some N are not P

Not No N are P = Some N are P

That's pretty simple, hm? These rules can easily be proven using two-circle [Venn Diagrams](#).

Want more? Here is a more [complete discussion](#) of the equivalence of two-term categorical statements (from my Guide to the Logical Fallacies).

Now... let's keep those categoricals straight, shall we?

Moncton, March 28, 2008

Educational Blogging

"I think it's the most beautiful tool of the world and it allows us the most magic thing..."

—Florence Dassylva-Simard, fifth-grade student

The bell rings, and the halls of Institut St-Joseph in Quebec City echo the clatter of the fifth- and sixth-graders. Some take their chairs in the more traditional classroom on the lower floor. Others attend to their projects in the large, open activity room upstairs, pausing perhaps to study one of the chess games hanging on the wall before meeting in groups to plan the current project. A third group steps up a half flight of stairs into the small narrow room at the front of the building, one wall lined with pictures and plastercine models of imagined aliens, the other with a bank of Apple computers.

This last group of students, eight or so at a time, fire up their browsers and log into their cyberportfolios, a publication space that Principal Mario Asselin calls a "virtual extension of the classroom."¹ This virtual space is composed of three sets of *weblogs*, or *blogs*: a classroom Web space, where announcements are displayed and work of common interest is posted; a public, personal communication zone, where students post the results of their work or reflection; and a private personal space, reserved for students' thoughts and teacher guidance.

Dominic Ouellet-Tremblay, a fifth-grade student at St-Joseph, writes: "The blogs give us a chance to communicate between us and motivate us to write more. When we publish on our blog, people from the entire world can respond by using the comments link. This way, they can ask questions or simply tell us what they like. We can then know if people like what we write and this indicate[s to] us what to do better. By reading these comments, we can know our weaknesses and our talents. Blogging is an opportunity to exchange our point of view with the rest of the world not just people in our immediate environment."²

The students at St-Joseph are reflective of a trend that is sweeping the world of online learning: the use of weblogs to support learning. And even though the world of fifth grade may seem remote to educators in the college and university system, these students, when they enter postsecondary education, may have had more experience writing online for an audience than writing with a pen and paper for a teacher. Such students will bring with them a new set of skills and attitudes.

Writes Asselin in his own blog, *Mario tout de go*: "The school administration's objective with this weblog initiative was to offer students and teachers a support tool to promote reflective analysis and the emergence of a learning community that goes beyond the school walls."³ The blogs fit the bill perfectly. "I see more than 2,000 posts and nearly 3,000 comments," says Asselin. "Because of that, I am able to name what they do and see where it comes from. I can also figure out the directions they are taking and how they do it."⁴

Institut St-Joseph is an unassuming, yellow-brick school on a tree-lined road in the west side of Quebec City. The students inside may be early adopters, but they are far from alone in their use of blogs. The phenomenon known as blogging, or weblogging, is sweeping the Internet. A February 2004 report published by the Pew Internet & American Life Project noted that at least 3 million Americans have created blogs, with similar numbers being seen worldwide.⁵ And schools have not been immune from this trend. While nobody can say for sure just how many students are blogging, inside the classroom or out, it seems clear that their numbers are equally impressive.

In his day job, Will Richardson is the supervisor of instructional technology at Hunterdon Central Regional High School in Flemington, New Jersey. But online, Richardson is known as one of the leading proponents of blogging in education and the maintainer of the *Weblogg-Ed* Web site. "More and more teachers and schools are starting to experiment with the technology as a way to communicate with students and parents," he writes. Blogs are used to "archive and publish student work, learn with far-flung collaborators, and 'manage' the knowledge that members of the school community create."⁶ And the number of educational bloggers is growing daily. The Educational Bloggers Network, sponsored by the Bay Area Writing Project and Weblogger.com, is a community of some 120 teachers and educators involved in blogging. The following announcement on the site, by San Diego State University's Bernie Dodge, is typical: "It's that time of semester again. Tonight I introduced blogging to my class of pre-service English and foreign language teachers." The result: twenty-eight new student blogs.⁷ This same pattern is being repeated in schools and universities across the United States and around the world.

In my own case, blogging evolved from three major directions. First, the blog that began as Stephen's Web (<http://www.downes.ca>) and that eventually became *OLDaily* originated as a better means for me to store bookmarks. Second, the blog that became *NewsTrolls* originated as a series of posts by Pasty Drone. Called *Media Rant News Trolls*, these were posted on the old Hotwired Threads. When eight of us, including Pasty and myself decided to leave the site in 1998, we adopted Pasty's format and name. And third, when I created *The Brandon Pages* site, about the city of Brandon, I created a blogging tool to announce new links and events.

Today, the weblog is frequently characterized (and criticized) as (only) a set of personal comments and observations. A look at the history of weblogging shows that this isn't the case. As Rebecca Blood observes: "The original weblogs were link-driven sites. Each was a mixture in unique proportions of links, commentary, and personal thoughts and essays." Bookmarks, rants and raves, news, events: all were fodder for the weblogger. Weblogs (so named in 1997 by Jorn Barger in his *Robot Wisdom* Web site) began to be recognized as such in 1999 when Jesse James Garrett, the editor of *infosift*, began compiling a list of "other sites like his." Garrett sent this list to *CamWorld*'s Cameron Barrett, who published it on his site. Soon after, Brigitte Eaton compiled a list of every weblog she knew about, creating the *Eatonweb Portal*.⁸ There is no doubt that these early lists were incomplete; weblogging was springing up around the Web more quickly than anyone realized.

Many writers assert that blogs came into their own only after the events of September 11, 2001. As Charles Cooper writes, "If you were scouring the Internet for news and context during those first terrible hours, you could have done a lot worse than eavesdropping on the free-wheeling mini-universe of Web logs chockablock with first-hand info and spirited commentary about what was going on. . . . For my money, some of the best stuff was being served up in this most unlikely venue."⁹

I myself spent the two days following 9-11 updating *NewsTrolls*. Although we had covered and commented on the tech boom, world events, and a presidential election, the events of September 11 brought home to me the immediacy of blogging. We ran ongoing coverage, submitted via SMS to my e-mail, as one of our own made her way from the dust and debris of New York's financial district to her home on the west side. Blogging not only allowed us access to the event; it made us *part* of the event. And with that, the form had indeed finally come into its own.

Barger's original definition of a weblog reads as follows: "A weblog (sometimes called a blog or a newspage or a filter) is a webpage where a weblogger (sometimes called a blogger, or a pre-surfer) 'logs' all the other webpages she finds interesting. The format is normally to add the newest entry at the top of the page, so that repeat visitors can catch up by simply reading down the page until they reach a link they saw on their last visit."¹⁰

The personal journal, also widely popular in the late 1990s, actually developed independently of weblogs. Personal journals, or online diaries, were described by Simon Firth as "direct, personal, honest, almost painful to read and yet compelling too," but by the time Firth's article in *Salon* was written in July 1998, personal journals were on the verge of extinction. "Many of the biggest journal 'fans' began online journals themselves, and soon everyone ended up mostly writing about each other. Some of them got famous, others got resentful."¹¹

The confusion between these two distinct forms is evident in the observations of commentators such as Catherine Seipp. "In general, 'blog' used to mean a personal online diary, typically concerned with boyfriend problems or techie news," she writes. "But after September 11, a slew of new or refocused media junkie/political sites reshaped the entire Internet media landscape. Blog now refers to a Web journal that comments on the news—often by criticizing the media and usually in rudely clever tones—with links to stories that back up the commentary with evidence."¹²

But this definition—which tries to characterize the blog by what it contains—seems to miss the point. Commenting on Seipp's statement, Meg Hourihan takes a different approach: "Whether you're a warblogger who works by day as a professional journalist or you're a teenage high school student worried about your final exams, you do the same thing: you use your blog to link to your friends and rivals and comment on what they're doing. Blog posts are short, informal, sometimes controversial, and sometimes deeply personal, no matter what topic they approach."¹³ The definitions of blogging offered by bloggers, as opposed to those offered by external commentators, follow this theme. Blogging is something defined by *format* and *process*, not by content.

A blog, therefore, is and has always been more than the online equivalent of a personal journal. Though consisting of regular (and often dated) updates, the blog adds to the form of the diary by incorporating the best features of hypertext: the capacity to link to new and useful resources. But a blog is also characterized by its reflection of a personal style, and this style may be reflected in either the writing or the selection of links passed along to readers. Blogs are, in their purest form, the core of what has come to be called *personal publishing*.

In the hands of teachers and students, blogs become something more again. The Web is by now a familiar piece of the educational landscape, and for those sites where personal publishing or chronologically ordered content would be useful, blogs have stepped to the fore. *Crooked Timber's* Henry Farrell identifies five major uses for blogs in education.¹⁴

First, teachers use blogs to replace the standard class Web page. Instructors post class times and rules, assignment notifications, suggested readings, and exercises. Aside from the ordering of material by date, students would find nothing unusual in this use of the blog. The instructor, however, finds that the use of blogging software makes this previously odious chore much simpler.

Second, and often accompanying the first, instructors begin to link to Internet items that relate to their course. Mesa Community College's Rick Effland, for example, maintains a blog to pass along links and comments about topics in archaeology.¹⁵ Though Mesa's archaeology Web pages have been around since 1995, blogging allows Effland to write what are in essence short essays directed specifically toward his students. Effland's entries are not mere annotations of interesting links. They effectively model his approach and interest in archaeology for his students.

Third, blogs are used to organize in-class discussions. At the State University of New York at Buffalo, for example, Alexander Halavais added a blog to his media law class of about 180 students. Course credit was awarded for online discussion, with topics ranging from the First Amendment to libel to Irish law reform. As the course wound down with a discussion of nude bikers, Halavais questioned whether he would continue the blog the following year because of the workload, but students were enthusiastic in their comments.¹⁶

Mireille Guay, an instructor at St-Joseph, notes: "The conversation possible on the weblog is also an amazing tool to develop our community of learners. The students get to know each other better by visiting and reading blogs from other students. They discover, in a non-threatening way, their similarities and differences. The student who usually talks very loud in the classroom and the student who is very timid have the same writing space to voice their opinion. It puts students in a situation of equity."¹⁷

Fourth, some instructors are using blogs to organize class seminars and to provide summaries of readings. Used in this way, the blogs become "group blogs"—that is, individual blogs authored by a group of people. Farrell notes: "It becomes much easier for the professor and students to access the

readings for a particular week—and if you make sure that people are organized about how they do it, the summaries will effectively file themselves."¹⁸

Finally, fifth, students may be asked to write their own blogs as part of their course grade. Educational Technologist Lane Dunlop wrote about one class at Cornell College: "Each day the students read a chunk of a book and post two paragraphs of their thoughts on the reading." In another class, French 304, students were given a similar exercise. Using a French-language blogging service called Monblogue, Molly, a business student, posted a few paragraphs every day.¹⁹

What makes blogs so attractive, in both the educational community and the Internet at large, is their ease of use. A blog owner can edit or update a new entry without worrying about page formats or HTML syntax. Sebastian Fiedler, a media pedagogy specialist at the University of Augsburg in Germany, has been monitoring the rise of blogs for a number of years. "Many lightweight, cost-efficient systems and tools have emerged in the personal Webpublishing realm," he writes. "These tools offer a new and powerful toolkit for the support of collaborative and individual learning that adheres to the patterns of contemporary information-intensive work and learning outside of formal educational settings."²⁰

The blogging tool is, at its heart, a form with two fields: title and entry—and the title field is optional. Learning Media Consultant Jay Cross captures the concept with his Bloggar tool. "Blog software comes with a personal Website for those who don't already have one. The software captures your words in dated entries, maintaining a chronological archive of prior entries. In the spirit of sharing inherent to Net culture, the software and the personal Websites are usually free."²¹ What needs to be kept in mind here is that with respect to blogging tools, anything other than the entry field is a bell or whistle. Since the essence of the blog is found in individual, dated entries, the essence of the blogging tool is the entry field.

Blogging software breaks down into two major categories: hosting services and installed applications.


Hosting services. A hosting service is a Web site that will give you access to everything you need in order to create a blog. It will offer a form for you to input your entries, some tools that allow you to create a template for your blog, and access to some built-in accessories. Your blog is hosted on the hosting service (hence the name), and the URL will typically reflect the hosting service's URL. In a way, blogging hosting services are very similar to the services that allowed people to host their own Web sites (services such as GeoCities or Angelfire) or their own discussions (services such as Yahoo! Groups or ezboard).

The best-known (and one of the earliest) hosting service is Blogger (<http://www.blogger.com>), founded by Pyra Labs. When the company was bought by Google early in 2003, it reporting having about 1.1 million users.²² The Blogger interface is not much more complicated than Jay Cross's Bloggar: the large field at the top allows you to submit an entry, while instructions and some options are provided in the lower pane (after you post, the help disappears, and you can view and edit your previous posts).

Another major hosting service is LiveJournal (<http://www.livejournal.com>), a name that speaks to the side of blogging that began as an online diary. Far more so than any other service, LiveJournal attempts to foster a community of users, a strategy that used to be reflected in its terms of use: "LiveJournal relies on the community it creates to maintain an enjoyable journaling environment. In order to encourage healthy community growth, new free accounts must be sponsored by a present member of LiveJournal." LiveJournal reports more than 3 million accounts, with about half that in active status. Other major blog hosting services include GrokSoup, Salon Blogs, and TypePad. Major international hosting services include FarsiBlogs, for Iranian writers, and BlogsCN, for Chinese contributors.

Installed Applications. A remotely installed application is a piece of software that you obtain from the provider and install on your own Web site. These systems are similar to Web-based applications such as ColdFusion or Hypermail. Because of this, the number of users is much lower, but those who do use them tend (arguably) to be more dedicated and more knowledgeable than those who use hosting services. Installed applications are also more suitable for institutional use, since access can be controlled.

Probably the best-known remotely installed application is Six Apart's *Movable Type* (<http://www.moveabletype.org>). As shown in the screenshot from the Learning Circuits blog back-end (figure 1), Movable Type offers numerous options for the blog author, including extended entries. Most school blogs use Movable Type. "We used this product because it is free for use by educational institutions such as schools," says the National Research Council's Todd Bingham, who with Sébastien Paquet has just completed work with Le Centre d'Apprentissage du Haut-Madawaska, an elementary school in northern New Brunswick, providing Weblogs to all its students and teachers. "In addition to its semi-open source nature, Movable Type is written in Perl and can be back-ended by a MySQL database system," Bingham adds. "Both of these products are also open-source in nature. This allows us to customize some of the features, rather than having to write something from the ground up. We were also able to set up an additional security system using this interface by using Linux's default security features. A private blog, viewable only by the teacher and a singular student, can be set up this way. This allows the student and teacher to have a private means of feedback, as opposed to the public blog open to the public."²³

Figure 1

In mid-May 2004, however, Six Apart changed its pricing strategy for Movable Type, dramatically increasing costs for sites with multiple blogs. This prompted a storm of protest from a blogging community fearful of even greater licensing changes, as typified by Mark Pilgrim's remarks: "Movable Type is a dead end. In the long run, the utility of all non-Free software approaches zero. All non-Free software is a dead end." And although Movable Type recanted, many bloggers moved to an open source blogging tool, WordPress (<http://wordpress.org/>).²⁴

Another major installed application, and one of the earliest available, is UserLand's Radio (<http://radio.userland.com>). This is an updated version of more comprehensive site-management tools such as Frontier and Manila. Instead of running on a Web server, Radio runs on the user's desktop and displays through a Web browser; blog entries are then uploaded to a Web site. In addition, "Radio includes a powerful newsreader that allows you to subscribe to all of the sites you like. Radio will automatically go out onto the Web and find new updates to sites like the NYTimes, the BBC, and weblogs that you subscribe to every hour."²⁵

UserLand's software was used to launch a high-profile blogging experiment, Weblogs at Harvard Law, which was created when UserLand's founder, Dave Winer, became a Berkman Fellow. Arising from a conference in November 2002 called "What Is Harvard's Digital Identity?" it was intended, at least in part, to establish "intellectual community" among "the University's" disparate schools and centers.²⁶ Launched in February 2003, it allows anyone with a harvard.edu e-mail address to create a weblog, and a hundred or so staff and students have done so, including Philip Greenspun, John Palfrey, and an anonymous blogger known only as "The Redhead."

Harvard's experience illustrates one of the pitfalls of hosting such free-ranging media. Though the university administration had intended not to interfere with blog content—sometimes a challenge, since staff and students can be openly critical—it was forced to step in when Derek Slater, a student, posted internal memos from Diebold Election Systems, an electronic voting-machine manufacturer, on his blog.

The memos suggested that the machines faced numerous problems, and the company threatened legal action against Slater and Harvard University.²⁷

Though the company retreated, the potential for conflict between a blog writer and an institution's administration remains. In addition to posting copyrighted or protected information, students can get into trouble for libelous content. For example, a Valley High School student in Nevada was reprimanded for writing, "Kill Alaina!" (a classmate he found irritating) and for making a vulgar comment about a teacher. In another case, a student at St. Martin High School in Mississippi was suspended for three days after using her blog to call a teacher "perverted."²⁸

Despite the risks, teachers and students alike feel the benefits make blogging well worthwhile, if for no other reason than that blogs encourage students to write. As Rosalie Brochu, a student at St-Joseph, observes: "The impact of the blogs on my day to day life is that I write a lot more and a lot longer than the previous years. I also pay more attention when I write in my blog (especially my spelling) since I know anybody can read my posts."²⁹

In one sense, asking why anyone would write a weblog is like asking why anyone would write at all. But more specifically, the question is why anyone would write a weblog as opposed to, say, a book or a journal article. George Siemens, an instructor at Red River College in Winnipeg and a longtime advocate of educational blogging, offers a comprehensive list of motivating factors. In particular, he notes, weblogs break down barriers. They allow ideas to be based on merit, rather than origin, and ideas that are of quality filter across the Internet, "viral-like across the blogosphere." Blogs allow readers to hear the day-to-day thoughts of presidential candidates, software company executives, and magazine writers, who all, in turn, hear opinions of people they would never otherwise hear.³⁰

The students at Institut St-Joseph learned about the communicative power of blogs firsthand. "In the beginning, students anticipated the audience in a restricted circle," notes Principal Asselin. "According to the comments about their work, they realized that a lot of people could react and be part of the conversation. Each student received more than ten comments related to their posts. They had not fully realized that the entire world could read them."³¹ Imagine the young students' surprise when, some time after posting a review of a circus on their blog, someone from the circus read the review and wrote back!

But perhaps the most telling motivation for blogging was offered by Mark Pilgrim in his response to and elaboration on "The Weblog Manifesto": "*Writers will write because they can't not write. Repeat that over and over to yourself until you get it. Do you know someone like that? Someone who does what they do, not for money or glory or love or God or country, but simply because it's who they are and you can't imagine them being any other way?*"³²

Pilgrim's moving declaration should be read as a cautionary note. The warning is not about bosses who don't want employees to write weblogs (though that danger exists), but this: *writing weblogs is not for*

everybody. In particular, if you feel no empathy, no twinge of recognition, on reading Pilgrim's words, then writing a weblog is probably not for you. This does not mean that you are not a part of the weblog world. It merely means that you participate in a different way.

And herein lies the dilemma for educators. What happens when a free-flowing medium such as blogging interacts with the more restrictive domains of the educational system? What happens when the necessary rules and boundaries of the system are imposed on students who are writing blogs, when grades are assigned in order to get students to write at all, and when posts are monitored to ensure that they don't say the wrong things?

After returning from a writing teachers' conference with sessions on blogging, Richard Long, a professor at St. Louis Community College, explained the issue this way: "I'm not convinced, however, the presenters who claimed to be blogging are actually blogging. They're using blogging software, their students use blogging software, but I'm not convinced that using the software is the same as blogging. For example, does posting writing prompts for students constitute blogging? Are students blogging when they use blogging software to write to those prompts?"³³

After three years of experimentation with his *Weblogg-Ed* blog, Will Richardson also expressed his doubts: "By its very nature, assigned blogging in schools cannot be blogging. It's contrived. No matter how much we want to spout off about the wonders of audience and readership, students who are asked to blog are blogging for an audience of one, the teacher." When the semester ends, "students drop blogging like wet cement." Richardson wants to teach students to write with passion, but he notes: "I can't let them do it passionately due to the inherent censorship that a high school served Weblog carries with it."³⁴

It seems clear that although blogging can and does have a significant and worthwhile educational impact, this impact does not come automatically and does not come without risks. As many writers have noted, writing a weblog appears in the first instance to be a form of publishing, but as time goes by, blogging resembles more and more a conversation. And for a conversation to be successful, it must be given a purpose and it must remain, for the most part, unconstrained.

One of the criticisms of blogs, and especially student blogs is that the students write about nothing but trivia. Examples can be seen all over the Internet. And how many students, when facing the blogging screen, feel like "Matt," who writes: "Now each time I warily approach writing a blog entry, or start writing it, or actually write it, I end up thinking 'what is the point?'—and, after all, what is?" When given their own resources to draw on, bloggers, especially young bloggers, can become frustrated and may eventually report having "committed the ultimate blogging sin of losing interest in myself."³⁵

As Richardson says, blogging as a genre of writing may have "great value in terms of developing all sorts of critical thinking skills, writing skills and information literacy among other things. We teach exposition and research and some other types of analytical writing already, I know. Blogging, however, offers

students a chance to a) reflect on what they are writing and thinking as they write and think it, b) carry on writing about a topic over a sustained period of time, maybe a lifetime, and c) engage readers and audience in a sustained conversation that then leads to further writing and thinking."³⁶

Good conversations begin with listening. Ken Smith, an English teacher at Indiana University, explains: "Maybe some folks write flat, empty posts or bad diary posts because they don't know any other genres (they just aren't readers, in one sense) *and* because [they] aren't responding to anything (that is, they aren't reading anything right now)." It's like arriving late to a party: the first act must be to listen, before venturing forth with an opinion. Smith suggests, "Instead of assigning students to go write, we should assign them to go read and then link to what interests them and write about why it does and what it means."³⁷

The jury is still out, but as Richardson suggests, "It's becoming more clear just what the importance of blogging might be." As Smith writes, "It is through quality linking . . . that one first comes in contact with the essential acts of blogging: close reading and interpretation. Blogging, at base, is writing down what you think when you read others. If you keep at it, others will eventually write down what they think when they read you, and you'll enter a new realm of blogging, a new realm of human connection."³⁸ But it is more than merely assigning topics to blog about. As Jeremy Hiebert, a Web designer and graduate student in Canada, comments, "I've seen evidence of this in courses with required e-portfolio or reflective journal elements. . . . As soon as these activities are put into the context of school, focused on topics the students are unlikely to care about much, they automatically lose a level of authenticity and engagement. These disengaged students (non-writers and writers alike) won't get the main benefits of true reflective learning no matter how good the instruction and tools are."³⁹

Despite obvious appearances, blogging isn't really about writing at all; that's just the end point of the process, the outcome that occurs more or less naturally if everything else has been done right. Blogging is about, first, reading. But more important, it is about reading what is of interest to you: your culture, your community, your ideas. And it is about engaging with the content and with the authors of what you have read—reflecting, criticizing, questioning, reacting. If a student has nothing to blog about, it is not because he or she has nothing to write about or has a boring life. It is because the student has not yet stretched out to the larger world, has not yet learned to meaningfully engage in a community. For blogging in education to be a success, this first must be embraced and encouraged.

From time to time, we read about the potential of online learning to bring learning into life, to engender workplace learning or lifelong learning. When Jay Cross and others say that 90 percent of our learning is informal, this is the sort of thing they mean: that the lessons we might expect to find in the classroom work their way, through alternative means, into our day-to-day activities.

Blogging can and should reverse this flow. The process of reading online, engaging a community, and reflecting it online is a process of bringing life into learning. As Richardson comments, "This [the blogging process] just seems to me to be closer to the way we learn outside of school, and I don't see

those things happening anywhere in traditional education." And he asks: "Could blogging be the needle that sews together what is now a lot of learning in isolation with no real connection among the disciplines? I mean ultimately, aren't we trying to teach our kids how to learn, and isn't that [what] blogging is all about?"⁴⁰

Notes

My thanks to the many educational bloggers who contributed to this article and without whom it could not have been completed: Will Richardson, Jeremy Hiebert, George Siemens, Todd Bingham, Rod Savoie, Mario Asselin, Mireille Guay, Dominic Ouellet-Tremblay, Florence Dassylva-Simard, Hugo De Larochelière, Jean-Philippe L. Côté, and Rosalie Brochu, and to all the rest of my friends in the blogosphere—you know who you are.

1. Mario Asselin, "Weblogging at the Institut St-Joseph," *Mario tout de go*, September 1, 2003, <http://carnets.ixmedia.com/mario/archives/002425.html>.
2. Visit the Institut St-Joseph public spaces online at <http://cyberportfolio.ixmedia.com/carriere/> or <http://cyberportfolio.ixmedia.com/>.
3. Asselin, "Weblogging at the Institut St-Joseph."
4. Mario Asselin, e-mail to the author, March 25, 2004.
5. See Amanda Lenhart, John Horrigan, and Deborah Fallows, "Content Creation Online," *Pew Internet & American Life Project*, February 29, 2004, http://www.pewinternet.org/pdfs/PIP_Content_Creation_Report.pdf; and "Content Creation Online," Pew Internet & American Life Project press release, February 29, 2004, http://www.pewinternet.org/PPF/r/77/press_release.asp.
6. Will Richardson, "Blogging and RSS—The 'What's It?' and 'How To' of Powerful New Web Tools for Educators," *Information Today*, January/February 2004, <http://www.infotoday.com/MMSchools/jan04/richardson.shtml>.
7. Bernie Dodge, "Birth of 28 New Bloggers," *One-Trick CyberPony*, January 20, 2004, http://home.earthlink.net/~bdodge/blog/2004_01_18_archive.html#107466556022679878, cited in the Educational Bloggers Network: <http://www.ebn.weblogger.com/>.
8. This short history and the quotation come from Rebecca Blood, "Weblogs: A History and Perspective," *Rebecca's Pocket*, September 7, 2000, http://www.rebeccablood.net/essays/weblog_history.html.
9. Charles Cooper, "When Blogging Came of Age," *CNET News.com*, September 21, 2001, http://news.com.com/2010-1071-281560.html?legacy=cnet&tag=bt_bh.
10. Jorn Barger, "Weblog Resources FAQ," *Robot Wisdom*, September 1999, <http://www.robotwisdom.com/weblogs/>.
11. Simon Firth, "Baring Your Soul to the Web," *Salon*, July 3, 1998, http://archive.salon.com/21st/feature/1998/07/cov_03feature.html.
12. Catherine Seipp, "Online Uprising," *American Journalism Review*, June 2002, <http://www.ajr.org/Article.asp?id=2555>.

13. Meg Hourihan, "What We're Doing When We Blog," *O'Reilly Web Devcenter*, June 13, 2002, <http://www.oreillynet.com/pub/a/javascript/2002/06/13/megnut.html>.
14. Henry Farrell, "The Street Finds Its Own Use for Things," *Crooked Timber*, September 15, 2003, <http://www.crookedtimber.org/archives/000516.html>.
15. Rick Effland, "The Treasure Fleet of Zheng He," *Rick Effland Blog*, April 4, 2004, http://apps.mc.maricopa.edu/blogs/asb222/reffland/archives/cat_something_to_share.html.
16. Alexander Halavais, "Media Law" course website, February 17, 2004, <http://alex.halavais.net/courses/law/>.
17. Mireille Guay, e-mail to the author, March 26, 2004.
18. Farrell, "The Street Finds Its Own Uses."
19. Lane Dunlop, comment, *Crooked Timber*, September 18, 2003, <http://www.crookedtimber.org/archives/000516.html>.
20. Sebastian Fiedler, symposium leader, "Introducing Disruptive Technologies for Learning: Personal Webpublishing and Weblogs," Ed-Media Meeting, June 24, 2004, <http://static.cognitivearchitects.com/gems/Seblogging/EdMediaSymposium.pdf>.
21. Jay Cross, "Blogs," *Learning Circuits*, April 2002, <http://www.learningcircuits.org/2002/apr2002/ttools.htm>.
22. Dan Gillmor, "Google Buys Pyra: Blogging Goes Big-Time," *SiliconValley.com*, February 15, 2003, <http://weblog.siliconvalley.com/column/dangillmor/archives/000802.shtml>.
23. Todd Bingham, e-mail to the author, April 14, 2004. See also Sébastien Paquet, "Weblogs Enter New Brunswick School," *Seb's Open Research*, April 16, 2004, <http://radio.weblogs.com/0110772/2004/04/16.html#a1539>.
24. Mena Trott, "It's About Time," *Mena's Corner*, May 13, 2004, Six Apart Web site, http://www.sixapart.com/corner/archives/2004/05/its_about_time.shtml#; Mark Pilgrim, "Freedom 0," *Dive Into Mark*, May 14, 2004, <http://diveintomark.org/archives/2004/05/14/freedom-0>; Mena Trott, "Announcing Pricing and Licensing Changes to Movable Type," *Six Log*, June 15, 2004, Six Apart Web site, http://www.sixapart.com/log/2004/06/announcing_pric.shtml.
25. "Radio UserLand v8.0.8," *PC World*, July 6, 2004, http://www.pcworld.com/downloads/file_description/0,fid,22898,00.asp.
26. Beth Potier, "Berkman Center Fellow Dave Winer Wants to Get Harvard Blogging," *Harvard Gazette*, April 17, 2003, <http://www.news.harvard.edu/gazette/2003/04.17/13-blogging.html>.
27. John Harvard's Journal, "Creating Community, On-line and Off," *Harvard Magazine*, vol. 106, no. 3 (January-February 2004), <http://www.harvard-magazine.com/on-line/010464.html>.
28. Lisa Kim Bach, "Internet Diaries: School Discipline Questioned," *Las Vegas Review-Journal*, November 10, 2003, http://www.reviewjournal.com/lvrj_home/2003/Nov-10-Mon-2003/news/22546246.html; "Miss. School Suspends Student for Calling Teacher 'Perverted' in Online Journal," *Student Press Law Center*, January 29, 2004, <http://www.splc.org/newsflash.asp?id=736&year>.
29. Mario Asselin, quoting the student, e-mail to the author, March 25, 2004.
30. George Siemens, "The Art of Blogging—Part 1," *elearnspace*, December 1, 2002, http://www.elearnspace.org/Articles/blogging_part_1.htm.
31. Mario Asselin, e-mail to the author, March 25, 2004.

32. Mark Pilgrim, "Write," *Dive Into Mark*, October 1, 2001 (no longer extant); "The Weblog Manifesto," *Talking Moose*, September 29, 2001, <http://talkingmoose.manilasites.com/2001/09/29>.
33. Richard Long, "Back from San Antonio," *2River*, March 28, 2004, <http://www.2river.org/blog/archives/000077.html>.
34. Will Richardson, "The Blogging in Schools Question," *Weblogg-Ed*, April 13, 2004, <http://www.weblogg-ed.com/2004/04/13#a1699>.
35. Matt, "Circle Limit II," *Walky Talky*, September 25, 2003, <http://walkytalky.net/archives/000062.html>.
36. Will Richardson, "Metablognition," *Weblogg-Ed*, April 27, 2004, <http://www.weblogg-ed.com/2004/04/27>.
37. Ken Smith, "CCCC Waves and Ripples," *Weblogs in Higher Education*, March 30, 2004, http://www.mchron.net/site/edublog_comments.php?id=P2636_0_13_0.
38. Will Richardson, "Reading and Blogging," *Weblogg-Ed*, March 31, 2004, <http://www.weblogg-ed.com/2004/03/31>; Smith, "CCCC Waves and Ripples."
39. Jeremy Hiebert, e-mail to the author, April 22, 2004.
40. Will Richardson, e-mail to the author, April 27, 2004.

How To Be Heard

Every man is the hero of his own story - [Dylan Hunt](#)

Writing a blog can be a lonely business. How many blogs have been started to languish with no readers day after day, week after week? Others, seemingly inexplicably, attract thousands of readers, hundreds of links.

What's the difference between them? After all, when you look at these popular blogs, you see nothing that you couldn't write yourself, *would* write yourself (if only you had some readers). Is it really just a matter of being part of the 'in group'?

The short answer is, partially, yes. People link to people they read, and they read people they know. Popular blogs build up a following over time, establish a level of trust (or at least a reputation for being interesting). Creating a popular blog is, for better or worse, like being popular at a party. It's pretty easy to be on the outside, a wallflower, wanting but never able to be a part of the conversation.

This guide will help you change that. It won't turn you into Instapudit or Scripting News. These are special cases; they got to their position by virtue of a push not available to most of us. But it will get you started. It will get you some readers, and make you a part of the conversation.

Plan

It is possible to launch a blog without a plan. But the probability of success is much lower. It may not appear to be so because of my informal style, but [OLDaily](#) was in planning for months before it actually launched. My first [abortive attempt](#) actually failed miserably, though it taught me some lessons. In the end I draw from another project I was working on, the [the MuniMall newsletter](#).

Purpose

Why do you want to write a blog? No, seriously! If you don't know *why* you're doing something, you won't know *what* you're doing.

The MuniMall Newsletter had a very clear sense of purpose. It was down-to-earth and pragmatic. It was a means to an end; the purpose wasn't merely altruistic, but rather, was in response to some specific corporate goals.

This will be true for you too (and you may as well be honest about it, at least with yourself). The purpose of your blog may be to make your name known, to attract people to your web site, and to show what you can do.

Most blogs display a certain amount of self-interest. Some blogs promote a social or political agenda. Other blogs are designed to promote a product or service. What you will notice is that, in all successful blogs, *there's something in it for the author*.

At the same time, there does need to be *some* altruistic purpose, some element of the blog that exists not to serve the author's interests but rather to serve the readers' interests. For, after all, if the reader gets nothing out of the blog, why would they read it?

Write out the purpose of your blog. This will be the core of your blog's [About](#) page. This should be the first page you write, and easily accessible from every page on your blog.

Content

What are you going to write about?

If you are not clear about what to write about, then your blog will be forever a blank page staring at you, challenging you to be creative, but resisting form or definition.

Many people start a blog thinking that they will simply write about whatever is important to them, and then after weeks of non-activity, find that nothing is so important that it deserves to be written. [OLDaily](#), for example, focuses very clearly not only on the topic area of interest - online learning - but also on the *type* of content that will be covered: "it reflects a rising trend, it describes a new approach to online learning, it recenters our thinking. Only items which look forward are included in OLDaily."

What are you going to write about? The best answer to this question is found in the question: *what do you read?* Check your bookmarks (if you don't have bookmarks, start bookmarking things that you have read and enjoyed). What's on your bookshelf? Make a list (no, really, *make a list*). What do you listen to, what do you watch on television, what types of movies do you watch, what do you talk about when you're with friends?

Collect all of this information and organize it - find not only the topic of the material, but also the *type* of material. When I look at my bookshelf, I find not only philosophy and programming texts but also a lot of science fiction, some political writing, historical non-fiction and essays on the mind. My bookmarks, meanwhile, reveal more technology sites and (not surprisingly) quite a few education technology and e-learning sites. Is it any wonder that my focus is, not merely on learning, but on the *future* of learning with technology?

The point here is that it is better not to focus on some specific topic, the way a university course does, but rather, to aim at some sort of *intersection* that touches on *all* of your interests. Anybody can write about e-learning, but *only you* can write about themes found in e-learning, romance fiction and skydiving. What would that look like? I haven't a clue - that's why I would need to ready your blog.

Don't just pick something and say "I'm going to write about that." Study yourself, and write about the things you're already thinking about.

Support

You are probably thinking, as you start your blog, that you can get your ideas, like Neil Gaiman, [out of your head](#). But Neil Gaiman is a liar (I mean that in the fondest and most respectful way). Your ideas don't spring from your brain fully formed like [Pallas from the forehead of Zeus](#). A lot goes into the creation of an idea (even short half-baked ones like the ones in OLDaily).

When you look at a blogger's website you'll find a lot of what I call support: on the side there may be a blogroll, on another page there may be a list of links and useful resources, after each post there will be a space for comments, and more. These supports aren't so much for the reader as for the blogger: they are the sources for his or her ideas. Look at the blogger's desktop and you'll find more supports: newsletter subscriptions, mailing lists, and more.

The point here is that all writing - even fiction writing - is to a large degree *reactive*. It has its origins in the prompts and stimuli that inform a person's everyday life. Readers that want to be writers recognize this, and organize and cultivate these supports. [Dave Pollard writes](#) that we need to learn how to observe the world. This means opening our eyes.

And then what? As Gaiman says, "You get ideas when you ask yourself simple questions. The most important of the questions is just, What if...? ... Another important question is, If only... And then there are the others: I wonder... ('I wonder what she does when she's alone...') and If This Goes On... ('If this goes on telephones are going to start talking to each other, and cut out the middleman...') and Wouldn't it be interesting if... ('Wouldn't it be interesting if the world used to be ruled by cats?')..."

Me? I have a vision of what the world will be like (or *should be like* - it's hard for me to distinguish). And I look at things as they are, and I ask, "How does this lead to that future world?" or "From the point of view of that future world, what's wrong with this?" My future world isn't fully formed either - but bit by bit through a thousand reflections, it is slowly being constructed in my mind and reflected in my writing.

Process

When are you going to write? A common complaint of would-be bloggers is that they have no time. Strictly speaking, of course, this is false: they have no less time than any of us, but they choose to spend it doing something other than blogging.

Sure, you could write a blog post whenever the feeling strikes you - that's how I write my articles, like this one. But without *some* discipline, one day will slide into a week, one week will slide into two, and then you'll have no readers because you aren't writing anything.

No good writing exists without a process (it is arguable that publishers create great writing not by the promise of a paycheck (which most writers disdain) but by the imposition of deadlines, which forces the writer to write something, anything, in order to avoid their wrath).

The MuniMall newsletter publishes once a week, whether or not there is any news. OLDaily publishes five days a week, taking week-ends off. Look at [Dave Pollard's site](#) and you'll see something *every day* - sometimes the content is a bit of a stretch, like yesterday's conversation with a mushroom. But the point here isn't to create something great every day - the point here is to *create* every day, and through that very process, in time, something great will emerge.

If you don't give yourself a schedule, you won't be successful. It's as simple as that. So plan ahead - what will your schedule be? How much time (and when) will you allocate *every day* (or nearly every day, or at the very least, every week) to write and publish?

Now you have a plan.

Design

This isn't really the place for a long thesis on design. Nonetheless, some words are in order, as the quality of a blog's design has an impact on readership. People don't read what they can't read, or find difficult to read. Design makes your content accessible, and that increases your readership.

Services like [Blogger](#) have taken almost all of the work out of design, offering as the service does a selection of quality templates for your blog. Other blogging services do much the same thing. Most blogs these days look more or less the same. That's good for you. It gives you a starting point. It also gives you thousands of examples to draw from.

Almost every blog out there uses one of the default designs from Blogger or some other blogging service, but if you look at the [most popular blogs](#) on the web, almost none of them do. This is not a coincidence. The predesigned templates are good, but good designers can do better.

What do the designs of popular blogs have in common?

- They are easy to read. That is probably the most crucial aspect. You will *never* get eyestrain reading them - the fonts are big and clear, and there's lots of white space.

- They load quickly. A slow page load is death for a blog. Web readers are notorious for their lack of patience. More than a second or two delay and you've lost half your readership.
- They have a unique visual identity. Even with advertising (which most top blogs have these days) nobody will confuse [BoingBoing](#) with [Daily Kos](#).
- They are light. Unless you are the [best page in the universe](#) you cannot get away with a black background and yellow text.
- They are simple. Most of the top blogs have a simpler design than any of the Blogger themes (and the Blogger themes are pretty minimal). Look at [Davenetics](#). It's hard to be much simpler than that.
- They use font sizes well. Not only is their body text large and easy to read, their headlines and page breaks stand out. Look at [Gizmodo](#). Large, clear (and short!) headlines, for quick and easy recognition.
- They use colour well. Each of the top blogs has a simple and recognizable colour theme. [Instapundit](#): black, red and grey. [Daily Kos](#): orange and white. Like sports teams.

When you plan your own blog design, you need to take each of these things into account. It is worth drawing some sketches on paper to see what it will look like (and if you can't design your own blog, it will be essential).

Font

The font is the typestyle used in the text of your blog.

With very few exceptions, you want your font to be *sans-serif* - fonts like Helvetica or Arial. They are a lot easier to read. Eight of the top ten blogs on Technorati use sans-serif fonts.

If you absolutely *must* use serif font, then you will have to pay attention to readability. Your font size will have to be larger, and you will have to provide extra leading between the lines (the way [Davenetics](#) does).

In your CSS, fonts should be defined in *point sizes*, not pixels or inches. A font of 11pt will look pretty much the same on every computer monitor, while a font of 11px will look too large on some and tiny on others.

Speed

There are three major things that slow down page loads:

- Amount of content. Do not put your life's work on your home page. Put the last week's work (or maybe the last month's work) on your home page. Sites like Blogger will manage this automatically (but not always well).
- Images. You should have images - they add a lot to your website. But remember that if an image is 28K then it will take a full second to load for someone with a 28K modem. And if it's 149K then the person will go to some other site long before your image loads.
- Doodads. Things like Java applets or Flash animations, content from remote sites (like Blogroller or Technorati), and other things that require a special action by your browser slow down page loads. Avoid these like the plague.

Colour Scheme

You are going to have a light background, a dark text colour, and a limited set of highlight colours. Live with it.

That said, you can have a distinctive look and feel. Begin by designing the graphical elements of your blog, and most especially, the blog banner. Note that the blog banner need not occupy the entire top of the page. It typically contains a distinctive image or graphic, and the name of the blog.

Then draw your site colours from your graphical element. The text should correspond to the darkest colour in the image - black, or close to black. The background should be almost the lightest (not pure white, if you can avoid it). Headlines and links comprise mid-range colours. Choose only two or three colours; more, and the sense of a colour scheme is lost.

Design

Web pages are typically designed in columns, and unless you are an *avant garde* artist, your decisions come down to:

- How many columns, two or three?
- Liquid layout or static layout?

The number of columns usually depends on your website as a whole. Typically, the large main column contains the weblog content, that is, the daily posts and (maybe) images. A smaller column contains links to the About page, archives, the RSS feed, and the blogroll. Thus, for most blogs, two columns is sufficient. However, if you have a large number of pages on your website containing articles, links and other resources, then you should have three columns: one for internal navigation, one for external navigation, and one for main content.

'Liquid layout' is layout that expands or shrinks with the browser window. Static layout stays the same size all the time. The advantage of liquid layout is that the blog does not look tiny on very large screens. However, it is difficult to design and some elements (such as photos) do not resize well. Consequently, many bloggers use a minimum static page width (usually about 760 pixels, just enough to fit the browser window on a typical 800x600 screen).

If You Can't Design Web Pages...

Then you are at a disadvantage. Fortunately the default templates in Blogger and other weblog engines are acceptable, if not outstanding.

If you can, get a friend or colleague who can design your site to do so for you; typically, site design is something that can be done once and then left alone for a while.

At the very least, you can start with a Blogger template and experiment with it. Or you can look at the source of a website you like and see how they designed it. Or you can draw from generic templates that are already out there - glish is an [extraordinarily useful](#) website.

Because, the nice thing about the web is, you can do anything you want with the simplest of tools. You can design your entire website from scratch with little more than Notepad (and maybe a beer). The web is, at its very core, accessible to anyone to design as they please.

Implement

Unless you have very good reasons for doing otherwise, use a blogging or content management tool of some sort.

This is not because it will make your blog easier to implement. In fact, if you are using your own custom design, and not the default templates, it can actually make it a bit harder. It means you will have to adapt your design to the templates supplied by the tool, which can be a time-consuming process. Rather, it is because the tools will offer better integration with the blogosphere as a whole than your custom-made website. It is beyond the ability of most people to program these themselves (not because they can't do it, but because it takes a certain amount of time to write these programs, and most people don't have that kind of time).

For example, when you create a new post, you want it to be picked up by a blog aggregation service. These services are searched by many people, and they also offer topical feeds, places where your content might stand out. Many services need to be pinged, as [described here](#). You can write this code yourself, or you can let your blogging tool do it for you.

In addition, many weblogs support a service called [trackback](#). Basically, what happens is that if you write a post about someone else's post, trackback sends a message to the other person, allowing that person's software to provide a link to your post.

These tools (and there are many of them) are constantly growing and changing. So unless you want to get involved in blog programming in a big way, use a tool.

But what tool? That depends on what you want to do and how important your blog is going to be to you. Ideally, what you want to do is set up an account with an internet service provider (ISP), get your own domain name, and use a blogging or content management package provided by the ISP. Why?

- you will get a URL that people can remember - like [www.downes.ca](#) (and a matching email, like [stephen @downes.ca](#))
- when you change schools, jobs, or whatever, your website address will not change
- you will have more control over the content of your website
- you will be able to add features that might not be supported by your institution or a blog hosting service

Do your homework. Ask people who you know about website hosting services (for those who wonder: my news website, [NewsTrolls](#), is hosted by [Jaguar LLC](#) for the last few years and I've been very happy with them; OLDaily, which attracts much more traffic, is self-hosted at NRC).

If you look at the demo [Jaguar control panel](#) you'll notice (in the right hand column) that several self-installing scripts (such as PHP-Nuke) are offered. These are quite literally one-click installations. This is typical for internet service providers, and you can shop around for a service that offers the blogging or content management system you want.

People in the field of educational technology may want to look at James Farmer's [incorporated subversion](#) hosting service. Like Jaguar, it offers a range of auto-loading software, not just blogging tools but also [Tikiwiki](#) for collaborative authoring and [Moodle](#), a learning management system.

If you choose to go this route, then I would recommend one of the following:

- [WordPress](#) - this is a straightforward blogging tool, very similar to Blogger or Moveable Type, simple, reliable and (these days) widely supported.
- [Drupal](#) - use this if you want more than just a blog, but also something like an online community with various resource pages, membership, and the link (note: most people won't become members of your site; they already have their own site - use Drupal for a collaborative project or if you want much more flexibility than a blogging tool).

What About Blogger?

Blogger is a great tool. Though there are many blog hosting services - [MSN Spaces](#), [Salon Blogs](#), [LiveJournal](#), [TypePad](#) and more, Blogger is far ahead of them (in my opinion) for one simple reason: speed.

Getting a Blogger account isn't as good as getting your own hosted service, because you are now completely dependent on the range of services offered by Blogger, and because your blog identity is now subsumed (to a degree) under Blogger's. But:

- It is free
- it is easy to use
- it is excellent for beginners

Honestly, though, survey the landscape (I've provided enough links here to get you started) and pick something you're comfortable with. Take the time; you'll be spending a lot of time on your site, so you should be comfortable with it.

Set Up Your Site

Once you've chosen your service, set up your template, create your 'About' page and any background pages you want to add, then try a few test posts to see what they look like.

Try to view your design on several systems and using several browsers (at the very minimum, test your design on both Internet Explorer and Firefox, and on Windows, Linux and Apple).

Go to a local cybercafe and test viewing your blog. While you're at the cybercafe, try entering a blog post (after all, you need to be sure you can post to your blog while you're travelling).

You are now ready to begin blogging.

Blog

For readers wondering how (as the title suggested) to be heard it may have seemed like it took a long time to get to this point, three sections and hundreds of words to be precise.

And if you have followed the suggestions above, it will seem like you've done a lot of work with no return at all thus far - after all, you haven't even written your first blog post yet, much less generated any readers.

But it is important to emphasize here that the work you've done to get ready will be reflected in every post, every word, that you type in the future. Your blog will look like the work of someone who has thought through what he or she wants to do and who has implemented that plan in a professional manner.

And your blog posts will not only seem better, they will *be* better.

But now you are ready to blog. It's time to get into the daily routine.

Generate an Information Flow

Set up an account in [Bloglines](#) (or some other [RSS aggregator](#)) and input the RSS feeds from sites you already read (here is the list of [sites I read](#) to get you started - if you click on the [export](#) you can generate the list as OPML, which you can [import](#) into your own Bloglines account).

Go to [Yahoo! Groups](#), set up an account, and search for groups that interest you ([Google Groups](#) too, though honestly, I don't use Google groups). Do a Google search for mailing lists in your areas of interest (use phrases from your About page in your search). Check major publications such as newspapers and magazines who may offer topic-specific email newsletters or RSS feeds.

The idea here is to set up a constant flow of information into your computer. Don't worry about being overwhelmed; you can always cull the useless sources later. Be ready, though, to read ruthlessly and aggressively. Scan headings and subject lines, immediately delete those that don't catch your interest. Linger only on items that definitely interest you.

Write Your Post

And pick one (at least one - I generally write six to eight short items, but my style is very different from most bloggers) to write about.

Keep in mind that you are not merely restating what the other person has already said. Your readers already have access to that content; they may already subscribe to that RSS feed (though if you cast your net widely enough, you will be able to introduce new things to readers, which always increases your value).

What you want to do now is to *add value* to the item you are writing about.

How? Well, you might want to consider the methodology outlined in one of my other papers, [Principles for Evaluating Websites](#). Read the article and offer an assessment of the resource. Is it reliable? Can it be trusted? Does your experience support what the author is saying?

Or, as suggested above, you can ask one of the 'what if' questions listed by Neil Gaiman. Suppose you read an announcement of a new product: what if every teacher used that product in the classroom? Would it work? Suppose you read about an [educational theory](#). What if we taught children that way. Would they learn?

As time goes by, and as you consider more items, you will settle into your own style. Some people relate resources to their personal experience in the classroom, other people evaluate developments according to whether they contribute to open source, still others assess resources for empirical support, theoretical soundness, or consistency with government policy.

The idea here is that you will over time develop a *critical stance* with respect to the material you read, a *frame of reference* that assists you in understanding, putting into context, and assessing resources. This is a good thing, and you should watch for it as it develops in your writing. *Don't force it*; just let it flow as you respond to individual items.

You may be thinking, "Well, that's not very original. All I'm doing is responding to things." That may be, but if you can't respond well, you have no hope of creating from scratch. In order to create even one thing worth reading, you need to establish a frame, and there's no shortcut to establishing a frame; you have to assemble it atom by atom over weeks, months and years of blogging.

Nobody said this would be quick.

Eventually...

Eventually, you should acquire something like the following habit: if it needs to be written, blog it. Taking notes at a lecture or seminar? Blog it. Responding to an email? Blog it (assuming you aren't giving away trade secrets). Got kudos for or a complaint against some product or service? Blog it. Don't just react to content online - use your entire life as raw materials for your blog. Always keeping in mind, of course, the purpose and content of your blog you designed at the outset.

This article, for example, didn't just spring up out of the blue because I felt one day like advising people how to be heard. It is the result of an email request asking the question, "How do I join the blog conversation?" Only, instead of writing a quick and not very good email response to one person, I decided to write a longer and more detailed article that would help many people.

Well, at least, I *assume* it will help (but as a blogger, you can't worry about that).

Market

Even if you do everything I've described above, nobody will read your blog.

Well, not nobody. You may get the occasional visitor from an aggregator search. But you won't really be participating in a community because, for the most part, the community will not know you exist.

Your Blog Is Your Identity

If you have followed the advice above, you have a good blog address, a URL that in some way expresses your name or who you are. A URL like downes.ca or internettime.com or incsub.org.

As a rule: *every place you would put your name, include your blog address.*

- If you have business cards, put your blog address on your card (and if you don't have cards, get some)
- If you send an email, make sure you have a proper signature that includes both your email address (because Outlook often hides them) and your blog address
- If you publish an article or essay, ask the publisher to include your blog address in the byline
- If you go to a conference, put your blog address on your nametag

And don't forget: *put your name on your blog.* Blog posts that cannot be attributed are much less likely to be cited by anyone (and if they're not cited, they're not read).

Register

Though your results will be mixed, it is probably a good idea to register your blog with some major aggregator services, especially [Yahoo!](#), [Technorati](#) and [Feedster](#).

When you register your blog, make sure you enter complete information. For example, some sites will want to know the location of your RSS or Atom feed (using Blogger? Your feed address is your website address plus 'atom.xml' - for example, <http://halfanhour.blogspot.com/atom.xml>). Create your description from your About page. Don't be wordy, but don't leave stuff out.

Robin good has a good list of places to submit your blog. <http://www.masternewmedia.org/rss/top55/>.

Launch

You only get to launch once, so do it right.

Make sure you have some good content ready for people to read - a couple week's worth of posts (you may already be attracting readers, so put a note on your site making it clear that you are in "prelaunch mode").

Pick a day that will be the actual launch of your blog.

On that day, create a 'launch' post that will introduce your blog to new readers. explain why you've created a blog and what readers can expect to find in it (this post will be a summary of your About page).

Write a *short* email announcing the launch of your blog. Describe the purpose and content of the blog, identify yourself clearly, and make sure the blog address is in the email (don't forget the 'http://' - otherwise, people won't be able to click on it).

Send this email to everyone you know. Apologize at the top of your email for "duplicate postings", and send it to every mailing list you subscribe to. Make it clear that this is a *launch*.
Is this spamming?

Emphatically, *no*.

First of all, you are sending *personal emails* to *people you know*, not a machine generated advertisement to a bunch of people you've never heard of.

Secondly, because they are your friends (or at least, people who read the same mailing list you do), they will *want* to know that you've started a blog.

Any possible lingering resentment will be tempered by the fact that you've sent your email only *once*.

Engage

To get people to visit your blog home, you have to visit them and leave some sort of calling card.

- When you read someone's blog, if you have something to add, leave a comment. Most blogs require that you include your email and blog address (and even if they don't, make sure you leave your address). But note: *don't add a comment unless you have something to add*. Think of a comment as a gift to the author, something nice you leave as a way of thanking them for their work. Even criticisms should be written this way.

- As you read posts in mailing lists, send a response from time to time. *Don't just say*, "Read my blog for a comment." Put your full comment in email (unless it's 20 pages long - but as a rule, you should not be writing many things that are 20 pages long). Your comment establishes your credibility, and if people find you credible, some (though by no means all) will follow the tactful link in your signature. Remember, your purpose here is to *add to the discussion*, not to promote yourself.

- From time to time, as appropriate, send emails to people who write magazine or newspaper article. Again, your purpose is to *give them something* they can use. Even a criticism should be worded in such a way as to suggest that, if they looked at something from a different point of view, they may find something new to say about it (and hence, be able to write a new article - after all, writers are looking for ideas too).

- After having written a few posts, new bloggers should send an introductory email to those bloggers they perceive to be leaders in the field. Your responses will vary, depending on the blogger, but don't take the responses (or lack of a response) personally. The more popular a blog, the more impact a link has. As a result, the most popular bloggers (the ones in the Technorati Top 100, say) get hundreds of such emails a day (people like me get a dozen a day).

- Go to conferences. I get more readers from conferences than anywhere else. Bring your (possibly new) business cards. Give them to vendors (who *will* read your blog, if only in the hope of finding a new customer). Give them to people you talk to at socials and parties. Participate in the discussions at conferences. If you give a presentation, put your blog address on the first and last pages of your slides, in *big* text (and leave the last page up while you are answering questions).

Be Generous

In your own blog, be generous to your sources.

You may not hear from readers if you are merely posting content without acknowledgement, but you can be sure there are bloggers out there saying, "Why doesn't he credit his sources?"

If you look at my blog, you will see the notation, "Via [some blog](#). This is to indicate, and credit, the source of the information. Credit the *actual* source, not the source you think you would have, or should have, gotten it from. Sure, you may read [Wired News](#) every day, but if you learned of a Wired News article first from [Abject Learning](#), give them credit.

Always link to any post you are discussing.

Why is this important? People who blog - the people you are trying to engage with, to get to hear you - typically scan blog aggregators for mentions of their own name. It's a bit like surveying the room in a party to see who's looking at you. And - just like at a party - when someone looks at you, don't look away, *look back* (and maybe smile, to show that you're a friend).

If you look at my Bloglines subscriptions, you will see a section called 'me and mine'. Here is where I scan the blogosphere for references to my own blog - [PubSub](#), [Technorati](#), [Feedster](#), [Blogdigger](#) and more. If you link to me, that means that (eventually) *I'll find you*.

Now the point here is, I'm not being egotistical (much) when I search for references to my own name and my own blog. Nor either will you be when you do the same. There is a very pragmatic purpose: people who link to me are more likely to be talking about things that interest me. They are good potential sources of content and ideas, blogs I might want to subscribe to, perhaps even write a post in response to.

Note that I won't link to everything, nor subscribe to everything. Neither will you; nobody has the time nor the energy for that. You may read dozens of things in a day and link to one. So you should be surprised - or hurt - if other people demonstrate the same common sense.

The main point here: if you are generous in acknowledging the (genuine) contributions other people make to your work, they are more likely to find you.

Highlight Certain Content

From time to time you will produce something of extraordinary importance or value. When this happens (and *only* when this happens) make a special effort to ensure people are able to find it.

After all, most people in the world don't read your blog and aren't so interested in your daily posts. But they *will* be interested in something extraordinary.

- Send an email to some prominent bloggers in your field advising that you have something of extra value. Keep the email short, offering a clickable link directly to the item and a short paragraph describing the content.

- Send an email to mailing lists in your field. For example, after writing this item, I [sent an email](#) to the WWWEDU mailing list. Why this list? because it is made up of more educators than, say, the RSS developers list (to which I did *not* send a notice) and more likely to find this article useful.

- If appropriate, make sure the people in your own organization are aware of your accomplishment. Send a similar email to 'all staff' in your division or department. This not only provides them with access, it gives management the information they need to market your department (or even better, *you*).

Things You Can't Help

I said at the outset that the advice in this article won't make you another Instapundit. You may have observed, as well, that I don't follow some of my own advice (I don't ping, for example, and I don't use trackback).

Blogging is a network phenomenon, and that means that it is subject to the principles of network dynamics. And one of the major principles is: the first people in the network have an unfair advantage. Because they were first, when other people search for things to read, they are more likely to be found. And as they are found by more people, the more likely they are to be found by other people.

Moreover, some bloggers can improve their position by getting a push. [Wil Wheaton](#), for example, isn't a great blogger; by any standards, he's average, maybe even less than average. But Wil Wheaton played

Wesley Crusher in Star Trek: The Next Generation. That's exposure (and built-in popularity) that most of us can't hope to get.

Indeed, if you look at the bloggers at the top of the Top 100, *all* of them got some sort of external push. Some of them have converted their popularity as magazine writers or authors into some blog goodness. Others (such as Dave Winer) created popular internet applications and became the default selection for many readers. Still others (such as [PowerLine](#)) used a political organization to gain popularity. Some [married well](#). Some belong to professional [blogging networks](#) that leverage each other.

These are rare things. They are distortions in the network. The bloggers that have benefitted from such assistance aren't better bloggers, or somehow uniquely qualified.

Audio and Video Blogging

The principles described above apply equally to audio and video blogs. However, there are some guidelines specific to audio and video related to content, production and length.

For example, one of the best video blogs I have seen is Andy Carvin's recent video [Strolling Through Kumasi Central Market](#). Why? Not only is the length reasonable, the content is *much* more than a head-and-shoulders shot of someone speaking. It is of lively and interesting content and, importantly, content most people can't see anywhere else (you will **never** see a video like this on the news, for example).

Audio

I read once ages ago (on a crusty typewritten sheet of paper) guidelines distributed by Canada's CBC for freelance radio documentary producers. It doesn't seem to be online but I remember large chunks of it:

- use a script, especially for intros and outros - the time it takes to think of what to say **really** shows up in audio and video
- background is important - try to capture not just the voices but the sounds expressive of the story
- don't use your 'radio announcer' voice -- use the 'talking to a friend about something interesting' voice
- use many voices -- interview people and capture recordings of people saying things (announcers, etc)
- watch sound levels -- ensure that the volume is the same throughout
- organize your tape into 'scenes' -- allow for a fade between scenes (in order to prevent jarring jumps in background noise)

Some other things about quality podcasts that I've noted:

- you need to create an audio 'look and feel' -- if you look at Adam Curry's [Daily Source Code](#), for example, you'll notice he has regular opening and closing dialogue and music. This is harder to do with video, but just as necessary. Note how Andy Carvin took time to add an opening title and credits
- use - and mix - multiple media - I recommend using the [Creative Commons music library](#) to best effect.

Video

Similar guidelines apply to video, as well as some guidance for the visual aspect of it. For example:

- frame an object or central scene (the 2:1 rule for visual content applies)
- capture motion, both on the part of the subject and on the part of the viewer (ie., the camera)

These aren't so much rules for promotion as rules for content. But that is part of what I was trying to say with this article - that paying attention to planning and quality content will make a big difference in how widely read and circulated your blog post is.

With respect to promotion specifically:

- podcast and vlog content is effectively invisible to search engines, therefore, provide a *good* text-based summary of the content -- the podcast or vlog should be *in addition* to a blog post rather than instead of it
- send email to people or organizations featured in the podcast or video to let them know it exists, as they are unlikely to know about it otherwise.

Revise

As time goes by, if you have followed the suggestions listed above, you *will* gain readers. I guarantee it. How many readers will vary depending on your choice of topics (let's face it: football is more popular than philosophy) and how well you express yourself. But you can pretty well count on a reasonable readership.

But don't kid yourself. After a year or so you may have dozens, maybe hundreds, of readers. That may not seem like much. But keep in mind, these are people who are reading what you write *every day*. And more importantly, these are people who will take what you have written and pass it on to *their* readers.

To put this in perspective: I have (about) five thousand readers (this is after being on the web for more than ten years, and blogging more than seven years). But any given article that I write (perhaps even this one!) can reach an audience of a hundred thousand readers - some of my works have reached such a large audience.

How? Well, simply, having a local network of even a few hundred regular readers is enough to get your content to spread into the blogosphere as a whole. Because people link, and those links are passed along, and so forth.

The size of your readership is not the size of your audience. Your footprint reaches well beyond the people who check in on you every day or every few days.

As you progress with your blogging, your readers will pull you along. When I first started, a couple hundred people signed up right away when I launched my blog. It was hard to write every day. But it's amazing how much easier it is to get motivated when you have many more readers.

Listen to your readers. Respect your readers. If they send you email, respond (even if it takes you 46 days, as it did for me to a New Zealand reader recently - I really am sorry about that). If they put comments in your blog, read their comments, and as appropriate, answer them. Always thank them - they are *giving* you something.

Listen to their suggestions. When a substantial number of readers say that maple brown is a bad background colour, plan to change it. If they complain that the font is too small, make it bigger. Don't change things reactively - wait for at least more than one complaint, because there's always someone who will complain. Plan your changes before you implement them. If you use new software, make sure it works before you implement it (I learned this first-hand).

That said - don't lose sight of who you are and why you're writing. Stay true to your purpose. You are not a commercial publication, that needs to pander in order to gain a wide audience. There is no reward to being the [Simple Life](#) of blogging.

And if it ever stops being fun, quit. Life is too short to do things you don't like doing.

Moncton, July 28, 2005

Your Career

Responding to [David Maister](#):

My experience is that if you leave your career in the hands of others, they will actively damage your career. So I think some of the points should be even more strongly worded.

For example: You write, "No one will tell you what experience you should be obtaining, let alone help you get it." Strictly speaking, this isn't true. They will recommend all sorts of experiences - company training courses, for example. But they will be the *wrong* experiences.

And you write: "Manage your own career. No one else will." Again, someone else will. They will tell you what you should do, what you are allowed to do, and what you should not do. In so doing, they will manage your career into the ground.

The point here, too, is that you should do *more* than what you were simply hired to do. *But not necessarily more for the company*. When you are at work, working on your career, you should understand that you are working, first, for your own benefit. Any benefit the employer gets out of it is an exchange of mutual value. And the employer should *never* get everything.

As they used to say to people climbing around the rigging on the high seas: one hand for the ship, one hand for yourself.

Moncton, February 09, 2007

Managing Your Blog Entry: 11 Better Tips

What is it about the writing of lists? Here we have yet another post that needs to be deconstructed because the author seems to have just slapped something together rather than thinking the topic through. In this case, it's journalist [Vincent Maher](#). *His stuff* is in italics, my reply in plain text.

A blog entry is a stub for conversation

One of the key ways to create a loyal audience for your blog is to create a community of readers who interact with each-other on your blog. This means that your blog entries should be structured in such a way that they start conversations. This means they need to be short and punchy, with a clearly defined point or set of points.

No it isn't. The point of a blog isn't to gather a loyal cadre of readers around you dutifully writing comments. And you certainly should not be writing your blog simply to entice the commenters. And if your readers aren't capable of reading anything other than short and punchy, are you sure you should be writing to them? Think this through. A blog entry isn't some place you create to prompt conversation.

A blog is a place where you **say something**. As for the commenters, they should be writing on their own blogs, where people can actually link to them.

Think about the perspectives of your audience

Getting the audience talking means you have to consider what their perspectives may be on the point you are blogging about and position your point accordingly. It doesn't make sense to waver from one point-of-view to another in your blog entry unless that's the point you want to make.

So, like, if you have conservative readers, you should position your posts around their perspectives, making sure (of course) not to actually state their perspective. If you are going "huh?" that's because this is really bad advice. Why should I align my writing to my readers' perspectives? Why shouldn't I waver around, consider different points of view, examine things from all sides. Sure, it makes it more difficult for these same conservative readers to respond with the snappy comeback that makes them read you day after the day. As for the rest, though, they appreciate the fact that I **don't treat readers like morons**. I mean, sheesh. One point of view?

Write tight headlines that encourage interest

Remember that many readers will be scanning your RSS feed along with many others, so the poignancy of your headline is critical. If the headline doesn't grab a reader's attention there is little likelihood they will click on it. (thanks to [Colin Daniels](#) for this one)

If the whole point of your blog is to attract readers, then you'll have to bait-and-switch them with catchy (but ultimately misleading) headlines. Like, say, *Sex Up Your Blog*. For the rest of us, though, we know that our readership will be looking at more than just the headline, though if they need some way to decide, our **informative and straightforward headlines** will give them a good guide.

Make points or lists and make them scan-friendly

Online readers don't like to read long columns of text unless your content is extremely compelling. A better way to get a series of complex points across is to create a list of key points that readers can scan, along with a description of each point. This will also help you structure your thoughts in a way that seems more lucid.

Well, you *could* write compelling content. Barring this, however, you will have to give ways for readers to comment on your posts without actually having read them. Of course, you are doing this at the expense of actually creating an argument or an explanation - short pithy descriptions are all you can manage. The rest of us, though, address readability issues with good design and leave the writing to take care of itself by writing **deep and compelling content** rather than tripe.

Link to the context

If you are blogging about something that other people are talking about, provide links to their conversations so you don't seem to be speaking out of context. Linking to other sites is a plus rather than a minus because it will help your readers understand where you're coming from.

Most people link to web pages and posts, not contexts. Of course, what the author means here is that you should link to the posts you're responding to. But he won't say that, because then you might have to link to sites that disagree with you, giving them Google juice and all that. To heck with it. **Link to the source** and let Google (and your readers) take care of the rest.

Quote indirectly and link

If you feel the need to quote other bloggers, don't take the easy route and cypypaste a blockquote unless there is something very specific about the original wording that you want to preserve. Rather, rephrase the quote indirectly and link it to the source.

If you quote somebody, quote them directly. And link to them, of course. By quoting directly you give the reader to see your point respond to theirs side by side. It's harder to respond to an actual argument rather than to a paraphrase, because you *still* have to interpret what they said, but you will now be held accountable for both your interpretation and your response. **Quote directly** and respond precisely.

Format long documents for print

If you have an essay with long paragraphs and an argument that needs careful development, rather make a PDF and provide a short summary of it on your blog with a link to the document.

If you have to format your long documents for print then your web page design is broken. PDF doesn't do anything you can't do in HTML, except it needs a special reader and bloats the files like crazy.

Design your pages for readability and forget the special PDF and print formats - those are for people who can't design.

Never delete anything

In blogger culture deleting something after people have read it or commented is a cardinal sin. Don't do it, rather post a correction on the original entry.

Delete spam, trolls, vile insults from right-wing attack dogs, and other garbage. Correct your spelling mistakes, fix your broken links, repair lost or gibbled images. Change your content, too, but be honest with your readers - if you update something significant, leave a note. Did you say something really stupid? You can delete it if you want to - if people want the original, there's Google Cache. Remember, it's your website - **keep your website clean and accurate** and don't let the commenters tell you what you can't do.

Troll the blogosphere for secondary conversation

If your blog entry is successful then other bloggers will blog about it. Use tools like Google Blog Search and Technorati to track what other bloggers are saying about your blog entry and update your blog with links to those conversations if they add to yours.

If you think of what other people write on their own sites as *secondary* then you have some serious rethinking to do. These other writers have taken the time to read what you've written and to expound in some length (usually much more than you can do in a comment). If you have a primary audience, *this is it*. If you are engaged in a conversation, *this is the conversation*. Yes, use these tools to find out what people are saying, not because you're tracking but because you genuinely want to **listen to what others say** about your comments.

Be active in your own conversations

Don't sit and watch the comments streaming in and do nothing, get in there! Unlike traditional journalists, the blogger's role is to steer and be part of the conversations they start.

The blogger's role is to blog. If it feels to you that this includes responding to comments, do that. If you'd rather make sure your responses are highlighted and indexed, respond in a new blog post. But don't let anyone convince you that you have to be some kind o chatterbug to be a good blogger. **Respond if you have something to say** and be a good listener otherwise.

Create buzz everywhere

Make sure there are lots of inbound links to your post. Find other blogs that are discussing the same issue, or your blog entry, and post comments with links to updated content or highlighting some of the perspectives put forward by your commenters.

People who do this sort of thing are called link spammers. Yes, in every community there's a group of people who do this, who link to each other endlessly, as though it builds some sort of blog juice. People like that demonstrate only that they don't know how the rankings work. Link as appropriate and be selective. The purpose of linking isn't to make your site popular, the purpose of linking is to **refer people to other sites**. If you don't understand this, you don't understand the web.

You see - it's like two points of view at work here. Overall, what we have is Vincent Maher who, despite writing in a new medium, still can't let go of those old media roots. To them, it's still all about accumulating as many readers as possible, about keeping them on your site, about pandering to your audience - about everything, in other words, except saying something meaningful and being honest to yourself.

I've seen a lot of former (and not so former) journalists go this route, in my own field as well as others. It's disturbing, because they (and sometimes others) think they gain some sort of credibility through popularity, as though if they grab a large enough *mass* of readers they will, by this fact, be important.

It's a chimera, of course. The keys to blogging (if not being Prom Queen) are honesty, integrity and meaningfulness. You will gain much more if you just write what you need or want to write and let the audience fall where it may. Even if you have only three readers, if you are able to connect with and really engage with them, then no number of hangers-on will replace them after the switch from substance to dross.

Moncton, October 13, 2006

Blogs in Education

What is a Blog?

A blog is a personal website that contains content organized like a journal or a diary. Each entry is dated, and the entries are displayed on the web page in reverse chronological order, so that the most recent entry is posted at the top. Readers catch up with blogs by starting at the top and reading down until they encounter material they're already read.

Though blogs are typically thought of as personal journals, there is no limit to what may be covered in a blog. It is common for people to write blogs to describe their work, their hobbies, their pets, social and political issues, or news and current events. And while blogs are typically the work of one individual, blogs combining contributions of several people, 'group blogs', are also popular.

While the earliest blogs were created by hand, blogging became widely popular with the advent of blog authoring tools. Among the earliest of these were Userland and [LiveJournal](#). Today, most bloggers use either Google's popular [Blogger](#) service or [WordPress](#). These services allow users to create new blogs and blog posts by means of simple online forms; the writer does not need to know any programming or formatting. As a result, blog aggregation services such as [Technorati](#) have reported that tens of millions of blogs have been [created](#).

Blogs are connected to each other to form what is commonly known as the 'blogosphere'. The most common form of connection is from blogs to link to each other. Blog authors may also post a list of blogs they frequently read; this list is known as a '[blogroll](#)'. Blogs may also be read through special readers, known as 'RSS readers', which aggregate blog summaries produced by blog software. Readers use RSS readers to '[subscribe](#)' to a blog. Popular web-based RSS readers include [Google Reader](#) and [Bloglines](#).

While blogs once dominated the personal publishing landscape, they now form one part in a much more diverse landscape. Many people who formerly write blogs are using social networking sites such as [MySpace](#) or [Facebook](#). Others use 'microblogging' services such as [Twitter](#). And blogs, which began as text-based services, have branched into audio blogs (also known as 'podcasts') and video blogs ('vlogs'). Authors typically upload a wide range of multimedia content such as art to sites like [Deviantart](#), videos to hosting services such as [YouTube](#), slide shows and PDFs to [SlideShare](#) and photos to sites like [Flickr](#).

Why Use Blogs In Education

Blogs are widely popular in education, as evidenced by the 400 thousand educational blogs hosted by [edublogs](#). Teachers have been using them to support teaching and learning [since 2005](#). Through years of practice, a common understanding has formed around the [benefits](#) of the use of blogs in education.

Because blogs are connected, they can foster the development of a learning community. Authors can share opinions with each other and support each other with commentary and answers to questions. For example, the University of Calgary uses blogs to create [learning communities](#).

Additionally, blogs give students ownership over their own learning and an authentic voice, allowing them to [articulate](#) their needs and inform their own learning. Blogs have been shown to contribute to identity-formation in students. (Bortree, D.S., 2005).

Further, blogging gives students a genuine and potentially worldwide [audience](#) for their work. Having such an audience can result in feedback and greatly increase student [motivation](#) to do their best work. Students also have each other as their potential audience, enabling each of them to take on a leadership role at different times through the course of their learning.

Moreover, blogging helps students see their work in different subjects as interconnected and helps them organize their own learning. Working with the teacher and informed by blogs authored by experts in the field, students can conduct a [collective enquiry](#) into a particular topic or subject matter creating their own interpretation of the material.

Blogs teach a [variety of skills](#) in addition to the particular subject under discussion. Regular blogging fosters the development of writing and research skills. Blogging also supports digital literacy as the student learns to critically assess and evaluate various online resources.

How To Use Blogging In Learning

Begin simply. [Most uses](#) of blogs in the classroom began with the instructor using blogs to post class information such as lists of readings and assignment deadlines. This fosters in the teacher a familiarity with the technology and with students a habit of regularly checking the online resource.

Lead by example. Before requiring students to blog, instructors should lead by example, creating their own blogs and adding links to interesting resources and commentary on class topics. This not only produces a useful source of supplemental information for students, it creates a pattern and sets expectations for when students begin their own blogging.

Read. Students should begin their entry into blogging by reading other blogs. Teachers should use this practice not only to demonstrate how other people use blogs to support learning but also to foster critical thinking and reading skills. Teaching how to respond to blog posts is as important as creating blog posts.

Create a context. Like the author facing a blank sheet of paper, a blogger will be perplexed unless given something specific to write about. Have students blog about a current issue, about a specific piece of

writing, or some question that comes up in the course.

Encourage interaction. Blogging should not be a solo activity. Encourage bloggers to read each other's works and to comment on them. Encouraging students to set up an RSS reader with each other's blogs will make reading and commenting a lot easier. Teachers, also, should subscribe to student blogs and offer comments, again setting an example of the expected practice.

Respect ownership. A student blog becomes important because it is a manifestation of his or her own work. However, to have this value, a student's ownership of a blog must be genuine. While reasonable limits or [codes of practice](#) need to be respected, student bloggers should have the widest latitude possible for personal expression and opinion.

Address issues immediately. The most significant danger to students online is posed by other students. In particular, [bullying](#) (or ragging) is a significant problem. It is important to spot instances of bullying as soon as they occur and to take steps to prevent further incidents. Teachers should educate themselves as online bullying can be invisible and hard to address.

Reference

Bortree, D.S. (2005). Presentation of self on the Web: an ethnographic study of teenage girls' weblogs. *Education, Communication & Information*, 5(1), 25-39 [Paywall - [http://taylorandfrancis.metapress.com/\(2cddl0eerkzf45pssrryqw\)/app/home/contribution.asp?referrer=parent&backto=issue,3,8;journal,3,13;linkingpublicationresults,1:107512,1](http://taylorandfrancis.metapress.com/(2cddl0eerkzf45pssrryqw)/app/home/contribution.asp?referrer=parent&backto=issue,3,8;journal,3,13;linkingpublicationresults,1:107512,1)]

Moncton, April 13, 2009

How To Write Articles Quickly and Expertly

Introduction: Four Types of Discursive Writing

From time to time people express amazement at how I can get so much done. I, of course, aware of the many hours I have idled away doing nothing, demur. It feels like nothing special; I don't work harder, really, than most people. Nonetheless, these people do have a point. I am, in fact, a fairly prolific writer.

Part of it is tenacity. For example, I am writing this item as I wait for the internet to start working again in the Joburg airport departures area. But part of it is a simple strategy for writing you essays and articles quickly and expertly, a strategy that allows you to plan your entire essay as you write it, and thus to allow you to make your first draft your final draft. This article describes that strategy.

Begin by writing - in your head, at least - your second paragraph (that would be the one you just read, above). Your second paragraph will tell people what your essay says. Some people write abstracts or executive summaries in order to accomplish this task. But you don't need to do this. You are stating your entire essay or article in one paragraph. If you were writing a news article, you would call this paragraph the 'lede'. A person could read just the one paragraph and know what you had to say.

But how do you write this paragraph? Reporters will tell you that writing the lede is the hardest part of writing an article. Because if you don't know what the story is, you cannot write it in a single paragraph. A reporter will sift through the different ways of writing the story - the different angles - and find a way to tell it. You, because you are writing an article or essay, have more options.

You have more options because there are four types of discursive writing. Each of these types has a distinct and easy structure, and once you know what sort of writing you are doing, the rest of the article almost writes itself. The four types of structure are: argument, explanation, definition, and description. So, as you think about writing your first paragraph, ask yourself, what sort of article are you writing. In this article, for example, I am writing a descriptive article.

These are your choices of types of article or essay:

Argument: convinces someone of something

Explanation: tells why something happened instead of something else

Definition: states what a word or concept means

Description: identifies properties or qualities of things

An **argument** is a collection of sentences (known formally as 'propositions') intended to convince the

reader that something is the case. Perhaps you want to convince people to take some action, to buy some product, to vote a certain way, or to believe a certain thing. The thing that you want to convince them to believe is the conclusion. In order to convince people, you need to offer one or more reasons. Those are the premises. So one type of article consists of premises leading to a conclusion, and that is how you would structure your first paragraph.

An **explanation** tells the reader why something is the case. It looks at some event or phenomenon, and shows the reader what sort of things led up to that event or phenomenon, what caused it to happen, why it came to be this way instead of some other way. An explanation, therefore, consists of three parts. First, you need to identify the thing being explained. Then, you need to identify the things that could have happened instead. And finally, you need to describe the conditions and principles that led to the one thing, and not the other, being the case. And so, if you are explaining something, this is how you would write your first paragraph.

A **definition** identifies the meaning of some word, phrase or concept. There are different ways to define something. You can define something using words and concepts you already know. Or you can define something by giving a name to something you can point to or describe. Or you can define something indirectly, by giving examples of telling stories. A definition always involves two parts: the word or concept being defined, and the set of sentences (or 'propositions') that do the defining. Whatever way you decide, this will be the structure of your article if you intend to define something.

Finally, a **description** provides information about some object, person, or state of affairs. It will consist of a series of related sentences. The sentences will each identify the object being defined, and then ascribe some property to that object. "The ball is red," for example, where the ball is the object and 'red' is the property. Descriptions may be of 'unary properties' - like colour, shape, taste, and the like, or it may describe a relation between the object and one or more other objects.

Organizing Your Writing

The set of sentences, meanwhile, will be organized in one of a few common ways. The sentences might be in *chronological* order. "This happened, and then this happened," and so on. Or they may enumerate a *set of properties* ('appearance', 'sound', 'taste', 'small', 'feeling about', and the like). Or they may be *elements of a list* ("nine rules for good technology," say, or "ten things you should learn"). Or, like the reporters, you may cover *the five W's*: who, what, where, when, why. Or the steps required to write an essay.

When you elect to write an essay or article, then, you are going to write one of these types of writing. If you cannot decide which type, then your purpose isn't clear. Think about it, and make the choice, before continuing. Then you will know the major parts of the article - the premises, say, or the parts of the definition. Again, if you don't know these, your purpose isn't clear. Know what you want to say (in two or three sentences) before you decide to write.

You may at this point be wondering what happened to the first paragraph. You are, after all, beginning with the second paragraph. The first paragraph is used to 'animate' your essay or article, to give it life and meaning and context. In my own writing, my animation is often a short story about myself showing how the topic is important to me. Animating paragraphs may express feelings - joy, happiness, sadness, or whatever. They may consist of short stories or examples of what you are trying to describe (this is very common in news articles). Animation may be placed into your essay at any point. But is generally most effective when introducing a topic, or when concluding a topic.

For example, I have now concluded the first paragraph of my essay, and then expanded on it, thus ending the first major part of my essay. So now I could offer an example here, to illustrate my point in practice, and to give the reader a chance to reflect, and a way to experience some empathy, before proceeding. This is also a good place to offer a picture, diagram, illustration or chart of what you are trying to say in words.

Like this: the second paragraph will consist of a set of statements. Here is what each of the four types look like:

Argument:

Premise 1

Premise 2 ... (and more, if needed)

Conclusion

Explanation:

Thing being explained

Alternative possibilities

Actual explanation

Definition:

Thing being defined

Actual definition

Description:

Thing being described

Descriptive sentence

Descriptive sentence (and more, connected to the rest, as needed)

So now the example should have made the concept clearer. You should easily see that your second paragraph will consist of two or more distinct sentences, depending on what you are trying to say. Now, all you need to do is to write the sentences. But also, you need to tell your reader which sentence is

which. In an argument, for example, you need to clearly indicate to the reader which sentence is your conclusion and which sentences are your premises.

Indicator Words

All four types of writing have their own indicator words. Let's look at each of the four types in more detail, and show (with examples, to animate!) the indicator words.

As stated above, an argument will consist of a conclusion and some premises. The conclusion is the most important sentence, and so will typically be stated first. For example, "Blue is better than red." Then a premise indicator will be used, to tell the reader that what follows is a series of premises. Words like 'because' and 'since' are common premise indicators (there are more; you may want to make a list). So your first paragraph might look like this: "Blue is better than red, *because* blue is darker than red, and all colours that are darker are better."

Sometimes, when the premises need to be stressed before the conclusion will be believed, the author will put the conclusion at the end of the paragraph. To do this, the author uses a conclusion indicator. Words like 'so' and 'therefore' and 'hence' are common conclusion indicators. Thus, for example, the paragraph might read: "Blue is darker than red, and all colours that are darker are better, *so* blue is better than red."

You should notice that indicator words like this help you understand someone else's writing more easily as well. Being able to spot the premises and the conclusion helps you spot the structure of their article or essay. Seeing the conclusion indicator, for example, tells you that you are looking at an argument, and helps you spot the conclusion. It is good practice to try spotting arguments in other writing, and to create arguments of your own, in our own writing.

Arguments can also be identified by their form. There are different types of argument, which follow standard patterns of reasoning. These patterns of reasoning are indicated by the words being used. Here is a quick guide to the types of arguments:

Inductive argument: the premise consists of a 'sample', such as a series of experiences, or experimental results, or polls. Watch for words describing these sorts of observation. The conclusion will be inferred as a generalization from these premises. Watch for words that indicate a statistical generalization, such as 'most', 'generally', 'usually', 'seventy percent', 'nine out of ten'. Also, watch for words that indicate a universal generalization, such as 'always' and 'all'.

A special case of the inductive argument is the *causal generalization*. If you want someone to believe that one thing causes another, then you need to show that there are many cases where the one thing was followed by the other, and also to show that when the one thing didn't happen, then the other didn't either. This establishes a 'correlation'. The argument becomes a causal argument when you

appeal to some general principle or law of nature to explain the correlation. Notice how, in this case, an explanation forms one of the premises of the argument.

Deductive argument: the premises consist of propositions, and the conclusion consists of some logical manipulation of the premises. A *categorical* argument, for example, consists of reasoning about sets of things, so watch for words like 'all', 'some' and 'none'. Many times, these words are implicit; they are not stated, but they are implied. When I said "Blue is better than red" above, for example, I meant that "blue is always better than red," and that's how you would have understood it.

Another type of deductive argument is a *propositional* argument. Propositional arguments are manipulations of sentences using the words 'or', 'if', and 'and'. For example, if I said "Either red is best or blue is best, and red is not best, so blue is best," then I have employed a propositional argument.

It is useful to learn the basic argument forms, so you can very clearly indicate which type of argument you are providing. This will make your writing clearer to the reader, and will help them evaluate your writing. And in addition, this will make easier for you to write your article.

See how the previous paragraph is constructed, for example. I have stated a conclusion, then a premise indicator, and then a series of premises. It was very easy to writing the paragraph; I didn't even need to think about it. I just wrote something I thought was true, then provided a list of the reasons I thought it was true. How hard is that?

In a similar manner, an **explanation** will also use indicator words. In fact, the indicator words used by explanations are very similar to those that are used by arguments. For example, I might explain by saying "The grass is green because it rained yesterday." I am explaining why the grass is green. I am using the word 'because' as an indicator. And my explanation is offered following the word 'because'.

People often confuse arguments and explanations, because they use similar indicator words. So when you are writing, you can make your point clearer by using words that will generally be unique to explanations.

In general, explanations are answers to 'why' questions. They consider why something happened 'instead of' something else. And usually, they will say that something was 'caused' by something else. So when offering an explanation, use these words as indicators. For example: "It rained yesterday. That's why the grass is green, instead of brown."

Almost all explanations are *causal explanations*, but in some cases (especially when describing complex states and events) you will also appeal to a *statistical explanation*. In essence, in a statistical explanation, you are saying, "it had to happen sometime, so that's why it happened now, but there's no reason, other than probability, why it happened this time instead of last time or next time." When people see somebody who was killed by lightning, and they say, "His number was just up," they are offering a

statistical explanation.

Definitions are trickier, because there are various types of definition. I will consider three types of definition: ostensive, lexical, and implicit.

An '*ostensive*' definition is an act of naming by pointing. You point to a dog and you say, "That's a dog." Do this enough times, and you have defined the concept of a dog. It's harder to point in text. But in text, a description amounts to the same thing as pointing. "The legs are shorter than the tail. The colour is brown, and the body is very long. That's what I mean by a 'wiener dog'." As you may have noticed, the description is followed by the indicator words "that's what I mean by". This makes it clear to the reader that you are defining by ostension.

A '*lexical*' definition is a definition one word or concept in terms of some other word or concept. Usually this is describes as providing the 'necessary and sufficient conditions' for being something. Another way of saying the same thing is to say that when you are defining a thing, you are saying that 'all and only' these things are the thing being defined. Yet another way of saying the same thing is to say that the thing belongs to such and such a category (all dogs are animals, or, a dog is necessarily an animal) and are distinguished from other members in such and such a way (only dogs pant, or, saying a thing is panting is sufficient to show that it is a dog).

That may seem complicated, but the result is that a lexical definition has a very simply and easy to write form: A (thing being defined) is a type of (category) which is (distinguishing feature). For example, "A dog is an animal that pants."

This sentence may look just like a description, so it is useful to indicate to the reader that you are defining the term 'dog', and not describing a dog. For example, "A 'dog' is defined as 'an animal that pants'." Notice how this is clearly a definition, and could not be confused as a mere description.

The third type of definition is an *implicit* definition. This occurs when you don't point to things, and don't place the thing being defined into categories, but rather, list instances of the thing being defined. For example, "Civilization is when people are polite to each other. When people can trust the other person. When there is order in the streets." And so on. Or: "You know what I mean. Japan is civilized. Singapore is civilized. Canada is civilized." Here we haven't listed necessary and sufficient conditions, but rather, offered enough of a description as to allow people to recognize instances of 'civilization' by their resemblance to the things being described.

Finally, the **description** employs the 'subject predicate object' form that you learned in school. The 'subject' is the thing being described. The 'predicate' is something that is true of the subject - some action it is undertaking, or, if the predicate is 'is', some property that it possesses. And the 'object' may be some other entity that forms a part of the description.

As mentioned, the sentences that form a description are related to each other. This relation is made explicit with a set of indicator words. For example, if the relation is chronological, the words might be 'first'... 'and then'... 'and finally'...! Or, 'yesterday'... 'then today'... 'and tomorrow'...

In this essay, the method employed was to identify a list of things - argument, explanation, definition, and description - and then to use each of these terms in the sequence. For example, "An argument will consist of a ..." Notice that I actually went through this list twice, first describing the parts of each of the four items, and then describing the indicator words used for each of the four items. Also, when I went through the list the second time, I offered for each type of sentence a subdivision. For example, I identified inductive and deductive arguments.

Summary

So, now, here is the full set of types of things I have described (with indicator words in brackets):

- Argument (premise: 'since', 'because'; conclusion: 'therefore', 'so')
 - Deductive
 - Categorical ('all', 'only', 'no', 'none', 'some')
 - Propositional ('if', 'or', 'and')
 - Inductive
 - Generalization ('sample', 'poll', 'observation')
 - Statistical ('most', 'generally', 'usually', 'seventy percent', 'nine out of ten')
 - Universal ('always' and 'all')
 - Causal ('causes')
- Explanation ('why', 'instead of')
 - Causal ('caused')
 - Statistical ('percent', 'probability')
- Definition ('is a', 'is defined as')
 - Ostensive ('That's what I mean by...')
 - Lexical ('All', 'Only', 'is a type of', 'is necessarily')
 - Implicit ('is a', 'for example')
- Description
 - Chronology ('yesterday', 'today')
 - Sensations ('seems', 'feels', 'appears', etc.)
 - List ('first', 'second', etc.)
 - 5 W's ('who', 'what', 'where', 'when', 'why')

Complex Forms

As you have seen in this article, each successive iteration (which has been followed by one of my tables) has been more and more detailed. You might ask how this is so, if there are only four types of article or essay.

The point is, each sentence in one type of thing might be a whole set of sentence of another type of thing. This is most clearly illustrated by looking at an argument.

An argument is a conclusion and some premises. Like this:

Statement 1, and
Statement 2,
Thus,
Statement 3

But each premise might in turn be the conclusion of another argument. Like this:

Statement 4, and
Statement 5,
Thus,
Statement 1

Which gives us a complex argument:

Statement 4, and
Statement 5,
Thus, Statement 1
Statement 2
Thus Statement 3

But this can be done with all four types of paragraph. For example, consider this:

Statement 1 (which is actually a definition, with several parts)
Statement 2 (which is actually a description)
Thus,
Statement 3

So, when you write your essay, you pick the main thing you want to say. For example:

Second paragraph:

Statement 1, and
Statement 2
Thus
Statement 3

Third paragraph:

Statement 4 (thing being defined)

Statement 5 (properties)
Statement 1 (actual definition)

Fourth Paragraph

Statement 5 (first statement of description)
Statement 6 (second statement of description)
Statement 2 (summary of description)

As you can see, each simple element of an essay - premise, for example - can become a complex part of an essay - the premise could be the conclusion of an argument, for example.

And so, when you write your essay, you just go deeper and deeper into the structure.

And you may ask: where does it stop?

For me, it stops with descriptions - something I've seen or experienced, or a reference to a study or a paper. To someone else, it all reduces to definitions and axioms. For someone else, it might never stop.

But you rarely get to the bottom. You simply go on until you've said enough. In essence, you give up, and hope the reader can continue the rest of the way on his or her own.

And just so with this paper. I would now look at each one of each type of argument and explanation, for example, and identify more types, or describe features that make some good and some bad, or add many more examples and animations.

But my time is up, I need to board my flight, so I'll stop here.

Nothing fancy at the end. Just a reminder, that this is how you can write great articles and essays, first draft, every time. Off the top of your head.

Johannesburg, September 13, 2006

Principles for Evaluating Websites

How do you know whether something you read on the web is true? You can't know, at least, not for sure. This makes it important to read carefully and to evaluate what you read. This guide will tell you how.

1. There Are No Authorities

Authorities used to be people you could trust. When you read it in the newspaper, for example, it was probably true. When a scientist reported a finding, you could count on it. But today, you can't trust the authorities.

Why not? There are many reasons, but here are some of the major ones:

- Authorities lie. Not all authorities, and not all the time, but frequently enough to mean you can't simply trust them.
- People impersonate authorities. A site may look like a newspaper or a government publication, but it might not be.
- Authorities are sometimes fooled. They may rely on bad data. They may be reporting something they heard.

Even if you trust the authority you are reading, you need to evaluate what they say for yourself. People don't always mean to mislead you, but they do.

This is the most important principle of reading on the internet. You must determine for yourself whether or not something is true.

2. What You Know Matters

If you saw the local grocery this morning, and then someone told you it burned down last night, you would know they were wrong because of what you saw. And you would probably say so.

You can depend on your own knowledge. And you should use this knowledge when you read websites. That doesn't mean that you cannot be wrong. But most people don't give themselves enough credit. They are too quick to assume that they must have been wrong.

Your own experiences matter. If someone says some software is easy to install, and you found that it wasn't that easy to install at all, don't simply assume that you can't install software. If it wasn't easy for you to install, it wasn't easy, and someone who says it is easy is wrong.

3. Keep Count

You can't check everything for yourself. Eventually, you will have to depend on what other people say. You can't simply assume that what they say is true.

The key here is trust. You need to learn who to trust.

The way you learn to trust someone is through repeated contact. They not only say things you know are true, they don't say things you know are not true. You need to keep track of this for yourself.

When a website says something, you need to ask yourself, have they misled me before? Websites usually follow a pattern; sites that are trustworthy generally stay trustworthy, while sites that mislead you once will likely mislead you again.

That doesn't mean you never question what they say. Always check what they say against your own experience. But if you don't know, depend on the sites you already trust rather than the ones you don't.

4. Facts and Appearances

Many people are very careful about appearances. Governments and businesses especially take great care to manage their image. Individual people, too, try to cast themselves in the best light possible. They do this because people trust people who look good. Politicians always take care to dress nicely. Con artists are often dressed in suits. Businesses spend a lot of money to make their buildings and their websites look nice.

People create appearances in words as well. For example, they often use adjectives and adverbs to suggest how you should feel about something. They also use loaded terms to suggest that something is good or bad. Compare the following:

"This respected software reliably saves your data in the most efficient format."

"This suspicious software misleadingly saves your data in a common format."

The first software sounds a lot better than the second software. But in fact, they do exactly the same thing!

In your mind, remove the adjectives and adverbs from any sentence you read. Convert any loaded terms to neutral terms (for example, convert a sentence like "He claimed..." to "He said...").

In other words, practice distinguishing the facts in a sentence from how they appear.

You may be tempted to distrust things that use a lot of adjectives, adverbs and loaded terms. And certainly you should be suspicious. But sometimes people just write that way; it doesn't mean they're lying. And sometimes people try to fool you by writing in plain and straightforward language.

The main thing is, find the facts. You can check facts. And just ignore the appearances.

5. Generalizations Are Often Untrustworthy

When you look at facts, you will see that there are two types: specifics and generalizations.

A specific is a statement about one thing, one person or one event. "John went to the store yesterday" is a specific.

A generalization talks about a group of things, many people, or a number of events. "John always goes to the store."

People use generalizations because generalizations help them predict the future. If you know that John always goes to the store, then you can predict that he will go to the store tomorrow. Generalizations also often explain why something happens. John knows the shopkeeper because he always goes to the store.

There are two types of generalizations:

A universal generalization talks about everything. When someone says "All dogs are animals", for example, they are talking about every single dog.

A statistical generalization talks about a number of things, but not all of them. When someone says "Most dogs are brown," they are talking about a large number of dogs, but not all of them.

It is important to keep in mind that most universal generalizations are false. Not always - after all, it is true that all dogs are animals.

But people often make universal generalizations that are false. And in fact, when you read universal generalizations on a website, you should be very skeptical.

Watch for the following words: all, none, only, never, always, completely. And words that mean the same sort of thing. These indicate a universal generalization. When people use them, ask yourself, is this true? Are there no exceptions? And if you know that there are exceptions, then the source is less trustworthy.

6. Absolutes Are Hidden Generalizations

People often make generalizations without realizing that they are doing it. And they might fool you into thinking that something is a fact, when it is actually a questionable generalization.

"The Chinese cannot be trusted." This looks like a statement of fact, doesn't it? But ask yourself, how many Chinese people is this person talking about? All of them? Most of them? There are a billion Chinese - how could this person possibly know that they cannot be trusted?

And of course, they can't. You have no reason to trust such a statement. And a person who makes such a statement is less trustworthy.

7. Statistics Are Often Misleading

As the truism says, "There are lies, damned lies, and statistics." People are often skeptical of statistics, and for good reason. There are many ways statistics can be used to mislead.

Statistics must be based on data. For example, for somebody to say that "most dogs are brown" they would have had to go out and actually count some dogs to see how many of them are brown. Statistics that are not supported with data should not be trusted at all.

Even if there is data, statistics can still mislead. There are two major ways statistics can mislead: The sample size is too small. If you know five Americans, and four of them are crooks, is that sufficient to conclude that most Americans are crooks? Of course not. There are 330 million Americans; you need to meet more than five before you can start making generalizations.

The sample is unrepresentative. If you wanted to know about Americans, and took your sample from a prison population, would you get a good result? Of course not - most Americans are not in prison, and are very different from prisoners.

Remember at the beginning of this article where I said that there are no authorities? When you look at the statistics produced by authorities, many of them break one of these two rules. What would you say about a scientist who surveyed 21 graduate studies and drew a conclusion about all people? Not much - but many papers that do exactly this are published.

Statistics are often misleading in ordinary writing as well. Often, they are disguised: a person might use words like 'most', 'often', 'many' or 'usually'. And their data will be suspect. A person might say, for example, "Most people are generous." How does he know? Because most of the people he knows are generous. But that's not good data at all!

Think about the generalizations you believe. Are they based on good data? What is the data? I said above that you should trust yourself - but you should always review your own beliefs, to make yourself more trustworthy.

8. Go to the Source

People say things about other things and other people. That's no surprise; you can't talk about yourself all the time. For example, a person might report about what someone else said, or about what some data shows.

They may not mean to mislead you, though sometimes they do:

They might have misread or misunderstood the original document. Heck, I do that myself.

They may have quoted something out of context. For example, I may have written, "If people vote the wrong way then we'll have private health care" and be quoted as saying "We'll have private health care."

They may be misrepresenting the original. People sometimes pretend that someone said something that they didn't, so they can make the other person look bad (that's called a straw man).

When you read something you always need to ask, are they talking about something else and especially what somebody else said or reported. If so, go to the source to find out for yourself what the other person really said.

If there's no link or reference to the source, don't believe it. And even more importantly, websites that don't offer links or references are less trustworthy.

If you can't find the original source, try searching for the same information. Other people may have seen the same source and reported on it themselves. They may have described it differently. You may never know exactly what was said, but if people on different sides of the same issue agree on what was said, then it's more likely to be true.

9. Motives and Frames Matter

Most content on the web is trying to convince you that something is true. That's why it's on the web in the first place.

Usually, what they want you to believe isn't just some isolated fact or data, but rather a whole collection of facts and data. They want you to see the world in a certain way. In philosophy, this is sometimes called a 'world view' while in linguistics this is called a 'frame'.

Here are some examples of frames:

- It's a dangerous world and we have a lot to fear
- Microsoft products cannot be trusted
- Our country is the best (most free, most democratic, most advanced, etc.)

Think about all the sorts of things that could lead you to believe any of these three statements. Think about other sorts of things that might also be frames. Think about the way you look at the world - you probably view it from a certain frame, whether or not you recognize it.

That's not bad in itself- we all have to have a way of looking at the world. But we need to choose this way of looking at the world for ourselves. That's why we need to understand what frames other people believe, so we know when we are being persuaded to look at the world one way or another.

That's why motives matter. A person's motive is the frame or worldview he or she wants you to accept. You need to know why somebody is telling you something as well as what they are telling you.

Websites that hide their motives are untrustworthy. They are trying to convince you of something, but they are trying to do it in a sneaky way, so that you can't make your decision for yourself. They think that if you just hear something over and over, and it all points to a certain way of looking at the world, that you will start seeing the world that way too.

If a website is sponsored by the government, but they hide this sponsorship, then they are hiding their motives. If a study is financed by a software company, but this financing is not revealed, then they are hiding their motives. If a news site is secretly sponsored by a religious organization, then the news site is untrustworthy. If an activist group is funded by the industry they are trying to change, then this group is untrustworthy.

They are not untrustworthy because what they are saying is false. They are untrustworthy because they are not being honest about why they are saying what they are saying.

10. Beware Misdirection

Have you even seen a political ad for one candidate that talks about the other candidate? Have you ever seen an advertisement about one product that only talks about another product?

These are cases of misdirection - they are trying to get you believe one thing by talking about another thing.

Misdirection is very common on the web. Sometimes it consists of misrepresenting the source, as discussed above. Very often, though, it consists of merely attacking the source.

You see this not only on discussion lists (where it is very common) but also on personal websites, corporate websites, political websites and even academic websites.

If a website is trying to convince you to believe one thing but actually talks about another thing, then the website is not trustworthy.

Summary

As I said in the second point, determining what to believe - or to not believe - is a matter of trust. You need to determine for yourself who to trust about what.

This is something you have to determine for yourself. Each time you look at a website, think of yourself as keeping score. When a website does something untrustworthy, take some trust away. When a website does something well, add some trust.

And it's something very personal. The better you get to know a website, the more easily you can determine whether or not to trust it. The website gradually acquires a track record with you. Just like a friend or an associate.

And finally, this is something that works best if you use diverse sources. Try to read points of view from different frames - after all, every frame has an element of truth to it. Don't just go with the flow, be ready to challenge and question everything - even yourself.

Examples

40 Things That Only Happen In Movies

Should you trust this site? The title should let you know that this is intended as humour. But if not, you should be alerted by the universals in this title. They are probably exaggerating to make a point.

Look at some of the assertions., "(In movies) any lock can be picked with a credit card or paperclip in seconds. " Well you know that this isn't true. People don't always pick locks in movies. Sometimes they can't even break the door down.

This site is funny. But you shouldn't trust it to tell you true things about the world.

Top Chinese general warns US over attack

This news article is offered by the Financial Times, a British news source with strong links to the British and American financial communities. The story reports that a Chinese general said that China would use nuclear arms if attacked.

Did the general say this? Probably. The general is named - Zhu Chenghu - and the place where he made the remark is also named - a function for foreign journalists (it would be better if they actually named the function and told us who else, in addition to the Chinese government, sponsored it). And a one-minute search in Google for 'Zhu Chenghu' links to other reports - from the BBC and the Times of India, for example - with the same information.

Is what the general said true? We have no way of knowing. Even the Financial Times article notes that Zhu is not a high-ranking official and that "Gen Zhu probably did not represent the mainstream People's Liberation Army view." Coverage elsewhere, for example in the BBC, reports that the Chinese government is "downplaying" the remark.

So now the key question is, why did the Financial Times run the article? The article is intended to shape our views even if we cannot know whether what was said was true. Does it make us fear China more? Do the British and American financial communities stand to gain if readers fear China or become more concerned about nuclear war? Does this article fit a pattern in Financial Times coverage of China?

In my opinion, this article, although an accurate report, makes the Financial Times a bit less trustworthy.

Iraqis March Against Terror

This article is found in a blog titled BlackFive. It tells us that about 1000 Iraqis in the city Qayarah, Iraq, marched against terror, and that "you probably haven't heard about it from Peter Jennings or Dan Rather. " The post includes a number of photographs of the demonstration taken by "Army Specialist David Nunn."

As one person commented, "Rather retired early in the year and Jennings has been off battling lung cancer for months." However, a search in Google shows that the protest was not covered by any major news outlet.

That a protest did happen seems evident from the pictures. Examination of the pictures, however, shows the banners to read "The juboor's tribe and its allies ask the coalition forces to release the highly-ranked officer Farhan Muthallak who was imprisoned by the coalition forces" in both English and Arabic. A Google search for "Army Specialist David Nunn" reveals no citations not associated with this particular story.

This story is very untrustworthy. It reports a protest for one thing as a protest for something else. The source of the photographs cannot be verified. It attempts, further, to discredit the news media, thus engaging in misdirection. The site (and other sites, for many sites ran this item) is much less trustworthy as a result of running this item.

It is worth noting - as demonstrated in the trackbacks - that this story has been widely circulated. This is common, even for untrustworthy stories. That is why it is important to read, not only numerous source, but also diverse sources. And to check the data for yourself.

Again, one should ask why such a blatantly misleading story achieved such wide circulation.

Secure RSS Syndication

This site suggests that there is a need for encrypted RSS feeds and demonstrates how it is done. The need expressed is the author's own, and two potential solutions are considered and rejected. The code used to generate the encryption is provided, along with samples of the encrypted data.

This article is very trustworthy. Very specific information is given, and in a form (via computer code) that can be directly verified by the reader. It should be noted that one argument ("Atom isn't finished") will cease to be true at a future point; if you were reading this article after Atom is finished you would want to check to see whether it satisfies the need as well.

This article is supportive of the idea that encrypted content syndication is a good idea. This suggests that the author may have an interest in promoting commercial applications of content syndication. But such a conclusion should not be drawn without looking at a large number of other items written by the same author.

Bastille Day

This is a Wikipedia article about Bastille Day. Readers should note that Wikipedia articles frequently change. This article was current at 11:40 a.m. EDT, July 16, 2005.

The article begins, "Bastille Day is the French national holiday, celebrated on 14 July each year" and provides some background. This information can be verified from numerous sources using a quick

Google search on 'Bastille Day'. Much of the background and information is substantiated by other sources.

The article next contains the comment, "Margaret Thatcher once said of the French 'who can trust a people who celebrate, as their national event, a jailbreak'." This statement does not tell us about Bastille Day. It is derogatory to the French. The source of the quotation is not given. This statement may be disregarded as vandalism. (It is worth noting that as of 11:47 a.m. the statement had been removed.) This article, with the exception of the one item noted, is trustworthy.

The Price is Right pricing games

This is a Wikipedia article about The Price is Right. Readers should note that Wikipedia articles frequently change. This article was current at 11:51 a.m. EDT, July 16, 2005.

The article lists a number of 'minigames' played on The Price is Right. Each game is described, with information about when it was played, how frequently it was played, and records, if applicable. Three external sources, including one from CBS, the producer of The Price Is Right, and one with screen shots of the games, are provided.

Readers who have seen The Price is Right can verify the game descriptions for themselves from their own experience. From my perspective (having seen many of the games) this article is very trustworthy.

The Flight of the Bumblebee

This is a video of a person playing Flight of the Bumblebee solo on guitar. The video is sufficiently detailed to show the fingering. The sound is a guitar sound. The tune is recognizable as Flight of the Bumblebee (people who have not heard this piece of music before should consult alternative sources to verify the title. This video is trustworthy.

Moncton, 16 July, 2005

Applying Critical Reasoning

A PLENK participant [wrote](#):

I have found that when I actually apply my critical knowledge skills - questioning the status of procedures, rules and arguments- I pay a price.

It may be that I am ignored or I am excluded. Or judged difficult. But the point is that even when we do enable our students to be users of knowledge - and encourage them to question, the pressure within society - our institutions - effects a price.

The price is usually economic and social exclusion.

Hmmm..What to do then?

I have three major pieces of advice, gleaned from years of hard experience of exactly what you describe here.

1. Pick your battles
2. Find your friends
3. Show results

The first means applying a judicious hand. My experience has been that most discussion in meetings, committees, etc., is unreasonable and illogical. People make decisions based on poor evidence, bad reasoning, bias and prejudice, and worse. It's ongoing and you can't stop it.

And, as you pointed out, these same people will react badly when you attempt to inject a note of reason into the debate. Your considered and fact-based reasoning will be interpreted as a personal attack. No matter what pains you take (and you should!) to depersonalize the matter, to make the subject of debate and not the person the locus of discussion, people will think it reflects badly on them if you disagree with them.

So, pick your battles. Let the small stuff go. A lot of what matters will not have any long term impact. Don't enter into arguments you can't win. Identify what is core for you, what you can't or won't back down on, and raise these issues consistently. You still won't be successful when you raise your points. But your consistency will serve notice to other people that on these points they need to be much more careful before leaping in.

The second applies to picking your friends. It's really hard to be the lone voice of reason at meetings or

in committees, and often, you don't have to be. Be observant and watch for people who apply reason and evidence to their considerations. Make it clear to them that by arguing in this way you will respect their position and potentially modify your position accordingly.

These are people you can work with, and people you should gravitate toward. If your current venue of activity contains no candidates, find alternative venues. Some venues (like, say, staff meetings) cannot be avoided, but if they are unproductive, invest your energy and commitment elsewhere - to subcommittees, special interest groups, professional associations - wherever you can find voices of reason.

You can also find friends by being clear about how you reach opinions, and by (in the parlance of fifth-grade mathematics) showing your work. This will identify you as a person who gives thought and consideration to logic and evidence, and hence identifies you to others as someone who can be worked with.

Finally, third, produce. As someone (Dan Rehak? Wayne Hodgins?) once said to me, "quality ships." It is all very well to have opinions, however well-versed, but the best evidence that you know things, can work through evidence, and evaluate priorities, is by "shipping", or actually delivering results.

Do the work. Write the background papers, do the interviews with clients, draft the policies, write the software, design the deliverable - whatever it is that your group or association does. In my work here at NRC the strongest argument I have for my opinions is that I deliver results. Disagreeing with me or blocking my work has a cost to the organization.

But also, by doing the work, you are providing tangible assistance to the people you are currently disagreeing with. You are providing a set of accomplishments they can hang their hats on, through their association with you. There's an old saying for public speakers, "love your audience," which really works, because if you are focused on how much you want to help and support your audience, you lose all your self-consciousness about speaking in public. It works in the office as well.

The person whose voice is most respected in any group is not the person who leads, or is smartest, or is even right. It is the person perceived by the rest to have the least self-interest, the person who is there to help rather than the person who is there to pursue an agenda or toot their own horn. A minute spent helping the other person achieve their ends (which are often not even in conflict with yours) is far more productive than a minute spent arguing with them.

I hope this helps.

Moncton, November 18, 2010

How Do You Know?

Part One

I have just finished a [presentation](#) to the British Council consisting of a video and a short discussion. I'm not happy with the result - partially because the process of producing the video seemed to be cursed (including one crash that wiped out hours of work - Camtasia has no autosave! who knew?) and partially because I didn't feel comfortable with the discourse.

The video production is one thing, and I can live with a more or less proficient video because it's part of the ongoing process of learning a new way to communicate. But I'm less sanguine about the discourse. I have a sense of what went wrong with it - I even talked about that a bit during the session - but still it nags at me with deeper issues still unresolved.

We weren't very far into the discussion when I made the comment that "if you're just presenting information, online is better than the traditional classroom." The point I was trying to make was that the unique advantage of the classroom is that it enables face-to-face interaction, and that it should be used for that, leaving other things to other people.

And so, of course, someone asked me, "How do you know?" Which stopped me - not because I don't know - but because of the utter impossibility of answering the question.

There are so many differences in community - the different vocabularies we use and the different assumptions we share, for example. For me to express point A in such a way that it will be understood the way I understand it, I need to work through a fair amount of background. But in a session like this - a 20 minute video and a few seconds of discussion - there was no way I was going to be able to accomplish that.

And this carries over to differences in epistemology. The question of 'How do you know' means different things to different people. In some cases, it's not even appropriate - if a football coach instructs a player, the player doesn't say "How do you know" because he knows that the coach isn't set up to answer questions of that sort (he'll say, "I depend on my experience" or some such thing, offering a statement that has no more credibility than the original assertion). In other cases, some sort of process or set of conditions is assumed - and this varies from discipline to discipline, community to community.

In this particular instance I was speaking at a conference on blended learning. So there's a certain perspective that has already been adopted, one that already says that the classroom should not be abandoned. Indeed, the classroom is like the baseline reference, and the role of ICT is to support by being what the classroom cannot be - being available at home, for example, or at midnight, or around the world. ICT is about *enhancing* learning, in the blended learning model. And this picture couldn't be

further from my own model if it tried. For me, it felt like going to a prayer meeting and talking about the role atheism could play in the devotee's life.

You see, from where I sit, blended learning is a bit like intelligent design. It's a way for people to keep hold of their traditional beliefs, to maintain the primacy of the classroom, the primacy of authority in education, the primacy of the information-transfer model of learning, and at the same time (because it's *blended*, you see) to appear as advocates of new learning technologies, including (as was the subject of the conference) Web 2.0. It's faith pretending to be science. While in my world, there is basically no role for the classroom at all. It's irrelevant.

To their credit, they were willing to let me have that, giving me room to reinvent the face-to-face interaction (which [I do believe in](#)) to allow full and proper play for Web 2.0 and ICT in general. But I am still faced with the fundamental questions: how do I explain what I mean, and how do I know (or show I know) it is true?

To take a case in point: I said "if you're just presenting information, online is better than the traditional classroom." What I thought I was making was a straight-forward assertion about the properties of the traditional classroom and the online presentation of information. I wanted to bring this out but found that I didn't have the words.

For example, information is transmitted online at much greater bandwidth than in a classroom. This is partially because a person standing at the front of the room can only speak at a certain speed. The words only come out so fast - and at a fraction of the speed they can be read (at least by most people). And in a classroom the instructor must attend to the needs of all students, which means there will be periods of 'dead air', where one student is being addressed at the expense of everyone else, who must sit and wait.

I wanted to say this, but I couldn't say this, because the audience must already *know* this - and yet, despite this knowledge, will *still* favour classroom delivery, which is why what I thought was a statement of fact - that "if you're just presenting information, online is better than the traditional classroom" - became a statement of opinion, that needed some sort of evidence. From *my* perspective, it was as though I had said "the sky is blue" and someone (who apparently believe there was no sky) asked my how I knew. How do you explain? How do you *argue*?

What could 'better' even *mean* in such a context?

Because my own statement - that "if you're just presenting information, online is better than the traditional classroom" - doesn't even make sense in the context of my own theory, because I do not support an information-transfer theory of education. I'm in the position where I'm trying to discuss the relative advantages of online and in-class learning, and trying to place myself into the context of the existing discussion, which works to a certain point, but which vaporizes when pressed in certain ways.

How do I know it is better? Well in this world there are certain outcomes to be expected, and means of measuring those outcomes, so that the relative efficacy of classroom instruction and online instruction could be compared, by conducting pretests and post-tests against standardized evaluations, using standardized curricula. And the best I could say, under such conditions, is that there is [no difference](#), based on 40 years of studies. Which they must know about, right?

All this is going through my mind as I seek to answer the question.

I consider the possibility that by 'better' he means 'more efficient'. Because here I could argue (with some caveats about production methods and delivery, the sort of things I outline in [Learning Objects](#)) that the use of online delivery methods is much cheaper than the very labour-intensive methodology of the classroom. That we are paying, for example, research professors (who don't even *want* to teach) very high salaries to accomplish something that could be as well done using multimedia.

So I concluded that he was looking for evidence of the usual sort - studies that showed knowledge was more reliably transferred (or at the very least, implanted) using ICTs than in classroom instruction. Probably such studies exist (you can find a study to support almost anything these days). But I am again hitting the two-fold dilemma.

First, our conception of the task is different. I had just come from reading and writing about [associative learning](#). "The result in the brain is strengthening or weakening of a set of neural connections, a relatively slow process." It's not about content transfer, it's about repeated exposure (preferably where it is highly salient, as this impacts the strength of the neural connection). The classroom plays almost no role in this; at best it focuses the student's attention, so that subsequent exposure to a phenomenon will be more salient.

This is (as so often happens) abutted directly against corporate or institutional objectives. The fact that trainers and teachers have certain things that they need to teach their students, and that this is generally non-negotiable (to me, this is a lot like the Senate legislating that the value of Pi is 3, but I digress). That evidently, and by all evidence, these objectives *can* be accomplished using classroom instruction, and that moreover, they *might not be* using ICTs.

The evidence, of course, is the set of successful exam results. One would *think*, with the experience of No Child Left Behind behind us, that we would be sensitive to the numerous and multifarious means of [manipulating](#) such results. I have [written before](#) about how such tests can't be trusted. About how the proposition that there can be (so-called) evidence-based policy [should not be believed](#). And I've linked to the [misconceptions](#) people carry with them about this. But [I can't shake](#) in people that belief that there is, after all is said and done, some way to measure whether one or the other is better.

The thing is, there is no definition of 'better' that we could define the parameters for such a

measurement, and even if there were, the determinates of 'better' are multiple and complex. A person's score on a test, for example, is subject to multiple and mutually dependent factors, such that you cannot control for one variable while testing for the others. Any such measurement will build into its methodology the outcome it is looking for.

The problem is - according to everything we seem to know - unless there is some way of *measuring* the difference, there is no way to *know* the difference. Even if we don't believe that "if it can't be measured, it doesn't exist," it must be that measurements give us *some* sense of what is better and what is not - that they can at least approximate reality, if not nail it down precisely. I don't agree - the wrong measurement can suggest that you are succeeding, when you are failing. Sometimes these wrong measurements are deliberately constructed - the phenomenon of greenhouse gas [intensity](#) is a case in point.

At a minimum, this position takes a good deal of background and analysis to establish. At worst, attempting to maintain such a position leaves open the charge of 'charlatan'. Responses like [this](#): "Each time I read a student's paper containing 'I think, I feel, I believe,' I am aggravated, acerbically critical, and given to outbursts of invective: 'Why do I care what you feel?' I write, roaring with claw-like red pen. 'This is not an emotional experience. Believe? Why would you think you can base an argument on unsubstantiated belief? You don't know enough to believe much of anything. Think? You don't think at all. This is mental masturbation. Without evidence you have said exactly nothing!'"

Am I a charlatan when I say things like "if you're just presenting information, online is better than the traditional classroom?" Even if I have nothing to personally gain from such statements, am I leading people down the garden path? It is very difficult, in the face of things like the British Council presentation, to suppose people are thinking anything else. "It's a nice line," they think to themselves as I stumble in front of them, attempting lamely to justify my lack of evidence, "but there's no reason I should believe it."

Which raises the question - why do I believe it?

I have made decisions in my own life. I have chosen this way of studying over that. I have chosen this way of communicating over that. I didn't conduct a study of which way to learn and which way to communicate. I operated by *feel*. There's no way of knowing whether I might not have been more successful if, say, I had stayed in the academic mainstream, published books and papers, assigned my copyrights to publishers, learned through classes and conferences and papers and lectures.

But, of course, *that* was never the decision I made. At no point did I sit down and say, I will eschew traditional academia, I will learn informally, through RSS and Gog-knows-what Web 2.0 technology, and (while I'm at it) I will embrace Creative Commons and lock publishers out of the loop. Indeed, I don't think I could have *imagined* all of that, were we to suppose some fateful day when such a decision would have been made. I made the decision one small step at a time, one small adjustment at a time, as

though I were surfing a wave, cutting, chipping, driving forward, each decision a minute adjustment, each characterized not by measurement, not by adherence to principle, but by *feel*, by *reaction*, by *recognition*.

This is important. George Siemens says that *knowledge* is distributed across the network, and it is, but *how we know* is irreducibly personal.

What does that mean? Well, part of what it means is that when we are actually making decisions, we do not *in fact* consult principles, best practices, statistics or measurements. Indeed, it is even with some effort that we refrain from playing the hunch, in cases where we (cognitively) know that it's a bad bet (and we walk away (and I've had this feeling) saying, "I know the horse lost, but I still should have bet on the gray," as if that would have made the difference).

Malcom Gladwell says, [make snap decisions](#). Trust our instincts. What this means is very precisely an abandonment of principle, an abandonment of measurement, in the making of decisions. It's the same sort of thing. My 'knowing' is the culmination of a lifetime of such decisions. I have come to 'know' that "if you're just presenting information, online is better than the traditional classroom" in this way - even though the statement is, in the context of my own theories, counterfactual. I know it in the same way I know that 'brakeless trains are dangerous' - not by any principle, not by any evaluations of actual brakeless trains, but because I have come to *know*, to recognize, the nature (and danger) of brakeless trains.

We sometimes call this 'the weight of experience'. And this is why my 'knowledge' differs from yours. Not because one of us, or the other, has failed to take into account the evidence. But because the weight of our respective experiences differs.

This gets back to the question of why 'presenting information' will not be 'successful' (let alone 'better') in my view. Recall that I said that the wrong measurement can suggest that you are succeeding, when you are failing. We can present information, and then test students to see if they remember that information. If they are successful on the test, then we say that they 'know' that information.

My experiences with my presentations is different. I can make a presentation - such as, say, today to the British Council - and walk away feeling that while the audience heard me, and while they could probably pass a test (I am a good presenter, after all, even on my bad days, and they are smart people, with exceptional memories), I would not say that they 'know' what I taught them. Wittgenstein says, "Somebody demonstrates that he *knows* that the ice is safe, by walking on it." These participants may leave the conference being able to repeat the words, but scarce any of them will change their practice, eschew the classroom, embrace the world of Web 2.0.

How can I say that they *know* my position, if all they do (all they can do?) is repeat the words? If they 'knew' my position, they would change their practice - wouldn't they? If they had the same *knowledge* I

had - which would have the same weight of experience I had - they they would naturally, without the need for convincing (or even training) make the same decisions I did. Without needing even to think about it. That's what Dreyfus and Dreyfus call '[expert knowledge](#)'. "He does not solve problems. He does not even think. He just does what normally works and, of course, it normally works." And it can't be obtained by measurement, it can't be expressed in principles, it can't be taught as a body of knowledge, and it can't be measured by answers on a test.

A presentation such as the one I gave at British Council this morning (or at CADE a month ago) isn't a transfer of information. People may acquire some words and expressions from me, but they won't acquire knowledge, because even if my presentation were perfect, it could not perform the repetition of instances required in order to create a weight of experience on a certain subject. The best I could do is to repeat a word or phrase over and over, in different ways and slightly different contexts, the way advertising does, or the comedian that kept repeating 'Via' ("Veeeeeeee.... ahhhhhhh").

A presentation is a performance. It is a demonstration of the presenter's expertise. The idea is that, through this modeling - through facility with the terminology, through demonstration of a methodology, through the definition of a domain of discourse (which will be reinforced by many other presentations on the same subject - if you hear Wittgenstein's name often enough, you come to believe he's a genius) - you learn what it is to be 'expert'.

A lecture won't impart new knowledge on older, more experienced listeners at all - it acquires the status of gossip, serving mainly to fill people in on who has been saying what recently, what are the latest 'in' theories or terms. The point of a talk on 'Web 2.0' is to allow people to talk about it, not to result in their 'knowing' it. With younger participants (interestingly the least represented at academic conferences, lest they be swayed by people other than their own professors) the inspiring demonstration of academic expertise serves as a point of departure for a lifetime of similar practices that will, in a generation, result in similar expertise (people did not become disciples of Wittgenstein because they *believed* him - it is very unlikely that they even *understood* him - but by the fact that he could (with a glance, it seemed) utterly demolish the giants in the field of mathematical philosophy).

I have spoken elsewhere about what sort of knowledge this is. It is - as I have characterized it elsewhere - *emergent knowledge*, which may be known by the fact that it is not *perceived* (ie., it is not sensory, the way 'red' or 'salty' are sensory) and it is not *measured*, but by the fact that it is *recognized*. It is a 'snapping to' of awareness, the way we see a [duck](#) (or a rabbit) or suddenly discover [Waldo](#).

'Recognition', in turn, amounts to the exciting of a set of connections, one that is (relevantly) similar to the current content of perception. It is a network phenomenon - the activation of a 'concept' (and its related and attendant expectations) given a certain (set of) input condition(s). When we present certain phenomena to the network, in the form of a set of activations at an 'input layer' of neurons, then based on the set of existing connections in the network, some neurons (and corresponding connections) are activated, while others remain silent; this present experience (sometimes) produces a response, and (in

every case) contributes to the set of future connections (one connection is subtly strengthened, another subtly weakened).

When presented with a certain set of input phenomena, you can *remember* - to certain degree. If given sufficient motivation, you can associate certain noises (or certain shapes) with each other. On being told, I can remember that 'Paris' is the 'capital' of 'France', and even repeat that information on a test (and moreover, remember who said it to me, and when, and under what circumstances), but I cannot be said to *know* unless I demonstrate (a disposition?) that if I want to see the President of France, that then I go to Paris. And this is not the sort of thing that is on a test - it is a sort of thing that allows a person to have 'learned' that Sydney is in Australia, and even how to book an airline ticket to Sydney, and not notice that they are [traveling to Canada](#).

How do I know? Because - by virtue of my experiences with traditional and online settings - if I were trying to support *knowledge* in a person, I would not turn to the classroom, but rather, some sort of practice, and even if I were (because of policy or the demands of corporate managers) trying to support *remembering* in a person, I would contrive to have it presented to them, over and over, in the most efficient and ubiquitous means possible, which today is via ICTs.

How do *you* know whether to believe me?

You don't. Or, more accurately, there is nothing I can provide you that will convince you to believe me if you are not already predisposed to believe me. The best I can do is to suggest a course of action (ie., a set of experiences that you can give yourself) such that, after these experiences, you will come to see the world in the same way I do. That is why my talk to the British Council (and to many other audiences) described just that, a set of *practices*, and not a set of theorems, or experimental results, or the like.

The practices I presented constitute (one way of describing) the practices I undertake in my own learning and development. The evidence, then, of whether these practices is in whether you believe that I have demonstrated my expertise. This, in turn, depends on your own sense of recognition - some people will recognize that I have achieved a certain degree of expertise, while others will leave the room with the verdict of 'charlatan'.

And what follows is a subtle dance - the [connectivism](#) George Siemens talks about - where you demonstrate your expertise and I demonstrate mine - and where each of us adopts some of the practices of the others (or rejects them, as the case may be) and where the connections between people with similar practices is reinforced, and knowledge demonstrated in such a community not by what it says (hence the fate of critical theory) but by what it does. This is the process (and I have explained [elsewhere](#) the properties of the network that will grant the process some degree of reliability).

Part Two

These are hard questions.

What I described [in my paper](#) is my best answer to the question 'how do I know', that is, it tries to explain how I (in fact) know things. It is therefore not a description of the criteria I should use to distinguish truth from falsity, nor how one person can convince another person of something.

Indeed, viewed as a system for determining the truth of something, the paper seems pretty ridiculous. Wealth of experience? Why should anyone trust that! Why is my wealth of experience any better than anyone else?

The problem is, the description of how we *in fact* learn things does not carry with it any sort of guarantee that what we've learned is true. But without such a guarantee, there can be no telling for ourselves what to believe or not to believe, no way to convince other people. It's like we're leaves blown about by the breeze, with no way to sway the natural forces that affect us.

Moreover, the problem is:

- the *is* no guarantee
- yet we *do* distinguish between true and false (and believe we have a method for doing so)
- and we *do* want to be able to sway other people

What complicates the matter - and the point where I deliberately stopped in the other paper - is that not everybody is honest about what they know and what they don't know. Sometimes there genuinely *are* charlatans, and they want to fool us. Sometimes they are simply mistaken.

There's not going to be a simple way to step through this.

I went immediately from the British Council talk, where I was trying to foster a point of view, to a session inside Second Life, where I played the role of the sceptic. Not that I think that the people promoting Second Life are charlatans. But I do think they are mistaken, and I do think some of the statements they make are false.

The fact is, even though there are no guarantees, we will nonetheless make judgments about truth and falsehood. It is these judgments - and the manner of making these judgments - that will sway the opinions of other people.

You can't *tell* people things, you can only *show* them.

Now even this statement needs to be understood carefully. It is true that we can tell people things, eg., that 'Paris is the capital of France' that they will remember - but it does not follow that they *know* this; they will need to see independent evidence (such as, say, newscasts from 'the capital of France'). Telling produces memory. Showing (and experience-producing processes in general) produces belief.

But now - even this needs to be qualified. Because if you *tell* something to somebody enough times, it becomes a type of proxy experience. So - strictly speaking - you *can* produce belief by telling - but not by 'telling' as we ordinarily think of it, but by a repeated and constant telling.

Additionally, we can make 'telling' seem more like experience when we isolate the person from other experiences. When the 'telling' is the only experience a person has, it becomes the proxy for experience.

It is worth noting that we consider these to be illicit means of persuasion. The former is propaganda, the latter is indoctrination. Neither (admittedly) is a guaranteed way of changing a person's mind. But it is reliable enough, as a causal process, that it has been identified and described as an illicit means of persuasion.

Let me return now to how we distinguish truth from falsehood.

This is not the same as the process of coming to know, because this process has no such mechanism built into it. The way we come to know things is distinct from the way we distinguish between truth and falsehood.

This may seem counter-intuitive, but I've seen it a lot. I may be arguing with someone, for example, and they follow my argument. "I agree," they say, "Um hm, um hm." But then I get to the conclusion, and they look at me and say, "But no..." It's this phenomenon that gives people the feeling they've been tricked, that I've played some sort of semantic game.

So there are processes through which we distinguish truth and falsity. Processes through which we (if you will) construct the knowledge we have. We see the qualities of things. We count things. We recognize patterns in complex phenomena. These all lead us, through a cognitive process, to assert that this or that is true or false.

Usually, this cognitive process accords with our experience. For example, we say that the ball is red because we saw that the ball was red. We say that there were four lights because we saw four lights. It is close enough that we saw that we came to know that the statement was true *because* of the experience. But - again - the process of knowing is separate from the process of distinguishing truth from falsehood.

There are general principles of cognition. These are well known - the propositional logic, mathematics, categorical logic, the rest (and if you want to see the separation between 'knowing' and 'distinguishing truth from falsity' then look at some of these advanced forms of logic - deontic logic, for example. We can use some such process to say that some statement is true, but because the process is so arcane to us, the statement never becomes something we 'know' - we would certainly hesitate before acting on it, for example).

There are also well known fallacies of cognition. I have documented (many of) these on my fallacies site. It is interesting to note that these are for the most part fallacies of distraction. What they do is focus your attention on something that has nothing to do with the proposition in question while suggesting that there is a cognitive link between the two. You come to 'know' something that isn't true, because you have had the experience.

Consider the fallacy: "If the plant was polluting the river, we would see the pollution. And look - we can see the pollution." We look, and we see the pollution. It becomes part of our experience. It becomes the reason we 'know' that the plant is polluting the river. No amount of argument - no amount of 'telling' (except, say, indoctrination) will convince us otherwise. We have to actually go to the plant and see that it is not polluting the river in order to understand - to *know* - that we were the victim of a fallacy.

There is a constant back-and-forth being waged in all of us, between what we 'know' and the things we say are 'true and false'.

That is why I say you can't 'tell' a person something. Merely convincing them (even if you can) to agree that 'this is true' is a long way from getting them to *know* it - getting them to believe it, to act on it, to make wagers on it.

So - convincing a person comes down to *showing* them something.

Often this 'showing' will be accompanied with a line of reasoning - a patter - designed to lead them to the 'truth' of what they are being shown. But the knowledge comes from the showing, not the patter.

Even with showing, there are no guarantees. 'You can lead a horse to water...' Even the experience may not be sufficient to convince a person. Any experience is being balanced against the combined weight of other experiences (perhaps the 'patter' is sufficient to sway people in some touch-and-go cases, by offering a coherence with other experiences - an easy path for belief to follow).

A great deal depends on the nature of the experience. Experiences can be vivid, can force themselves on us. They can be shocking or disturbing. Images of violence capture our attention; images having to do with sex capture our attention. Our attention, even, can be swayed by prior experiences - a person who has spent a lifetime around tame tigers will react very differently on seeing a tiger than a person who has only known them to be dangerous carnivores.

'Convincing' becomes a process of pointing, a process of showing. Sometimes what a person is told can direct a person where to look (in this piece of writing I am encouraging to look at how you come to have your own knowledge, to see how it is the result of a separate track from how you come to see things are true and false). Sometimes the experiences can be contrived - as, say, in a simulation - or the senses fooled. Some media - especially visual media - can stand as substitutes for experience.

We can have experiences of abstract things - the weight of experience just is a way of accomplishing this. The logical fallacies, for example - on being shown a sufficient number of fallacies, and on seeing the *fallaciousness* of them, we can come to have a knowledge of the fallacies - such that, when we experience a similar phenomenon in the future, we experience it *as fallacious*.

Convincing becomes a matter of showing, showing not just states of affairs in the world, but processes of reason and inference. If I can show actual instances of inference, how a person comes to believe, comes to know, this or that, then it becomes *known*, and not merely believed, by the viewer. If I can show my reasoning *process*, then this process can be known (after being experienced and practiced any number of times) by the learner.

'Expert knowledge' is when a person not only remembers something, but when a person has come to know it, has come to know the processes surrounding a discipline. Such knowledge is often ineffable - the knower can't even enumerate the (true or false) statements that constitute the knowledge, or that led to the knowledge. What a person *knows* is distinct from what a person *says* is true or false. It is not truth that guarantees knowledge. It is knowledge that guarantees truth.

Moncton, 19 June, 2007

Having Reasons

Semantics is the study of meaning, truth, purpose or goal in communication. It can be thought of loosely as an examination of what elements in communication 'stand for'.

Because human communication is so wonderfully varied and expressive, a study of semantics can very quickly become complex and obscure.

This is especially the case when we allow that meanings can be based not only in what the speaker intended, but what the listener understood, what the analyst finds, what the reasonable person expects, and what the words suggest.

In formal logic, semantics is the study of the conditions under which a proposition can be true. This can be based on [states of affairs](#) in the world, the meanings of the terms, such as we find in a [truth table](#), or can be based on a [model or representation](#) of the world or some part of it.

In computer science, there are [well-established methods](#) of constructing models. These models form the basis for representations of data on which operations will be formed, and from which views will be generated.

David Chandler [explains](#) why this study is important. "The study of signs is the study of the construction and maintenance of reality. To decline such a study is to leave to others the control of the world of meanings."

When you allow other people to define what the words *mean* and to state what makes them *true*, you are surrendering to them significant ground in a conversation or argument. These constitute what Lakoff calls a "[frame](#)".

"Every word is defined relative to a conceptual framework. If you have something like 'revolt,' that implies a population that is being ruled unfairly, or assumes it is being ruled unfairly, and that they are throwing off their rulers, which would be considered a good thing. That's a frame."

It's easy and tempting to leave the task of defining meanings and truth conditions to others. Everyone tires of playing "semantical games" at some time or another. Yet understanding the tools and techniques of semantics gives a person tools to more deeply understand the world and to more clearly express him or her self.

Let me offer one simple example to make this point.

We often hear people express propositions as probabilities. Sometimes these are very precisely expressed, as in the form "there is a 40 percent probability of rain." Other times they are vague. "He probably eats lettuce for lunch." And other times, probabilities are expressed as 'odds'. "He has a one in three chance of winning."

The calculation of probability can be daunting. Probability can become [complex](#) in a hurry. Understanding probability can require understanding a [probability calculus](#). And there is an endless supply of related concepts, such as [Bayes Theorem](#) of prior probability.

But when we consider the *semantics* of probability, we are asking the question, "on what are all of these calculations based?" Because there's no simple answer to the question, "what makes a statement about probabilities true?" There is no such thing *in the world* that corresponds to a "40 percent chance" - it's either raining, or it's not raining.

A semantics of probability depends on an *interpretation* of probability theory. And there are some [major interpretations](#) you can choose from, including:

1. The logical interpretation of probability. Described most fully in Rudolf Carnap's [Logical Foundations of Probability](#), the idea at its heart is quite simple. Create 'state descriptions' consisting of all possible states of affairs in the world. These state descriptions are conjunctions of atomic sentences or their negations. The probability that one of these state sentences is 'true' is the percentage of state descriptions in which it is asserted. What is the possibility that a dice roll will be 'three'? There are six possible states, and 'three' occurs in one of them, therefore the probability is 1 in 6, or 16.6 percent.
2. The frequentist interpretation of probability. Articulated by [Hans Reichenbach](#), the idea is that all frequencies are subsets of larger frequencies. "Reichenbach attempts to provide a foundation for probability claims in terms of properties of sequences." This is the basis for inductive inference. What we have seen in the world in the past is part of a larger picture that will continue into the future. If you roll the dice enough times and observe the results, what you will discover (in fair dice) that the number 'three' appears 16.6 percent of the time. This is good grounds for expecting the dice to roll 'three' at that same percentage in the future.
3. The subjectivist interpretation of probability. Articulated by Frank Ramsay, "The [subjectivist theory](#) analyses probability in terms of degrees of belief. A crude version would simply identify the statement that something is probable with the statement that the speaker is more inclined to believe it than to disbelieve it." What is the probability that the dice will roll 'three'? Well, what would we bet on it? Observers of these dice, and of dice in general, would bet one dollar to win six. Thus, the probability is 16.6 percent.

Each of these interpretations has its strengths and weaknesses. And each could be expanded into more and more detail. What counts, for example, as a 'property' in a state description? Or, what are we to make of irrational gamblers in the subjectivist interpretation?

But the main lesson to be drawn is two-fold:

- first, when somebody offers a statement about probabilities, there are different ways of looking at it, different ways it could be true, different meanings we could assign to it.
- and second, when such a statement has been offered, the person offering the statement *may well be assuming* one of these interpretations, and expects that you will too, even in cases where the interpretation may not be warranted.

What's important here is not so much a knowledge of the details of the different interpretations - first of all, you probability couldn't learn all the details in a lifetime, and second, most people who make probability assertions do so without any knowledge of these details. What is important to know is simply *that they exist*, that there are different *foundations* of probability, and that any of them could come into play at any time.

What's more, these interpretations will come into play not only when you make statements about the probability of something happening, but when you make statements generally. What is the foundation of your belief?

How should we interpret what you've said? Is it based on your own analytical knowledge, your own experience of states of affairs, or of the degree of certainty that you hold? Each of these is a reasonable option, and knowing which of these motivates you will help you understand your own beliefs and how to argue for them.

Because, in the end, semantics *isn't* about what some communication 'stands for'. It is about, most precisely, what you believe words to mean, what you believe creates truth and falsehood, what makes a principle worth defending or an action worth carrying out.

It is what separates you from automatons or animals operating on instinct. It is the basis behind having reasons at all. It is what allows for the *possibility* of having reasons, and what allows you to regard your point of view, and that of others, from the perspective of those reasons, *even if* they are not clearly articulated or identified.

The whole concept of 'having reasons' is probably the deepest challenge there is for connectivism, or for

any theory of learning. We don't want people to simply to react instinctively to events, we want them to react on a reasonable (and hopefully rational) basis. At the same time, we are hoping to develop a [degree of expertise](#) so natural and effortless that it seems intuitive.

Connectivist theory is essentially the idea that if we expose a network to appropriate stimuli, and have it interact with that stimuli, the result will be that the network is *trained* to react appropriately to that stimuli. The *model* suggests that exposure to stimuli - the conversation and practices of the discipline of chemistry, say - will result in the creation of a distributed *representation* of the knowledge embodied in that discipline, that we will literally *become* a chemist, having internalized *what it is to be* a chemist.

But the need to 'have reasons' suggests that there is more to becoming a chemist than simply developing the instincts of a chemist. Underlying that, and underlying that of any domain of knowledge, is the idea of being an *epistemic agent*, a knowing knower who knows, and not a mere perceiver, reactor, or doer. The having of reasons implies what Dennett calls the [intentional stance](#) - an *interpretation* of physical systems or designs from the point of view or perspective of reasons, belief and knowledge.

We could discuss the details of having and giving reasons until the cows come home (or until the cows follow their pre-programmed instinct to follow paths leading to sources of food to a place designated by an external agent as 'home'). From the point of view of the learner, through, probably the most important point to stress is that they *can* have reasons, they *do* have reasons, and they should be reflective and consider the *source* of those reasons.

Owning your own reasons is probably the most critical starting point, and ending point, in personal learning and personal empowerment. To undertake personal learning is to undertake learning for your own reasons, whatever they may be, and the outcome is, ultimately, your being able to articulate, examine, and define those reasons.

Moncton, Tuesday, July 06, 2010

How Memory Works

This is a summarization of a paper by Eric R. Kandel on the molecular and synaptic basis for memory, [Genes, synapses and memory storage](#). Kandel won the 2000 Nobel Prize for this work. I was moved to write this after listening to a segment of [CBC's Ideas program](#) discussing the nature of learning and memory. At the end of the paper, I draw the inferences from Kandel's work to my own.

The problem of memory has two major parts:

The *systems* component, which concerns "where in the brain memory is stored and how neural circuits work together to create, process, and recall memories. "

The *molecular* component which studies "the mechanisms whereby synapses change and information is stored"

The systems component - a history:

1865 - Pierre-Paul Broca identifies speech production with a specific area of the brain.

1876 - Carl Wernicke identifies language comprehension with a *different* area of the brain and suggests that complex behaviour requires the interaction of different brain areas.

1929 - Efforts to localize memory fail; Karl Lashley formulates the Law of Mass Action: "the extent of a memory deficit is correlated with the size of a cortical lesion but not with the specific site of that lesion."

1938 - Wilder Penfield localizes specific memories in epileptic patients (this was the subject of a ['Heritage Minute' video](#) in Canada - "I smell toast burning").

1957 - Scoville and Milner localize memory formation in the medial temporal lobe and show there are multiple, functionally specialized memory systems in the brain.

The idea that there are multiple memory systems in the brain has a long history in the philosophy of psychology:

early 1800s - French philosopher Maine de Biran argues memory can be subdivided into different systems for ideas, feelings and habits

early 1900s - William James divides memory into distinct temporal phases

1913 - Henri Bergson distinguishes between conscious memory and habit

1949 - [Gilbert Ryle](#) distinguishes between 'knowing that' and 'knowing how'

(1956 - [Michael Polanyi](#) - tacit knowledge (this isn't mentioned by Kandel))

1969 - [Jerome Bruner](#) describes 'knowing that' as a memory with record and 'knowing how' as a memory without record

Scoville and Milner's studies of H.R., a patient who had the medial temporal lobe removed, yielded three major findings:

There was a short-term memory unaffected by the loss of other memory functions.

There was a long-term memory of events prior to the operation.

H.R. could form *some* long-term memories after, but denied doing so.

This established the distinctions postulated by the philosophers. This distinction between types of long-term memory is now characterized using the terms:

implicit - corresponding to 'knowing how', is habitual, unarticulated, and not recorded

explicit - corresponding to 'knowing that', is cognitive, articulated and recorded

The molecular component



Kandel started by looking at the hippocampus but decided to focus on the simplest possible case, the marine snail [Aplysia](#).

Why study Aplysia?

- It is smart (for a snail) - it can create both short-term and long-term memories
- It is simple - it has only 20,000 neural cells
- Then neural cells are quite large, and hence easy to study
- It is possible to map in detail the synaptic connections between cells with each other and with sensory and motor systems.

What they found (this is the key finding):

Short-term storage for implicit memory involves functional changes in the strength of *pre-existing synaptic connections*.

Long-term storage for implicit memory involves the synthesis of new protein and the *growth of new connections*.

This protein synthesis required to convert from short-term to long-term memory was developed early in evolution and hence preserved through all life forms, and is a general mechanism, responsible for both explicit and implicit memories.

Learning in pre-existing synaptic connections

Let's look at this in detail:

There are two major types of conditioning:

habituation - an animal perceives a sensation as innocuous and ignores it

sensitization - an animal perceives a sensation as noxious and tries to defend itself or flee


And two forms of learning:

non-associative - an animal habituates or sensitizes to a single stimulus

associative - an animal habituates or sensitizes to a pair of unrelated stimuli

In order to understand how the animal learns, therefore, "one needs in particular to work out the pathway whereby the sensory stimulus of the reflex leads to a behavioral response."

In the short term, habituation is represented by the weakening of the synaptic connection, and the resulting decrease in the release of glutamate, while sensitization is represented by the strengthening of the synaptic connection, and the corresponding increase in the release of glutamate.



Kandel doesn't include the diagram [from Mann](#) at right in the paper, but it nicely illustrates the process. The little blue dots represent the release of glutamate.

Kandel's description of this process (pp. 34-35) provides the chemical basis for [Hebbian \(associative\) learning](#):

"Two events need to happen simultaneously: glutamate needs to bind to the postsynaptic nmda receptor, and the postsynaptic membrane needs to be depolarized substantially... This coincident activation of the nmda receptor and postsynaptic depolarization only occur when the weak siphon stimulus (cs) and the strong tail shock (us), are paired together."

Three major lessons are drawn from this work:

1. learning can lead to changes in the strength of connections (synaptic strength)
2. a single connection can participate in several types of learning
3. each of the three simple types of learning - habituation, sensitization and classic conditioning - gives rise to both short-term and long-term memory, depending on the number of repetitions


The growth of new connections

History of the distinction between short-term and long-term memory:

1885 - Herman Ebbinghaus identifies two phases while learning nonsense syllables

1941 - Zubin and Barrera, 1941 note the distinction in people hit in the head

1960s - Louis Flexner and his colleagues identify a biochemical difference between them; long-term memory requires the synthesis of new protein during the consolidation phase



What's important is that there is a *genetic basis* for both the synthesis of the protein and for the consolidation phase.

Kandel notes, "Aplysia and Drosophila [a type of fruit fly] share some of the same genes and proteins for converting short- to long-term memory... creb has a role in learning in Drosophila that is similar or identical to its role in Aplysia, demonstrating striking evolutionary conservation. "

The mechanisms through which the proteins - [CREB-1](#) and [CREB-2 \(aka ATF-2\)](#) - interact with the nucleus are complex and diagrammed ([from Mann](#)) at right.

In combination with other factors (such as, in the fruit fly, the the loss of a cell adhesion molecule), the interaction with the nucleus stimulates genes that results in the production of new synaptic connections.

Explicit memory storage

Explicit memory is more complex because:

- it involves conscious participation in the memory recall
- it doesn't depend on a simple stimulus; it usually depends on several sensory cues

Based on studies of mice, the [hippocampus](#) appears to play a major role in explicit memory. The hippocampus is basically a set of interconnected neural cell fields. It acts as a clearing-house for sensory input. Plasticity (the growth of new connections) has been discovered at all levels of the hippocampus. And the creb proteins appear once again to be implicated in the production of new connections.

Other work has demonstrated the plasticity of sensory systems. For example, experiments in kittens have demonstrated plasticity in the visual system. Cortical plasticity has also been demonstrated in adult monkeys. "These several studies suggest that long-term memory storage lead to anatomical changes in the mammalian and even the human brain much as it does in Aplysia."

A good example is the work done correlating the growth of synaptic connections and *place memory*. There are cells in the hippocampus, called pyramid cells, that are place cells - they fire when we occupy a certain place in our environment. So these cells form a cognitive 'map' of the environment. Various manipulations can lead to *remapping*, in which all the place cells change.

Consequences

a. Learning and Memory

At this point we reach the end of Kandel's paper. What are we to make of these discoveries? What lessons should we draw?

For me, it requires a clarification of a comment that I have made on several occasions recently: *learning is not memory*. Kandel does draw a distinction (p. 31): "Learning refers to the acquisition of new

information about the world and memory refers to the retention of that information over time." But what does that mean?

Learning is a *semantic* process. It is *about things*. It has *meaning*.

Memory is a *syntactic* process. It is a set of mechanisms. It may or may not have meaning.

This is a difficult distinction because the two are so frequently found in the same location. Pyramid cells, for example, that contain a 'map' of the environment, are created through a process of remembering, as a result of the changes of synaptic connections in the hippocampus, but also represent (via the sensory impressions that cause those changes) distinct places in the environment.

Nonetheless, the two are not the same. It should be clear from this work that it is possible to create memories that have no semantic content. It should be clear that through manipulations of the physical process we can create *meaningless* memories.

This, in turn, tells us a lot about the *reliability* of synaptic networks, and hence, of networks in general. In reliable networks, the mechanisms that cause the creation of connections between neurons are *meaning-preserving*, that is, they represent *memories*, and not merely manipulations of the process.

(I am being careful about how I state this, because there will be different accounts of what constitutes 'meaning-preserving').

This suggests:

1. Approaches to testing that test for learning, and not merely memory: such testing will be *individual-centric* (like the environment maps in the hippocampus) and not standardized (which is more likely to reflect syntactic manipulations).
2. Approaches to teaching which are based on creating semantic connections with the world, through the production of meaningful experiences, rather than syntactic manipulations of memory, such as memorization and rote

But this needs to be studied further. What constitutes meaning-preservation? It is *not* (as I'll show below) *truth*-preservation. But what is it? How do we measure for meaning, and not just syntactic compliance? Can knowing *how* learn help us determine *what* we learn?

b. Practice and Reflection

Again, as noted previously, learning is the result of repeated experiences of the same (or similar) type; the neural connections required for long-term memory will not be created without this repetition.

Learning is therefore not simply the *presentation* of information to an individual. It is not simply the transfer of a fact from one person to another. At best, this process could create only a short-term memory. In order to activate the neural connections necessary we need to stimulate the production of creb proteins, which happens only through repetition.

Advertisers, of course, know this, which is why they repeat brand names, jingles and phone numbers over and over. Seasoned politicians also know this, which is why the best orators employ catching phrases that will be repeated over and over, as in the video [Yes We Can](#) (maybe one of the best political advertisements ever).

As I have said before, *learning is not content*. Learning is something *over and above* the presentation of semantically meaningful information to a person. To learn, one does not simply 'acquire' content, one *grows*. To learn is a *physical* act, not a merely mental act.

Again, though, we want to look at this more closely. For example, what constitutes a repetition?

For example: the need for repetition would seem to suggest that a lecture would be a poor form of teaching, since it does not produce repetition. But:

- Can we style lectures such that the repetition is contained in the lecture?
- Can people listening to lectures create repetition through the use of different modalities, such as taking notes, live-blogging or summarizing?
- Can we create repetitions through the conduct of lecture-related tasks, such as projects or problem-solving based on the contents of lectures?
- Does *learning for ourselves* stimulate the production of the repetitions required for memory?
- Is there a connection between *semantic content* and repetition - does learning in authentic contexts increase the probability of remembering?

I would suggest that the answer to each of these questions is 'yes'. But they are the sorts of things that bear further investigation.

c. The nature of knowledge and inference

There is a persistent school of thought in both the philosophy of psychology and also in educational theory that suggests that cognition is based on logical and linguistic rules, that there is a *logical syntax* that governs learning and cognition.

Examples of this range from the postulation of Chomsky's [deep grammar](#) to Fodor's [language of thought](#) to Hempel's [H-D model](#) of the sciences. The proposition is essentially that meaning-preservation is tantamount to truth-preservation, where truth-preservation is as is well understood from logic and mathematics.

But what we learn here is that learning is *associative*, not propositional. That the mechanisms that

govern this process are not expressions of truth-preservations, but are - *at best* - expressions of meaning preservation, where meaning has to do with sensory perceptions and states of affairs in the environment rather than abstract principles of logic and mathematics.

I have expressed this in the past as follows:

Our old understanding of logic, inference and discovery is based on *universals*:

- rules, such as rules of inference, or natural laws
- categories, such as classifications and taxonomies

Our new understanding, through, is based on *patterns recognition*:

- patterns, such as the activations of similar sets of neurons
- similarities, such as the perception of similar properties in nature

That is not to say that these universal principles play no role in our understanding. It means, rather, that we need to see them in a new light:

These principles represent 'convenient fictions', not underlying principles of nature

These principles are *learned* - they are not innate

There's a lot more work to be done here. The nature of inference based on patterns and similarities is poorly understood. It is one thing to say things like 'an understanding of learning based on simple causation is mistaken' and quite another to describe the complex mechanisms that actually occur.

We need to dig into the logic of similarity, following the work of people like [Tversky](#) and [Varela](#), to conjoin this with our understandings of [social network theory](#) and [graph theory](#).

Moncton, February 09, 2008

How The Net Works

Originally published in CEGSA RAMpage Magazine, and as a summary of my talk given to CEGSA, '[How the Net Works](#)'.

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 transmission 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](#).

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.

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."

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. ([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 - possibly most of what we know - is 'tacit'. That means it is 'ineffable'. 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.

- knowledge is also irreducibly 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.

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](#)' with [Tarski semantics](#).)

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. 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). Or any of the numerous 'out of the blue' experiences described by [Tom Haskins](#).

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](#)'. 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](#)'. We [need presence](#) to help develop cognitive skills and to feel the supportive environment that supports growth. 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 than 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 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, if 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 current sites. What you won't find are things like dropdown menus, fancy icons, image maps and the other arcana 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 through the development of networks in computer science, an approach called '[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.

- [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 instructing 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.

- 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 or 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 practice is to improve performance by receiving feedback. They then reflect on what they have experienced and practiced - this is (somewhat) analogous 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](#). "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](#)." It is this wisdom we see not only in the audience 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](#). They are like 'jumping to a conclusion' 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, case studies and stories
- Things that help us practice - such as games, sandboxes, job aides, simulations and environments
- 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 its 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](#) communication 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.

Moncton, October 18, 2007

An Operating System For The Mind

The core of the [opposition](#) to what are being called "21st century skills" is contained in the following argument: "Cognitive science teaches us that skills and knowledge are interdependent and that possessing a base of knowledge is necessary to the acquisition not only of more knowledge, but also of skills. Skills can neither be taught nor applied effectively without prior knowledge of a wide array of subjects."

In response, I [pose](#) this question to the defenders of this 'base of knowledge', "why is a common core necessary for the teaching of skills, and why is testing of that core necessary." And specifically, "the question isn't whether skills can be taught in isolation, but rather whether they must be taught in the context of some common base of knowledge and whether students ought to be tested on the basis of that knowledge."

The point I am making may seem difficult to understand intuitively, because it seems to suggest that you don't need deep knowledge of any discipline. A commenter, for example, says, "we want them to think critically, to criticize, analyze, and apply. So we say, draw on the theories of developmental psychology, in particular Erikson and Arens to learn more about individuals in their early twenties. A Based on this understanding, develop a plan to deal with risk taking behavior in junior colleges. Don't they need to first have a 'base of knowledge' in who these theorists are?"

The argument - and it's a reasonable argument - is basically, "you need to know psychology to do psychology; 21st century skills don't give you some kind of short-cut to being able to do psychology without some sort of deep understanding of the subject." And the 'core knowledge' people can take their argument a step further: you can't learn psychology without first learning various other things - the knowledge of psychology *builds* on these things, and that's why we need a grounding in core knowledge.

This form of argument is very common, and you'll find it repeated over and over - people need to know about bones to study medicine, people need to know about the elements to study chemistry, people need to know about history to study politics. Stated this way, the argument seems plausible, and the people promoting 21st century skills look like shysters, promoting something that will leave people unable to work in any discipline, let alone become psychologists, scientists and engineers.

In response to this line of reasoning, let me be upfront about saying a few things:

First, it isn't impossible to teach people facts. Quite the opposite is the case - we understand, and can prove (and have proved, over and over) that we can teach facts very simply and easily, through repetition, rote, memorization, practice examples, worked examples, and more. People can memorize

the alphabet, the multiplication tables, the Koran, whatever. A great deal of our education today in fact turns on this very proposition: it consists of the teaching of facts, and the testing for recall of those facts.

Second, it isn't wrong to teach facts. Or (perhaps more accurately) to learn facts. Having an easy memory recall of a body of facts will serve a person well in life. Knowing the multiplication tables, knowing the capital of France, knowing that carbon and hydrogen are elements (and that plastic is not) will be useful in a wide range of areas. Teaching children facts is a great shortcut, *the* great shortcut, in human development.

Third, we need facts to do stuff. We need to know about psychology, about Freud and Jung and maybe Erikson and Arens, in order to do the job. We need to know about navigation and aerodynamics and where the brake lever is in order to fly an airplane. As anyone who drive knows, you need to know the rules of the road, the meaning of signs, the location of the steering wheel, in order to drive. To do *anything*, you need to *know stuff*.

Not only do I make these statements, I would say that any person who is an advocate of 21st century learning also makes these statements. I can't imagine anyone seriously proposing any sort of educational reform who does not agree with these statements. This is important, because it means *it isn't sufficient* to respond to advocates of 21st century skills by saying 'we need facts'. Everybody has already agreed with that. That's why I pose my question, above, more precisely: do we need *these specific* facts? Do we need a *common core*?

The reason I pose these questions in particular is that, while it is necessary (and possible) to teach facts to people, it comes with a price. And the price is this: facts learned in this way, and especially by rote, and especially at a younger age, take a direct route into the mind, and bypass a person's critical and reflective capacities, and indeed, become a *part* of those capacities in the future.

When you teach children facts *as facts*, and when you do it through a process of study and drill, it doesn't occur to children to question whether or not those facts are true, or appropriate, or moral, or legal, or anything else. Rote learning is a short circuit into the brain. It's direct programming. People who study, and learn, that $2+2=4$, *know* that $2+2=4$, not because they understand the theory of mathematics, not because they have read Hilbert and understand formalism, or can refute Brouwer and reject intuitionism, but because *they know* (full stop) $2+2=4$.

I used the phrase "it's direct programming" deliberately. This is an analogy we can wrap our minds around. We can *think* of direct instruction as being similar to direct programming. It is, effectively, a mechanism of putting content into a learner's mind as effectively and efficiently as possible, so that when the time comes later (as it will) that the learner needs to use that fact, it is instantly and easily accessible.

Interestingly, that's how many people used to think of electronic information systems - as mechanisms

into which you input facts to facilitate easy discovery and retrieval. The computer, or online systems (such as Minitel) were visualized as giant electronic libraries, with the sum total of the world's knowledge at our disposal. The computer was thought of as some sort of electronic book, radio or TV. A place where we could, whenever we wanted, get facts. Put the facts into the system, and come up with some mechanism to get them out. Like Newell and Simon's "General Problem Solver" of 1957.

We know now - and, indeed, have probably always known - that an education based strictly and solely in facts is insufficient. The reasons are legion, but I will focus on six major points:

First. There are more facts in the world than anyone could know, which means that we need to be able to *find facts* that we do not already know. This is the first facet of literacy, the ability to read, view or listen (etc).

Second. As time passes, facts change, and so we need the capacity to know when *facts change* and to be able to update our own knowledge of these facts. We need to be able to learn - that is, to *change* the previously existing state of our knowledge.

Third. And as the number of people, and the amount of information, in the world increases, we need some mechanism for selecting which facts we will be exposed to, and how to filter out *irrelevant facts*. We need to be able to determine what is *salient* or *important* to ourselves and to others.

Fourth. Even more critically, not every bit of information presented to you in life will be a fact, and you need some mechanism to detect and reject *false representations of facts*. We need, in other words, some mechanism for *comparing* and *assessing facts*.

Fifth. Additionally, we need to know which, of the many facts we have in our possession, constitute a *basis for action*. We need some sense of, and mechanism for, *agency* in the world, a sense that we can not only receive, input and assess facts, but that we can *create facts* in the world.

Sixth. Finally, we need *the capacity to act*, which may mean some physical activity, or may mean some communicative activity, a set of abilities we can place under the heading of *empowerment*, as they involve not only the physical capacity to undertake an act, the knowledge which informs that act, but also the willingness to undertake it, the belief that one is entitled to act, and the faith that one's acts can have an impact on the world.

These six elements (and there may be more) constitute what people have taken to calling '21st century skills'.

One might ask: why '21st century'? Because, after all, haven't these skills already been important? Haven't we always had to have the capacity to learn, assess, and act? And, of course, we have. But what has happened recently is, first, there has been a proliferation of new skills in these areas that arise as a

result of 21st century technology, and second, the importance of these skills, relative to a basis for learning in facts, has dramatically changed in recent years.

Let me address the second of these points first.

First. In a world of a million facts, if you learned a hundred thousand you could get by with your basic education and a good library. But in a world of a trillion facts, your education of a hundred thousand facts is pitifully small, probably irrelevant, and no library is large enough to hold all the facts. You need a new skill, a way to access the facts you need from an *ocean* of facts, and the tools of a person who used to just dip from the well will be insufficient.

Second. Our world of a trillion facts, moreover, is much less static than the old world of a million facts. New facts come into being all the time. It used to be, if you knew who married whom, you knew. Now you need a scorecard. It used to be, there were basic foundational elements. Now there are quarks and muons and who knows what showing up in the scientific press tomorrow. It used to be the case that planets were discovered, and in all of human history, this had happened nine times. Now new planets are discovered every week and our understanding of what is a planet has changed. You need new skills to keep track of how what you know *has changed*, and the skills of a person who simply accumulates facts are insufficient.

Third. It used to be you needed to know about your home and community, and maybe a bit about your state and your country, basics of mathematics and language, farming or mechanics, and that was about it. Now you need to know about the prince of pears in Japan, growing unrest in Mongolia, the basics of everything from electronics to finance to algebra, and *still*, most of the information you are presented with every day is *irrelevant* to you. You need new skills to be able to select and prioritize information, and the skills of a person who just watched and learned are not enough.

Fourth. Our world of a trillion facts consists consists of many new *types* of fact that must be assessed in new ways. In the ancient world, reason consisted of the syllogisms and geometry and numbers, and was used by a small number of philosophers. In the renaissance we extended our understanding of propositional logic and discovered the calculus, and these were used by courtly scholars. In the 20th century we discovered that logic and mathematics are part of the same system, added probability and statistics, and developed imaging and video technology, and these were used by scientists and professors. By the 21st we have learned about networking, programming, photoshopping, topos and memes, visual literacy, and more, and these (along with all the rest) need to be used by everyone. In the 21st century, there are *more types* of reasoning, and they must be used by *more people*.

Fifth. Just as in the 20th century everyday people learned skills not even imagined in previous years (like, say, driving a car), the 21st century is seeing an accelerated need for new skills. Just imagine, in the first decade we have seen the need for a host of skills related to managing one's presence in a social network, finding information using a search engine, and more (much more). We need to be able to turn

our knowledge into these and other sorts of skills very quickly. And more and more people need to be able to learn these skills.

Sixth. The rise of internet technology has corresponded with a rise of activism and agency. In the world of even a few decades ago, the mass of people did their jobs, did what they were told, and exercised their options - *if* they lived in a democracy - through the vote. Today, people manage much more of their own lives (and clamour for even more). Almost everyone lives in a democracy, and even if people don't vote, they are increasingly involved in community activities, not mere socializing outside the family but organizing, participating, creating, lobbying and more. The skills we need in order *to simply act* are far more than what used to be required, and are needed by far more people.

So these 21st century skills are more numerous and needed by more people than ever in the last few years, with no sign of this trend changing or even slowing down. What of their relation to facts?

Not so long ago, pretty much every bit of information a person needed in his or her life could be taught as a fact, which basic mechanisms - such as literacy - being used to make up the difference. Spending a lot of time teaching facts could be justified, because people needed basic knowledge to survive in an industrial world, needed to be able to understand the basics of language and literature, science and mathematics, and - crucially - not much more. And anything that detracted from that learning made a person *less* able to cope in society. These useless 'soft' skills might help with their hobbies and avocations, but they wouldn't help them get a job or do well in their career.

Today, the situation has completely turned around because of the six factors identified above. People need such greater capacities in literacy, learning, prioritizing, evaluation, planning and acting. And as their need for these dynamic skills and capacities increases, their need for facts decreases. Indeed, the more these skills are needed, the more the teaching of facts *as facts* actually *impairs* the teaching of these skills. The more *static* our teaching, the less *dynamic* the learner can be.

Let's return to the computer system analogy. Let's imagine we are designing a computer system. We have a certain amount of memory, a lot, enough for our purposes but not so much that every fact in the world can be stored on it. Just like a human brain. And we have certain expectations of our computing system - that it can help people lead their lives, for example, that it will be useful in the current work environment, that it can help improve productivity and make decisions.

We *could* simply fill it up with facts. That's what we did with books: a book is basically a system that does nothing but store facts. Assuming we had some way to access those facts, it would be a superb resource. A library on a desk. If we had *some* mechanism for querying what it knows - basic literacy - we would have a system that could do a lot of useful things.

But it wouldn't do what we wanted, would it? We would want, *at the very least*, a computer system that could add new facts. We want some way of writing to memory and to (perhaps) change existing

memory. And, in fact, once we started along these lines, we discovered that we want our computer systems to do a *lot* more than to simply store facts. So much so, in fact, that the storage of facts in the computer became a *secondary activity*.

Think about it for a second. Our computer system, the one we use to do finances, when it came out of the box, *knew nothing about finances*. The brand new hardware, when it came off the assembly line, that we use to read Shakespeare, *knew nothing about literature*. The computer, which now performs advanced engineering, when it was first delivered, had only the basic rudiments of mathematics, and knew nothing about applied math, geometry or physics.

Even more importantly: if we had, while we were building this computer, programmed into it the knowledge of finances, literature and mathematics, it would have been a *less useful* computer, not a more useful computer. Nobody would build a financial system into the operating system (much less the hardware); it would be obsolete with the next year's tax laws and would double the amount of time (and cost) it took to set it up. It would be nice to receive a computer with the basic works of literature pre-installed, but only if we were actually planning to read them; otherwise, it's simply a waste of disk space and an unneeded addition to the cost.

And even more significantly: even if we *could* program all of this knowledge, all of these facts, into the computer ahead of time, we have no idea how we could actually *use* these facts in day-to-day operations. Sure, you could provide any fact you were asked, but then what? How do these facts combine to form a computer game? A term paper? A funding proposal? A new house? Sure, you may need facts to do any of these things - but *which* facts? In *what order*? It makes no sense to store everything you could possibly need ahead of time, and even if it did, it makes no sense to store these facts without *first* understanding how they might be used, where they might be used, when then might be used.

And what we discover when we think about it this way is that it's not simple whether or not we need facts that is important, but also, *what format* the facts are in that is equally important, if not more important.

Again, suppose you are contemplating a computer purchase. You know that you want a tool to help you do your job and the rest. You decide ahead of time that it will need facts, because you are convinced that you can't do your work without facts. So you go out and buy a bunch of books.

Then, when you get your computer, you will feel cheated. The computer is stupid. It can't do anything. And when you need a fact, you go to your books, not the computer. The computer may have a lit of *skills*, you may say, but what I need in order to do my work are *facts*. And so you continue to maintain that the best way to support your work is to use a printing press and to store your facts in books.

What you have done, of course, is not only to use an older, less efficient, system for storing and

retrieving facts, you've cheated yourself out of a way to accomplish a whole range of new tasks and activities. You've cheated yourself out of the very *possibility* of mastering these new skills. Think about the problems you've created by depending on a library, by depending on an information system in which facts are *impressed* on a storage medium:

- you have to buy new books to get new information, an ongoing and expensive activity
- your books don't update, and you have no real way of knowing when any bit of a book is out of date
- you have no good means of choosing which books to buy; you can handle your local bookstore, but the thought of a library with a trillion books is frightening
- you have no way of knowing whether something in a book is true or false
- you have no way to move beyond 'book learning', and nothing in the book tells you when you should do something (your actions are underdetermined by your knowledge; should you believe the sceptic, who says there is no floor, or the alarmist, who says the building is on fire?)
- you can't develop skills; despite reading all about 'bicycle riding' you still fall over

You need, in other words, need to acquire facts in a format *appropriate* to your knowledge system.

That's why, when we design computers, *first* we build the hardware, then we install the operating system, then we install application programs, *and only then* do we add the data - the *facts* with which we expect our computer to work.

The same principle applies in education and learning.

Take driving, for example. If our knowledge of how to drive depended on a set of facts, then at a certain point it would become impossible, because while we could teach people how to drive on common streets and in common situations, as we drive further and further away from home, in newer and different vehicles, our knowledge becomes less relevant, until eventually we are simply unable to drive. If, instead of focusing on the 'facts' of driving, we think of driving as an activity or *skill*, then we are able to adapt, and develop new abilities, and new knowledge, mastering the ability to drive in strange places as we progress.

Or take mathematics, for example. If we just need basic mathematics - operations, algebra, geometry, trigonometry, and calculus - then we could simply learn the facts and we're fine. But if we envision actually working with mathematics, and extending our knowledge of mathematics well beyond these basics, then our method of learning by adding facts will make it harder and harder to progress, and beyond a certain point, progress will become impossible. If, however, mathematics is taught, not as a

set of facts, but as a *skill*, then advanced mathematics becomes more like new terrain over which we are navigating, rather than new stuff we have to memorize.

21st century skills are, in short, *an operating system for the mind*.

They constitute the processes and capacities that make it possible for people to navigate a fact-filled landscape, a way to see, understand and acquire those facts in such a way as to be relevant and useful, and in the end, to be self-contained and autonomous agents capable of making their own decisions and directing their own lives, rather than people who need to learn ever larger piles of 'facts' in order to do even the most basic tasks.

And even more: when we understand facts in this way, when we understand that facts are like data, then we obtain a more sophisticated understanding of the nature of facts themselves. Because, throughout this discussion, we have been using the word 'facts' uncritically, as though they represent some atomic basic, a state of understanding below which we cannot delve. But *in fact*, we can see, through our newly acquired critical capacities, that our relation with facts is *much more contingent* than previously supposed.

And it is *that* that the common core people object to. Not simply that people can be taught new skills and capacities. But rather, that these skills and capacities result in an understanding that facts are *created*, and that they, too, can create facts. That facts are not beyond questioning, and that facts not only should be questioned, they *must* be questioned. The common core people want the means and the ability to implant *unquestioned truths* into the minds of children, and this in an environment where the possession of unquestioned truths becomes to be more and more of a handicap, an impediment, a barrier to personal growth and prosperity.

They want to *use* children to promote their own political agenda, rather than to enable children to have lives, beliefs and faiths of their own.

What we have learned - what we are understanding, uniquely, in the 21st century - is that the nature of facts is very different from anything we thought before:

First. Facts are not simply read, they are not simply expressed in language, and they are not independent of the means in which they are expressed. The very same truth, expressed in a poem, says something different than expressed in a photograph; the very same truth, told to your girlfriend, means something very different when expressed to your grandmother. There are not only many languages, there are many forms of expression, and literacy involves not only reading books, but reading faces, photos, idea, omens and portents, signs, between the lines, and much, much more.

Second. Facts change. There is no simple hierarchy of facts, with some facts being universally true in all cases (because the same fact, represented differently, becomes a different fact, meaning something

different). Our understanding of the world changes, it shifts and weaves about like a riverbed, it grows and it shrinks as we look more or less closely at a state of affairs. At any given time, we only have a point of view, a perspective, a way of seeing a fact, never the whole thing, a fact-in-itself, which means that even if there is a fixed state of affairs in the world, we each can at best attain knowledge only of a part of it, and that other people, inevitably, will understand only a *different* part of it.

Third. Some facts are important and some facts are not. Some facts are *salient* - vivid, impressive, animated - and have an impact on how we see the world, how we categorize things, how we decide whether things are similar or different. In some cases, it may make no difference in our lives whether or not some fact is the case (and we, therefore, have no way of knowing whether it is true), while in other cases, the fact of the matter may be of the utmost importance to us. And different facts are important to different people, and there is no single set of facts - none - that is important to everybody.

Fourth. There is no easy way to determine what is a fact and what is a misrepresentation, but there are ways, and these ways are accessible to everybody. When somebody tells you that '1+1=3', it doesn't matter whether you know psychology, or physics, or engineering, or whatever it is that they're talking about, you can still know that they are misleading you in some way. Though every domain, every discipline, every *person* may have their own specific knowledge, their own way of talking, there is no need for *you* to accept statements from that domain, discipline or person, simply because you are not know every fact. Detecting deception is a *skill* and you need it just as much as your computer needs to be able to detect malware and viruses.

Fifth. You need to be able to *decide*. To consider the state of affairs, consider the data you have at your disposal, and to take some action. This involves some decision-making mechanism that is *not* rote performance, but rather, a complex exercise of calculation that involves the entire mind, and not simple rules or principles of action.

Sixth. You need to have the *capacity to act*. The skills that you need to create output, to communicate your thoughts and beliefs and desires, to transport oneself, to manipulate the environment, to interact and participate in the wider social order in a meaningful way. This involves not only the performance of one's job but also participation as a citizen in society, a creator and contributor to culture, the creation and raising of a family, and the nurturing of one own physical, mental and spiritual life.

We are in a period of transition. We still to a great degree treat facts as things and of education as the acquisition of those things. But more and more, as our work, homes and lives become increasingly complex, we see this understanding becoming not only increasingly obsolete, but increasingly an impediment.

Today - surely we've seen enough evidence of this! - if you simply follow the rules, do what you're told, do your job and stay out of trouble, *you will be led to ruin*. It's like sitting on a log floating in a river: it works for a while, and seems like the safest place to me, but all the while, you're approaching a


waterfall. Whether it be a financial crash, the degradation of the environment, war and terrorism, or even something as simple as a car accident or family crisis, you will need more and more the ability to keep yourself afloat in troubled and rapidly changing circumstances, and an abundance of facts will not help you, it will instead sweep you over the waterfall.

Moncton, September 19, 2009

Personal Knowledge: Transmission or Induction?


I'm going to use an oversimplified example from electricity to make a point. I still think there is a deficiency in the personal knowledge management model being discussed in various quarters. Let me see if I can tease it out with the following discussion.

[Harold Jarche](#) points to a diagram [Silvia Andreoli](#) adds to his [last post](#) on personal knowledge management. Here it is:




Now the activity happening at the centre is becoming more sophisticated, with an expanded list of processes happening, to convert data into knowledge. I don't want to focus on the particular types of activity - that's just mechanics. I am more concerned on what might be called the 'flow' of information from data to knowledge.

So let me strip down the details and present an abstract version of the model.




In a nutshell: does the data *itself* become knowledge, or does the data lead to *something else* becoming knowledge? Let me use my electrical analogy to make the point.

In what might be called the 'naive model' (not disparagingly) we have a direct circuit from input (data) to output (knowledge). The purpose of the process in the middle is to filter, transform, reshape, and otherwise improve the data, but ultimately, to pass it along. Like this:



Now presumably, what is happening here is the data is coming in from outside the person and the knowledge is being stored or in some way impressed in the head or mind; there may in addition be an output in the form of a transmission or creative act, producing the freshly minded data as publicly accessible 'knowledge'.

But I'm not at all sure this is the correct model. I don't think there is a direct flow from data to knowledge. My model looks more like this:



What we have here is a model where the input data *induces* the creation of knowledge. There is no direct flow from input to output; rather, the input acts on the pre-existing system, and it is the pre-existing system that produces the output.

In electricity, this is known as induction, and is a common phenomenon. We use induction to build step-up or step-down transformers, to power electric motors, and more. Basically, the way induction works is that, first, an electric current produces a magnetic field, and second, a magnetic field creates an electric current.

Why is this significant? Because the inductive model (not the greatest name in the world, but a good alternative to the transmission model) depends on the existing structure of the receiving circuit, what it means is that the knowledge output may vary from system to system (person to person) depending on the pre-existing configuration of that circuit.

What it means is that you can't just apply some sort of standard recipe and get the same output. Whatever combination of filtering, validation, synthesis and all the rest you use, the resulting knowledge will be different for each person. Or alternatively, if you want the same knowledge to be output for each

person (were that even possible), you would have to use a different combination of filtering, validation, synthesis for each person.

That's why personal knowledge is *personal*. Each person, individually, presumably attempting to approximate the production of knowledge output exhibited by other people who know (where this knowledge output may be as simple as the recitation of a fact or as complex as a set of expert behaviours in a knowing community), must select an individual set of filtering, validation, synthesis, etc., activities.

And probably, the best (and only) person who can make this selection is the person him or her self, because only the person in question knows and can make adjustments to the internal circuit in order to produce the desired output. That doesn't mean we can't suggest, demonstrate, or in other ways mediate these adjustments.

So why do I think the induction model is more likely to be correct than the transmission model?

What characterizes induction is a *field shift*. Though we can track the flow of energy from input to output (which is why no causal laws are broken) the type of energy changes from electrical to magnetic and back. Hence, the *carriers* of the energy, the individual electrons, never connect from beginning to end.

A similar sort of field shift happens in knowledge transmission. When we think, we convert from complex neural structures to words. We output these words, and it is these words that constitute (in part) the data that enters the system (other forms of data - other audio-visual inputs, are also present). This data, in the process of becoming knowledge, is not stored as the physical inputs (we do not literally store sounds in our brains) nor even echoes of them.

Rather, what happens is that, as the cascading waves of sensory input diffuse through our neural net, they have a secondary, inductive effect of adjusting the set of pre-existing neural connections in the brain. It is this set of neural connections that constitutes knowledge, not the set of signals, however processed and filtered, that interacted upon them.

At a certain gross level this should be pretty obvious. When we examine the brain, we do not detect sounds or images, nor even (beyond the most basic sort) echo-like constructions or neural arrangements that correspond to them. Nor either do we detect sentences, syntical-like structures, or anything similar. Therefore, whatever knowledge is, it *has* undergone a field shift.

But we do *infer* to the existence of such structures, and we infer to them on the basis of what appear to be obvious productions of knowledge. Not only can we write text, draw pictures, and speak descriptions, we have actual memories and dreams that have the *same phenomenal qualities* as those we experienced in the first place. This would not be possible (to echo Chomsky) were they not stored in the

mind in the first place. Would it?

Only if there are no field shifts. But if there *are* field shifts (which would explain why we cannot observe in the brain what we so obviously experience in the having of one) then the production of dreams, memories, verbal utterances, and other behaviours constitutes the reverse field change. It's like converting the magnetism back to electricity again.

Our dreams, memories, thoughts and behaviours aren't stored in the brain and the re-presented. *They are built from scratch again* as a result of the functioning of the neural network. A memory isn't the *same* experience which is had a second time. It is a *new* experience.

That's why we misremember, have fanciful dreams, see things as we want to see, and all the rest. When we are recreating the phenomenal experience, this recreation is affected by all number of factors, all the other elements of the neural net, the configuration of the pre-existing circuit.

We can draw numerous lessons from this, and I have drawn them in other posts. That we do not remember 'facts', for example. That knowledge is 'grown' through the growing of our neural network, rather than accumulation or construction or any of the theories that do not incorporate a field shift. And the rest, which I won't reiterate here.

Why this is important for the present purposes is that it changes our approach to the sorts of activities postulated to take place in personal knowledge management, the filtering, validation, synthesis and all the rest. Because we now have two points of view from which we can regard these activities:

- from the perspective of the *content* on which they operate, or
- from the perspective of the *person* that is doing the operating.

To put the distinction very crassly, we could say that, on the one view, the content constitutes the knowledge, while on the other hand, we could say that the *operation* constitutes the knowledge. This latter view, which can be classified under the head of [operationalist](#) theories of knowledge, is more representative of the inductivist approach.

The paradigm case here is mathematical knowledge. In what does a knowledge of mathematics consist? A typically realist interpretation of mathematics will say something like, "there are such thing as mathematical objects, and there is a set of facts that describes those objects, and mathematical knowledge consists of the acquisition, or at the very least, the internalization, of those facts."

An operationalist interpretation of mathematics, by contrast, remains silent on the question of the existence of mathematical objects, and interprets mathematical knowledge as corresponding (for lack of a better work) to the operations typical of mathematics. The number 'four' is tantamount to an act of

counting, "one - two - three - *four*." The act of addition is tantamount to the act of putting one pile of beans in the same place as another pile of beans, and then counting all of them (a short though critical account of Kitcher-Mill can be found [here](#)).

When we place the locus of knowledge, not in the content, but in the person, then the content becomes essentially nothing more than the *raw material* on which the learning practice will occur. What matters is not the semantical referent of the input data, but rather, the *act* (or operation) of filtering, validation, synthesis, etc., that takes place on that content.

When we say something like "words are things we use to think" we should understand this in the sense of "paint is something we use to imagine" or "sand is something we use to tell time". Time is not in the sand, imagination is not in the paint, and thought is not in the words. These are just raw materials we use to stimulate an inductive process - we can generate a field shift from thought to sand to thought again.

Even when you are explicitly teaching content, and when what appears to be learned is content, since the content itself *never persists* from the initial presentation of that content to the ultimate reproduction of that content, what you are teaching is not the content. Rather, what you are trying to *induce* is a neural state such that, when presented with similar phenomena in the future, will present similar output. Understanding that you are train a neural net, rather than store content for reproduction, is key.

Moncton, Monday, April 05, 2010

Virtues Education

Re: [What It Takes to Make a Student](#)

By Paul Tough, New York Times

This was an interesting article, well worth the read.

A couple of things are happening here. First, and at surface level, the author is attempting a defense of the No Child Left Behind (NCLB) legislation, even though even he admits, near the end of the article, that it will probably fail. He does this by focusing not on the mechanics of the legislation - standardized testing and penalties for 'failing schools' - but rather on the intent: reducing the gap between students from poor and minority households and their contemporaries from white middle class households.

By the measures taken to demonstrate achievement (which are at no point questioned by the author) NCLB is failing. On almost all measures there has been virtually no movement in the gap in achievement (and it is worth noting that, even by this measure, the achievements of students from the two groups have been *falling* in tandem).

The second, and more interesting, part of the article is the author's examination of *how* he thinks NCLB can be saved. First, he argues, we need to recognize the value of the objective; we need to move from saying "these kids can't be saved..." to "these kids can be saved only if...". Then the question becomes, what's working?

This is where the study of the different parenting styles between upper income and lower income households becomes significant:

- children of higher income parents are spoken to more frequently
- they hear a larger vocabulary of words
- they hear more encouragements, and fewer discouragements
- their parents are more involved in their lives

The suggestion here is twofold: first, that children of higher income parents get a better foundation for their education, and second, and much more interestingly, they acquire a different set of cultural values than those of poor parents, a set of values that is more valued in society, and it is the having of these values that gives them a significant advantage.

I think there is merit to this observation.

Take smoking and exercising, for example. My observation is that people with higher incomes are much more likely to exercise and stay in shape, and much less likely to smoke. Consequently, a person from a

lower income home, even if very well educated, will enter society with these different habits. They will remain in lower class social circles (the smokers congregating outside, rather than the exercisers at the club).

Many cultural mannerisms are a lot more subtle. The story looks at how polite people in the different groups are and how they listen to other people. But the range of phenomena is there for the studying: dress and hair style, eating habits, musical preferences, speaking style and accent, and the more.

"The real advantages that middle-class children gain come from more elusive processes: the language that their parents use, the attitudes toward life that they convey. However you measure child-rearing, middle-class parents tend to do it differently than poor parents — and the path they follow in turn tends to give their children an array of advantages."

Leaving aside the criticisms of that for the moment, we can see now how the schools singled out by the author apparently succeed. In a word, they are immersive. They occupy much more of a child's time, and they require the child's total attention while they are there. They deal with all aspect of the child's life. They treat learning as immersion into a cultural community as much as (or even more than) the acquisition of a certain set of facts. And they explicitly instruct the children in those aspects of this culture they may have missed in their home life.

Some of these result in almost immediate improvements in outcomes. The 'slant' listening method, for example (sit up, listen, ask questions, nod and track the speaker). I notice when I speak that this is how people listen to what I'm saying, and I employ these cues to adapt my talk as I go along (interestingly, the style of listening I learned was rather different - "active listening", where you actually restate and echo what was said, was much more to my style, much more immersive, though considered 'disruptive' in a classroom).

A lot of the cultural discipline has to do with work habits, this based on work showing "self-discipline scores were a more accurate predictor of G.P.A. than the I.Q. scores by a factor of two." The trick to 'attitude', of course, is how to get it. The author writes, "Duckworth's paper connects with a new wave of research being done around the country showing that 'non-cognitive' abilities like self-control, adaptability, patience and openness — the kinds of qualities that middle-class parents pass on to their children every day, in all kinds of subtle and indirect ways — have a huge and measurable impact on a child's future success."

A third part of the article looks at how the successes of these schools to tout the benefits of segregated learning and charter schools. "KIPP, Amistad and North Star were embraced by advocates from the right who believed in the whole menu of conservative positions on education: school choice, vouchers, merit pay for teachers." But the more important lesson, he writes, is "the effort that would be required to provide those students with that education."

This is significant. Because in one sense, what these programs show is something we have always known: that if you immerse a child into a culture, they will generally adopt the beliefs and values of that culture. And it follows that if the culture is one that values learning, and behaviours that lead to achievement in learning, then the child will, all other things being equal, learn to achieve in learning.

But, first of all, this is not sustainable. As the article notes, the level of education provided in these schools is much higher, much more intensive, than in traditional schools. Students are in class 60 percent longer, and the instructors work 15 hours a day. This means the students cannot work part time, which excludes the poorest of them. And it will not be possible to require teachers to work such hours (or to pay for the staff to provide an equivalent level of support). Charter schools are premised (in part) on the idea of getting more out of teaching staff but eliminating their unions, but experience is beginning to show that better working conditions, not worse, are required to get the most out of teachers.

Secondly, though, we need to question the wisdom in the long term, and across a wider scale, of what amounts to cultural subversion. The methodology involves what is essentially a separation of the students from their own culture, placing them into a substitute culture. This creates the tension between home and school captured in such films as 'Dangerous Minds'. More, it also carries the connotation that one culture - the white middle class culture - is 'good' while the minority and lower class cultures are 'bad'.

My observation is that in Canada we have taken a different approach to obtain the same outcomes. This approach is characterized in an advertising campaign that ran frequently on all television channels (though the program seems to have ended with the current government). "Babies' brains do not grow by themselves," the announcer says. "You need to help them grow. Talk to your baby. Sing to your baby. Play with your baby." What is significant is that in the background we hear people talking in their own language, singing their own songs. (You can also see the British equivalent online, <http://www.literacytrust.org.uk/talktoyourbaby/index.html>).

What this program recognizes is that children are *already* immersed in a cultural environment, and while there is nothing particularly wrong with that environment, the children need supports that they are not obtaining from their family and social milieu. So, instead of removing the children from their culture entirely, they are attempting to foster the behaviours and attitudes that will lead to improved learning. Instead, in other words, of taking the students to the learning, it takes the learning to the student. And instead of rebuilding a culture from the ground up, it supplements existing cultures.

This is more difficult because it doesn't address the second part of the equation - that society values some cultural traits more than others. For while in some cases this valuation is reasonable and rational - valuing self-control and openness, for example, promotes safe and trusting environments - in other cases - valuing straight over gay, for example, or short hair over long - the valuation creates barriers that exist only for the purpose of exclusion and defense of privilege. Tacitly and passively accepting the

valuations made by the dominant culture creates an environment of permanent privilege, for no matter how much a person adapts, they will never be able to 'fit in'.

For this reason, creating an environment that fosters greater achievement by the children of lower income and minority families also involves creating an environment where the values favoured by society as a whole are not cultural values per se but rather values that are culturally neutral. To create, in other words, to foster a common cultural syntax, rather than a common cultural semantics. To, in other words, value certain behaviours, mechanisms and protocols, but to encourage independence and diversity with respect to belief, style and appearance. Where 'cultural syntax' is a description of the mechanisms that enable each other to interact, while 'cultural semantics' describes the mechanisms that give meaning and value to individual cultures.

Cultural syntax, unlike cultural semantics, can be treated as (and thought as) a 'game': "Middle-class Americans know intuitively that 'good behavior' is mostly a game with established rules; the KIPP students seemed to be experiencing the pleasure of being let in on a joke." It is a set of movements and actions that you undertake, a protocol, a process - a game, in other words - but emphatically *not* an abandonment of cultural value.

What this involves in practice, then, is not the segregation and isolation of children in what amount to 'cultural factories' but rather an understanding of and adaptation of the environments children already inhabit. It involves, in other words, the 'seeding' of these environments with the knowledge, tools and values children need in order to become effective and capable learners.

We have already seen this effort undertaken with significant effect in previous societies. Propaganda and advertising work. That is why propagandists and advertisers spend billions of dollars every year. That is why, in every moment of every day, we are subjected to a barrage of commercial messages. That is why, through the watching of television programs, sporting events and musical presentations each generation acquires a particular set of shared icons and images, a particular vocabulary, a particular ethos. That is why it is even *possible* for people like Drucker to talk about a 'net generation'. They (or, at least, the white middle class version of 'they') are the product of what we have taught them, through our society and our media.

In order to address the problems of inequality in education, we need to begin by recognizing that learning occurs not merely in the hours in front of a teacher but also in the many more hours spent watching television, playing video games, surfing the internet and inhabiting society generally. And that in order to foster a learning culture it will be necessary to take the learning to *where the children are* and to seed those environments with exemplars of the values and virtues desired.

This is *possible* but may in some societies be very difficult. In all societies commercial media more or less predominates, and the messages contained in these media are fashioned in such a way as to promote the consumption of the products and services offered by the sponsors. From time to time good models

and good values appear (one thinks of the work of people like Bill Cosby or Henry Winkler) but on the whole this is the exception rather than the rule. It is not necessary to cynically observe sponsors would prefer that their customers remain uneducated, rather merely to observe that public education simply isn't one of the things sponsors think about.

This is not a situation that will be reversed through public policy. The government cannot ordain that the environments of all its children will be converted into environments supportive of a learning culture, with positive role models. Even programs such as 'talk to your kids' and only have a limited effect. Only total control of the media would even make this possible, and the dangers inherent in total media control are now apparent to us all. It will not be possible to 'manage' or 'construct' a pro-learning culture.

What is needed is a constant, persistent and positive reiteration throughout all of media of those practices and values needed in order to foster a learning culture. People need to demonstrate their commitment to learning, and to do so in a public manner, and in such a way that they will be emulated by those who follow. Children - and their parents and communities - need to be encouraged and supported in the development of their own learning cultures and environments. We need to put into the hands of these parents the knowledge and tools that will increase their capacity to teach, and the models they can point to in support of this practice.

There is no quick fix - and that is perhaps the greatest disservice this article does. There will not be a turning around of society in a few years. Not because there is a lack of will, and not because there is a lack of resources, but because the attempt to engineer, rather than grow, a learning culture is fraught with difficulties and dangers the planners cannot imagine. Paul Tough is quick to say academics do not examine schools to see what's going on inside. I have stood inside former residential schools in Canada, however, and taught their former students, and can speak of the damage caused when you attempt to replace one culture with another.

Indoctrination, however much it appears to work in the short term and on a small scale, is not the answer. Empowerment is, and only ever has been.

Free Learning and Control Learning

Text of my presentation to SURF Education Days, 13 November 2007, Utrecht, the Netherlands. Slides, audio and video are also available.

1. Introduction

I don't have fancy slides today. I don't have nice pictures or anything like that. I've spent some time over the last few days looking at a paper by Paul Kirschner, John Sweller and Richard E. Clark, which describes the "failure of constructivist discovery problem-based experiential and inquiry based teaching."

For those of you who are familiar with my work you'll know that a great deal of the work that I've done is constructivist discovery problem-based experiential inquiry based teaching. And so this sort of paper is published in Educational Psychologist, which I guess is an important journal. It's been widely cited.

This sort of paper and the criticisms associated with that sort of paper raise questions. I do work in educational theory and I also do work in software development and I'm never completely sure that I'm doing the right thing. I raise questions. Is the work grounded in research? Is this really the way learning happens? And so I question myself and I question myself not just because it's a good thing to do, but just because I'm not positive.

So I worry when I see papers like this. I worry that maybe the foundations of the work that I and other people are doing is not well-grounded, but then I look at a paper like this and I realize that what really needs to be done is that these arguments need to be drawn out. They need to be made explicit and it needs to be shown very clearly and not ambiguously why these are not good arguments and that is going to be my task today.

And for you who are listening to this task today, what you'll be able to take away from this, I believe, is, first of all, an outline an idea of the theory of learning that underlies learning with blogs and wikis and other web 2.0 technologies, but also, more importantly, a way to respond to people who think that instructional technology and the use of the Internet in online learning ought to be nothing more than the presentation of instructions to students telling them what to do. That is control learning. That is the old way of learning. The web way of learning, Web 2.0 learning, is the new way. It is free learning and that is what I advocate.

2. Connectivism (Free Learning)

Sometimes this theory is known under the heading of connectivism. Connectivism is a name that was coined originally by George Siemens. Connectivism is essentially, and this is my take on it not necessarily George's, the theory that knowledge and learning can be described and explained using network principals.

Now what do I mean by that? What I mean is: knowledge itself, to know something, is to be organized in a certain way, to have a certain pattern of connectivity in the mind, a certain neural connection in the mind. To learn on that theory is therefore to acquire that pattern of organization. To learn is therefore not to have things pushed in your head but to grow and to develop in a certain way and specifically to grow and develop in such a way that you are able to recognize patterns in the environment.

Connectivist learning theory, therefore, is based on the theory of how networks learn, that is to say, how networks grow, how networks develop, how networks form structures of connections between neurons. There are four major ways in which networks grow. And I'm not going to say that these are the only ways, these are all the ways, that this is the definitive statement. But these are ways that we have observed through history that networks grow.

One way is simple Hebbian association. What that means is that if two neurons fire at the same time and don't fire at the same time, a connection tends to be drawn between them. That's it. Very simple!

The second way is accidental association. If two neurons are beside each other a connection tends to join them.

The third way connections are formed is back propagation and this comes from the theory of connectionism in the field of computer science. The networks form their connection and then feedback is sent into the network according to the output that the network produces. If the network produces good output the connection will be reinforced. If the connection produces bad output the connection will be broken.

And then finally, Boltzmann learning, which is a theory based on thermodynamics, which says essentially that connections will tend to form at the most stable configuration. If you think of it as like throwing a stone into a pond, the water will settle out. Well connection in brains work in much the same way, according to this theory, and the brain settles out.

The main thing to understand here is that connectivist learning theory is about how connections form in the brain, and for that matter how connections form in networks generally. Because connectivists talk about not simply networks in the brain, they also talk about learning networks in society at large, networks of people in society who are connected to each other. The two theories work out to be two parts of the same theory.

In connectivist pedagogy, therefore, to teach is to model and demonstrate. To teach is to present experiences to people so that they can begin to form these connections in their mind. And then to learn is to actively form these connections by practicing, by repetition, and by reflecting on that practice.

Both of these imply what might be called participation in an authentic community of practice. The idea here is that to learn is to put oneself in a situation where you are practicing in the way that whatever discipline you are in is practicing. For example, you learn physics by doing physics. You learn how to take care of forests by going to a forest.

The role of the teacher in this model is to practice one's work in an open manner. This has been a challenge, I think, for pretty much all of society, but the idea here is that instead of doing your work in secret in back rooms without being open about what you do, you do your work in an open and transparent manner so that people can see what you're doing.

In preparation for this talk, for example, I created a summary of the Kirschner, Sweller and Clark paper and I put that on my Website and I collected notes from other people who wrote about this paper and I put it on the website and what I wanted people to see is how I go about assembling my thoughts in order to prepare for a talk like this.

To work, on this theory, is to engage in a community. Most of us, when we do our work, in whatever profession, don't do it simply all by ourselves. We are involved in a community of practice. We have shared ways of doing things. We have a shared vocabulary, a shared understanding of what constitutes success, a shared understanding of how we test for that success. And it's to be openly reflective to think about what we're doing in this community and to think about whether it is the best way to do it, why we're doing the things that we're doing.

In this model the role of the learner is to themselves in some sort of environment like this. It may be the actual community of practice itself, which is what I recommend, or it may be a simulation of that community, perhaps a role playing game, perhaps electronic performance support systems, perhaps the actual community of practice itself. You can imagine all the different ways a learner can place themselves into one of those communities of practice. It is to observe the way people who are successful in that practice actually conduct their practice and it is to be reflective to engage in conversations about that practice.

You can see why Web 2.0 and Internet technologies like blogging and wikis and things like social networks play such an important role in these theories. These are technologies that make this sort of activity possible. These are technologies that make it possible for a person to practice their discipline in an open manner. These are technologies that make it possible for a community of practice to develop on a worldwide basis connecting people from many different countries together. They make it possible for the learner to observe experts, to interact with experts, and to learn by doing.

3. Instructivism (Control Learning)

The other view of the world is known as, well, there are different names for it, but I will call it control learning or instructivism, and it is the approach that's characterized in the Kirschner, Sweller, and Clark's paper.

It's a model of learning, and especially on line learning, that is what we might call traditional on line learning. What this learning is, is learning based in the theory of learning objects and based in the theory of the learning management system.

And the model here is that our on line learning environments basically emulate the practices and the processes of the content in a traditional classroom. Learning objects contain content in the core, and pedagogy that is wrapped around that content and the idea of this traditional learning is that the content must be explicitly instructional. The content must actually guide the student or the learner through a series of instructional steps.

The idea of learning object theory is that the learning object needs to have, for example, things like learning objectives. It needs to have guided practice. It needs to have assessment of some sort.

The basis for this theory, the basis for this approach and where it is practiced the most is in environments like the U. S. military where the Sharable Courseware Object Model or SCORM was developed. And SCORM is basically a model of personal learning where you go step by step by step through the learning manual and you follow the instructions that they tell you to follow, you do what they tell you to do and that is what is supposed to produce learning.

The learning management system and we're familiar with the learning management system, we're familiar with systems like blackboard, for example, or even open source systems like Moodle and Sakai. The intention of these programs is to present this material step by step by step.

In the introduction (to this talk) the Buntine oration that I gave in Perth is mentioned. In that talk, there are three major locuses or loci of control; three places in which control of the learning process, control of the learning content is exerted in this traditional picture.

One is in content packaging. And if you think about content packaging what content packaging is, is you take a bunch of learning objects, put them all together, put a wrapper around them and then you compress them with a zip archive or something like that, thereby making them useless even to a browser. Content packaging is a picture of learning as a package that you put on a shelf like a book in a library.

Federated search, second. Federated search is not like Google. Federated search, which is the approach

that is recommended by the people who created SCORM by advanced distributed learning is a mechanism where when you search you search this library, this library, this library, and this library and that's it. You search only from recognized authoritative sources who have 'the knowledge', whatever that happens to be, authoritative sources, I guess like Educational Psychologist, the journal.

And then, third, learning design, which is a third major component, which actually has its origins here in the Netherlands with Rob Koper's educational modeling language or EML is a mechanism for stepping you through the presentation of learning materials. Now Learning Design will put you in roles and Learning Design will branch and present different materials in different circumstances, but it's still that theory of presenting material, presenting material, presenting material, and the theory behind that is if you are told what to do you will learn. And that's the basics, the basis of instructivism.

4. The Argument

And so now I turn to the paper, that I intend to criticize in this talk; why minimal guidance during instruction does not work and the corollary of that is why maximal guidance does work.

And what the authors do in this paper is they set up two alternatives and on the one side, say the authors, are those people advocating the hypothesis, people like me, that people learn best in an unguided or minimally guided environment.

(Just as an aside every time I quote from them I've very sensitive to their use of language. Their use of language is very often loaded or prejudicial and I'm trying to be careful not to be drawn aside by that. "In an unguided or minimally guided environment." What does that mean?)

Generally defined as one in which learners, rather than being presented with essential information, must discover (now they throw that word discover in there very deliberately because they want to tie it to discovery learning) or construct (there you go with constructivism right) essential (notice that word 'essential' - that word has a whole philosophy that comes with it, the whole Saul Kripke essentialism view of the world that there are certain innate natures of things) information for themselves. So that's the bad thing.

On the other side, this is their view now, are those suggesting that novice learners should be provided with direct instructional guidance on the concepts and procedures required by a particular discipline (we'll come back to that) and should not be left to discover, (notice how it's changed a little bit from the first presentation) those procedures (again a very careful word) by themselves.

Just as an aside before I get into the main criticism, I was uncertain, and of course they do not discuss in the paper, what a novice learner is. There's a certain sense in this paper that a novice learner is any person who has not learned what needs to be learned. And if you read the paper a certain way the

paper says if you already know what is being taught you have no problem being taught it, but if you don't know, then you have to be instructed, properly so called. That would be an unfair reading of the paper, but they leave it open because they don't tell us what a novice is. Is a novice a baby? Is a novice a 10 year old? Is a novice a first year student in a college or university? We don't know.

So what do they mean by minimally guided learning? Everything! The old theories of discovery learning, problem based learning, inquiry learning, experiential learning, and of course, the most recent thing, constructivism and, if they had thought of it and felt so inclined, they would have included connectivism and on line learning probably all in the same breath.

There are two assumptions, they say, to this unguided kind of learning. The first assumption is that students should be challenged to solve authentic problems (notice the scare quotes) or inquire information, acquire knowledge (again notice how that is phrased, 'acquire knowledge,' I'm going to get it from here and I'm going to put it in here) in information rich settings. And they say (notice again the loaded term, the assumption right, as though we're just making it up) the assumption is that having learners construct their own solutions, whatever that means leads to the most effective learning experience. And then the second assumption, they say, is that the non-guided people assume that knowledge can best be acquired through experience based on the procedures (the discipline). We're going to come back to that.

Now their argument, this is their main argument here, it's a nice categorical syllogism. I like categorical syllogisms because they're so easy to work with. Any instructional procedure (i.e., ours) that ignores the structures that constitute human cognitive architecture (now there again, 'human cognitive architecture', as though the mind is like a house) is not likely to be effective. There, you can see where this is going, right? Minimally guided instruction appears to proceed with no reference to the characteristics of human cognitive architecture. So we're just one of those things. And indeed, what they mean now (we'll come back to this later in the talk) by 'human cognitive architecture' is the characteristics of working memory, long-term memory and the relations between them. Okay, fine. We'll come back to that. Thus, they conclude, minimally guided instruction is unlikely to result in effective learning.

You see why I certainly worry about papers like this is because what they're saying basically is the approach that people who are talking what 2.0 learning, 2.0 blogs, wikis, social practice, communities practice and all that, has no grounding in the theory of how the mind works. And that's a very serious charge. Turns out to be false but it's still a very serious charge.

So this is their credo, their manifesto, a.k.a. their conclusion: after a half century of advocacy of minimally guided learning (people like me) it appears there is no body of research supporting the technique. Now, off in the distance you can hear the howls and the wails of protest from the people who have been studying this stuff for 50 years and have found that it works, but we'll leave that aside.

Insofar that there is any evidence, they argue, it almost uniformly supports direct strong instructional guidance rather than constructivist based minimal guidance. Not only is non-guided instruction less effective (they even talk about this for a bit) it may produce negative results.

And you ask, “How can that be?”

5. The Reality Check

Well let’s do a reality check first about their conclusion. Their argument, first of all, is simply inconsistent internally. These aren’t major issues, but it’s a bit of a problem.

On the one hand when they’re busy criticizing the minimally guided research they say instructors can’t apply it, they always cheat. They always do some scaffolding. They won’t let students discover things for themselves. They’re always suggesting, telling them what to do and all of that. Okay, fair enough.

But on the other hand, they say that minimally guided instruction is failing. Well either they do it and it doesn’t work, which is bad, or they don’t do it. You can’t say both. You can’t say they’re not doing it, and it doesn’t work.

And then also the section later on in the paper they talk about how minimally guided learning, discovery learning, constructivism, and problem based learning especially, are used in 50 medical schools in the United States. Now I hadn’t heard any particular criticisms about the quality of doctors in the United States. Maybe that’s just a fluke.

But they examine this and they say the strongest criticism they find in those doctors, because the doctors turn out to be fine, but in their paper, they recommend, “unnecessary procedures,” and it strikes me knowing the American medical system, that the last cause of the recommendation of unnecessary medical procedures is the kind of learning that we do. This is the most litigious society in the world. If you drop your pencil you will be sued. That is why they recommend unnecessary medical procedures (whatever those might be) not because they were taught to use problem-based learning.

But be that as it may. Their conclusion is simply not plausible. It doesn’t make sense. It’s not believable. We know that people learn using problem based learning and inquiry based learning. There’s a huge body of research and Hmelo-Silver, Duncan, and Chinn cite numerous studies in their response to the Kirscher, Sweller and Clark paper. And even without that research we know that people learn without guidance because we see it all the time. We have the evidence of our own senses.

Nobody went to school to learn how to build the Internet. Nobody was instructed (‘they said well first you get a website, get it on the web, get some HTML’). People discovered that all by themselves, and it turns out that the mechanisms of computers are things that even small children can learn by

themselves.

Many examples, I'll just point to one. It's called the 'hole in the wall gang' and what they did in cities in India is they would take a computer, they put it literally in a hole in a wall and so they make it assessable to the children in the community. These are children in India, so it's not like they're growing up working on their laptops at home. And so they get this computer. They're not instructed in any way. They learn how to operate it, they learn how to program it and they learn all kind of things about this computer without being told to do anything. So we can see that this works.

Even more to the point, instructivism is a kind of learning by telling; it's a kind of learning by giving people the information, the concepts, the facts. But we know that people have to learn by practice. Learning is not being told. Learning is doing.

Examples are all over the place. Deanna Kuhn writes that we can hope to impart the smallest fraction of knowledge in any science. How could we possibly teach science by teaching facts when there is not literally, not enough time to put all the facts in people's heads? Even if we were putting them in one after another every second of every day, there's too many facts in science. We have to do it a different way.

Think about what you would want in medical practice. Would you want a doctor who was told about medicine or a doctor who practiced medicine?

Their argument is based on a straw man. Inquiring learning, problem based learning, are not examples of 'minimally guided learning' and again, Hmelo-Silver, Duncan, and Chinn talked about this at some length. They are based on the process of scaffolding, they are based on direct instruction when needed. Indeed, my criticism of problem based learning and inquiry learning is the instructor is too involved. I think there's too much instructing happening in these kinds of learnings and that there should be less. But that's a separate argument for a different day.

Their argument is a false dilemma and this is the easy and obvious criticism. They are offering (remember at the very beginning where I presented their argument) the choice between minimally guided instruction or strong instruction. And it begs the question who is doing the choosing, doesn't it?

This is the example I like to use for this. Imagine my first visit to Rome. I remember getting off the train and walking out and I'm in Rome, in the train station. I have no idea what I'm doing and so a number of things present themselves here, right? I could take a guided tour. I could get a map and walk around myself. I could just walk around aimlessly and never find my hotel or... these are all things where I am choosing what to do, right? And the alternative is, to be kidnapped and to be told where to go.

Now the alternatives are not 'being kidnapped' or 'being lost'. There's all kinds of ground in between, all kinds of ways that a person can receive guidance that is not in the form of direct instruction. And so

their argument is a straw man, yet again.

People who are minimally instructed are in no sense cast adrift. So in preparing this paper for today, nobody told me how to prepare this paper, but that doesn't mean that I'm sitting all alone. I put out messages in e-mail, I log onto websites, I got all kinds of information back, really helpful useful information that mapped out the territory for me, told me about resources I hadn't found, pointed to me through objections. I got lots of guidance when I was preparing this talk, but I was not instructed.

And this is a general criticism. This isn't just me. There's a social dimension to much learning and Miles Berry points to that in his criticism of the article. A lot of learning, even traditional learning in the classroom, contains a large social dimension as you interact with the other people in your class.

6. Scientific Practice

There are some deeper misunderstandings in this paper as well and I want to explore them. Let's turn to Kirschner, Swellers, and Clark's explanation of why learning turned out to be the way it is. And their explanation is inexplicably US-based, but we'll leave that aside.

They identify the curricular reform that happened after Sputnik (it's kind of neat, Sputnik happened and then I was born; I came into the world roughly the same time as Sputnik, so I am a product of these reforms, maybe that's what really scares them). And what the reform is, they say, and it's repeated throughout the paper, is that it's based on the assumption that knowledge is best or can only be acquired or learned through experience that is based on the procedures of the discipline. And they repeat this four or five times, I didn't count them, in the article.

And so they're saying the assumption here is, if you want to learn physics you should practice physics the way a physicist does, which is what I said at the beginning of this talk. And this has led, they say, to this unguided project work (I did lots of projects when I was in school: Ecuador, the Danube River...) and a rejection of instruction (this is the key phrase, it's repeated several times in the paper), a rejection of instruction based on the facts, laws, principles, and theories that make up a discipline's content.

That's a pretty common view, isn't it? There's probably fewer advocating wikis or weblogs or something like that, Wikipedia. People are saying, "what about the facts, what about the laws, what about the content that actually is the discipline, physics, or mathematics or whatever?"

And they say it may be, they say it is, an error to assume that the pedagogical content of the learning experience is identical to the methods and practices of the discipline being studied. So what they're saying is that the basic fundamental assumption of my own theory is in error. And it's a mistake to assume that instruction should focus exclusively on application (well it's one of those weasel words, 'exclusively on', it's not what they mean, what they mean is it's a mistake to say that instruction should

be application).

Well what do they think? How do they think science works? Because if it's a mistake to adopt that method, then the nature of that method is pretty important, don't you think?

Well, happily they explain it in a couple of places for us. One place they explain it is when they discuss Kolb (1971), and Kolb and Fry (1975) and they present a process where a person carried out an action and sees the effect, and then they see this and they understand this effect and begin to anticipate the consequence, and as a consequence of that, they generalize. They understand the general principle.

It sounds like discovery learning. And they extend this to other types of unguided or minimally guided instruction. It's one of two major components of problem-based learning, they say in their paper: explicit teaching of problem solving strategies in the form of the HD (hypothetical-deductive) method. Barrows & Tamblyn 1980. 1980 is kind of significant because (Thomas) Kuhn was only about 1974 or so.

And then teaching, it's the same principle. Teaching of the basic content in the context of specific case or instance. So again the same sort of set up here. You have the specific and as a student you're supposed to generalize. That's the law. And there's the problem. That sort of way of going about learning, they argue, might not be the most appropriate way to solve problems. In fact, it's really difficult, especially in clinical settings, especially in "information rich environments." It's really difficult to come up with generalization. You've got all kind of convenient hypotheses (that reminded me of Chomsky when he talks about the poverty of the of stimulus). There's so much information it's really hard to pick the right generalization.

But the thing is, real science - the stuff that real scientists do in real labs - is not the hypothetical deductive method. Hasn't been since the '50's. It was developed by Carl Hempel, and the ink wasn't even dry on the page when Hempel and other logical positivists like A.J. Ayer were being criticized all over the place. Karl Popper, right off the bat, not verification but falsification, which is very much not the HD method. And numerous others. Kuhn - Thomas Kuhn. Lakatos. Laudan. Feyerabend. All pointing to the fact that scientists in practice do not practice the HD model. Nobody does it.

What is science? Science according to Thomas Kuhn (The Structure of Scientific Revolutions) is a community process. The process is not argumentation and it's specifically not inductive argumentation. It's explanations. The criteria for explanations are theory based or theory-bound and include things like simplicity, parsimony, testability. And explanations aren't stand-alone facts. They depend on your expectations; they depend on your theories. As Bas van Fraassen says in The Scientific Image, the explanation of something depends not only what caused something to happen, but also in your expectation of what could have happened instead. Obviously.

7. The Prestige

Let's turn now to the core of the Kirschner, Sweller and Clark paper, which is the 'cognitive load argument'. This is where you're going to give you cognitive psychology and tell you, here's how learning works.

So their theory is based on the theory of long term and short term memory. They're not going to be interested in sensory memory (I find that a very interesting statement I'm not going to linger on) and the manner in which our cognitive structures are organized. You've heard this before, right? You have short term memory, you have long term memory. Short term memory is the stuff that we are consciously aware of. Long term memory is the stuff we aren't.

Long term memory is - and these are their words - a 'massive knowledge base' and anyone who understands cognitive structures knows how bad a statement that is. If you looked at the structure of the mind, it does not resemble a knowledge base at all. And I won't get into the details of it, but neural structures and databases are two very different things.

And they say, you're skillful in an area because your long term memory contains huge amounts of information, so it's a theory knowledge based on quantity, piled fact on fact on fact. And then they cite – astonishingly – DeGroot's work on chess expertise. But, you know, I play chess. And there's this fiction that chess players who are really good chess players can predict ahead, and they can keep all these different board positions in their head. But you cannot defeat a computer like that. The computer will always predict ahead further than you. But people still can beat computers because they visualize, they recognize, successful formations. They don't memorize a whole bunch of chess positions; they see what is going to work and what is not going to work.

This is Kirschner, Sweller and Clark again: they say the alternate means of instruction is justified by this cognitive theory of long term memory. And the aim of all instruction they say is to alter (nice word there, 'alter') long term memory. If nothing has been changed in long term memory they said nothing has been learned. It harkens back to the logical positivist principle of verificationism, that a difference that makes no difference is a difference at all. Or in Bergmann's formulation, in the dark, all cows are black.

And then the other side of the theory is that working memory, short term memory (and we've seen the research on this) has limited capacity, has limited duration, only lasts for a few minutes and is restricted to a small number of elements - the famous number seven elements. We've done lots of work on that. Seven elements doesn't mean seven digits because if you taught to cluster things you can actually remember more than seven. That's going to form an important part of their theory.

Here's where the prestige comes in (that's from magic, right, the movie, you do your trick and then you do the unexpected thing and it all comes together and all the crowd goes ooh ahh). Here's where it all comes together.

Most learners can, to use their words, construct knowledge. People can actually spot the generalizations. And to do this you must construct a representation. Now this representation, they say, is equally good whether you're using full information or partial information. So if you have the full information you're really not losing anything, but if you have partial information, part of your short term memory is occupied in trying to search for possible theories or hypothesis. And because you're searching for what theory this data could fit into, you're overloaded, your short term memory is overloaded and you are not able to focus and actually think about what you're supposed to be learning.

And so, constructivist – because it's making you search for these theories - is not a good prescriptive instructional designed theory. It's too hard.

8. Personal Knowledge

Well it's all in the search, isn't it? And this is the putative 'worked example effect', where you're doing a search for the hypothesis (imagine a logic problem, "oh what principle am I going to apply, am I going to apply transposition, composition or whatever?"). So you're searching these hypotheses, and it's unrelated to learning, you're trying to find the best fit, but if you're studying a worked example, the theory has already been picked for you and you're looking at what has been done. You're learning by being told.

The problem is, problem solving - real human problem solving - doesn't work that way.

We don't do this mental search of our internal hard drive and try to find the right theory. That's not how it works at all. We don't have a whole bunch of general principles or theories stored in our head. Rather we looked at the data and we recognize patterns in that data. It's a process of seeing rather than searching. This is why is so interesting that at the beginning of this they dropped perceptual memory as though it was completely irrelevant. But perceptual memory is totally relevant because perceptual memory is a process of recognition.

There's even in the literature there's discussion of this (I'm not just not making this up). Stephen Kosslyn, for example, image-rotating examples, the theory where you're going through algorithms and processes does not explain reaction time when you're asking people to visualize the rotation of images. We work on a sub-symbolic level, not a symbolic level. Cognition is based on a process of pattern recognition at that sub-symbolic level.

And we know this. There has been a lot of discussion about experts and schemas and how an expert will acquire a frame – to use Lakoff's term – or knowledge of organization. And they even admit, the authors even admit, cognitive research has shown that to acquire expertise in a domain the learners must acquire the necessary schemata.

But what are schemata? What is this 'picture of domain' that a person must have. They (Kirschner, Sweller, Clark) think it is facts, laws, principles, and theories that make up the discipline's content. But this is simply wrong. It is an incorrect understanding of science. And they ought to know that.

Bloom's Taxonomy talks about types of learning that are not facts, data, and all theories. Or Michael Polanyi, his book *Personal Knowledge*, talking about personal knowledge, the difference between 'knowing that' and 'knowing how'. And what's important about Polanyi's work is that he says, the bulk of our knowledge, even our conceptual knowledge, is ineffable. That means it cannot be represented in words, which means that a statement (which is a theory) is not a good expression of that knowledge, and that a law is not a good expression of that knowledge. The configuration of connections in your brain, that is a good representation of that knowledge.

Knowing a discipline is knowing the practice of that discipline. It's learning to think like a scientist or a forester or a hockey player, and learning to recognize, to see the way they see, to speak the way they speak their words, to judge the way they judge.

And that's what we are producing in the read-write web. The read-write web, the web in which we talked about what we're doing, we reflected, we practiced, is developing these new kinds of literacies, these new ways of people being able to express how people think - that is, how they talked, how they practiced, how they judged, how they evaluated. There's a nice picture of this, the revision of Bloom's verbs that includes verbs tied to the social infrastructure skills and abilities that help people learn and help people grow new mental configurations, grow in knowledge in this environment.

The pedagogy in this environment is based on personal learning. It's based on the acquisition, the developing, the growth of capacities and aptitudes - like recognition - rather than laws and facts and theories and data. I tried to talk about this a little bit in a paper I presented the last time I was here in Holland called *Things You Really Need to Learn*. And I talked about the general principle, how to predict consequences. How to stay healthy. How to live a meaningful life. These are the core principles, the core things that people need to learn, not facts and data.

I've talked about the semantic principle, the mechanisms, the ways we design our networks in order to make the most reliable system for recognizing patterns in the environment and these principles I put under the headings of Autonomy, Diversity, Openness, and Interactivity. And I've talked about the principle of personal learning, the idea here that you develop your own learning yourself. You make your own learning and the way you make your own learning is you find the relevance in the environment. And that has to do with similarity or salience and Amos Tversky, interaction, communicating with your community, and then usability, just being able to speak the language, use the interactions.

That's the theory that is really supported by, if you will, the cognitive architecture. And so that's my refutation of Kirschner, Sweller and Clark and my presentation of the alternative connectivist theory of

knowledge. And I thank you very much for your time.

1. Paul A. Kirschner, John Sweller, Richard E. Clark. 2006. Why Minimal Guidance During Instruction Does Not Work: An Analysis of the Failure of Constructivist, Discovery, Problem-Based, Experiential, and Inquiry-Based Teaching. *Educational Psychologist*, 41(2), 75-96. Lawrence Erlbaum Associates, Inc. http://www.cogtech.usc.edu/publications/kirschner_Sweller_Clark.pdf
2. George Siemens. 2005. Connectivism: A Learning Theory for a Digital Age. *International Journal of Instructional Technology and Distance Learning*, Volume 2, Number 1, January, 2005. http://www.itdl.org/Journal/Jan_05/article01.htm
3. Stephen Downes. 2007. Kirschner, Sweller, Clark (2007): Summary. Half an Hour. <http://halfanhour.blogspot.com/2007/11/kirschner-sweller-clark-2006-summary.html>
4. Stephen Downes. 2007a. Kirschner, Sweller, Clark (2007): Readings. Half an Hour. http://halfanhour.blogspot.com/2007/11/kirschner-sweller-clark-2006-readings_12.html
5. Advanced Distributed Learning. Website. <http://www.adlnet.gov/>
6. Stephen Downes. 2004. Buntine Oration: Learning Networks. Delivered to the at the Australian College of Educators and the Australian Council of Educational Leaders conference in Perth, Australia, October 8, 2004. <http://www.downes.ca/post/20>
7. Cindy E. Hmelo-Silver, Ravit Golan Duncan, and Clark A. Chinn. 2006. Scaffolding and Achievement in Problem-Based and Inquiry Learning: A Response to Kirschner, Sweller, and Clark (2006). *Educational Psychologist* 42(2), 99-107. http://www.cogtech.usc.edu/publications/hmelo_ep07.pdf
8. Mark Rickards. 2005. Teach-yourself computing for kids. BBC News, may 2, 2005. http://news.bbc.co.uk/1/hi/world/south_asia/4498511.stm
9. Deanna Kuhn. 2007. Is Direct Instruction an Answer to the Right Question? *Educational Psychologist* 42(2), 109-113. http://www.cogtech.usc.edu/publications/kuhn_ep_07.pdf
10. Miles Berry. 2006. Why minimally guided instruction does not work. Weblog. <http://elgg.net/mberry/weblog/124841.html>

The Science of Learning

As [Clive Shepherd writes](#), "cognitive neuroscientist Dr Itiel Dror of Southampton University. Itiel is becoming a bit of a celebrity amongst the e-learning community in the UK as someone who avoids the grand theories of learning and concentrates instead on practical tips based on what we know about the brain and how it works (assuming we really do and this I must place on trust)."

Of course, the sceptical side of me says that this is something like saying that so-and-so can tell us best how to win an auto race because he's a mechanic. There is, indeed, a distinction between knowing how something works and knowing how best to use it.

Reading through the points in this summary, they seem *sort of* right, but not exactly right. Let me clarify them.

- *The brain is a machine with limited resources for processing the enormous quantity of information received by the senses. As a result, attention is extremely selective and the brain must rely on all sorts of shortcuts if it is to cope effectively.*

My response: no criticism of this; it seems to be about right.

- *Teachers/designers can adopt two strategies to reduce the risk of learners experiencing **cognitive overload**: provide less information (quantitative approach) or take much more care about how this information is communicated (qualitative approach).*

Well, you see now, this approaches the problem not from the nature of the brain but rather from the nature of the *information*. And when we look at the information as nothing more than a pile of stuff to be processed by the brain, then sure, these are the ways to deal with it.

But the other way to look at it is to, as promised, look at it from the perspective of neuroscience. What does the *brain* do in cases of cognitive overload? This is important because, if we know how the brain will adapt, we know *how* to shape our information (if at all).

This is the subject of the next few points, so I'll continue.

- *It is easier for a person to focus their intention on the desired point if there is minimal noise (other information) surrounding it. Reducing noise also reduces context, so a balance needs to be struck.*

I assume he meant 'attention' and not 'intention'. In any case, I'm sure there are all kinds of tests proving this, but I will point out that the *nature of the subject* is a much more significant variable.

People are able to focus on things even in the most extreme of circumstances if they are sufficiently interested. That's how you can have kids playing video games even while the house is burning down around them (I guess that's the sort of 'context' that would be important). By contrast, if you aren't really interested in what you are doing, the least amount of noise distracts you.

We know this because (as was just stated above) we know that the brain is extremely selective and filters out stuff that isn't important.

Perspective matters. From the teacher's point of view, the content (lessons or curriculum) is constant, while the level of background noise is the variable. From the point of view of the learner, however, the content is also variable. That's why you get two very different interpretations of the same phenomena.

- *Overload can be reduced by **grouping** items/steps (what Itiel calls 'chunking'). Grouping can be accomplished by placing people/objects/events into categories, or by compressing a number of procedural steps into one, automatic action. Visually you may separate items by space, size or colour. Learners will naturally employ grouping as a strategy, although they may do this inappropriately and the process requires effort. Better for the designer/teacher to present material ready grouped.*

This is a good strategy and one I have recommended elsewhere to help people write academic essays easily and proficiently (and without notes, but I digress). I find it interesting, though, that he used the DE design term 'chunking'. Maybe reading something other than neuroscience?

Yes, there are different types of groups. Groups that make sense conceptually, especially if linked to a larger framework, are better (I would add that colour is rarely, if ever, a part of such a framework).

But is it better for the teacher to present the material already grouped? How does that follow? If the intent is to have the student learn the information (ugh, bad terminology) then we *must* ask, is it the *groups* that aide remembering and understanding, or the *process of grouping* that does this? If it's the latter, then presenting the information already grouped may help the teacher remember, but will do nothing for the student.

Because, as I noted above, it is better if the groups align with a pre-existing conceptual framework, it is better then if the *student* does the grouping, because that way the process allows the student to connect, in an organized way, new knowledge with existing knowledge.

- *A side effect of grouping is that once the action is completely familiar (that old 'unconscious incompetence' phase), the individual finds it hard to explain how they do it; they lose control*

over the process because it has become automatic (so old hands may not always be the best teachers?). Grouping is essential to our functioning, but there are obvious dangers, i.e. unhelpful stereotyping.

Here there seems to be a confusion between grouping, as in the sense of classifying different perceptual entities into types, and grouping, as in the sense of combining several activities into one. Now this isn't necessarily bad (I have said elsewhere that learning to read is similar to learning to ride a bicycle) but can be very misleading unless carefully explained.

I think it would have been better to present them separately.

There are mental *processes* that can become automatic. Add 1+1 for example. One of these processes is 'categorization'. You look at a bunch of things and automatically associate some with the others, based on habitually formed patterns of association. In some cases, such as grouping people by colour, this sort of automatic association can be inappropriate.

There are also mental processes that constitute sequences of steps. The steps involved in a logical derivation, for example. So a process that actually involves multiple steps may be performed by an experienced logician as though it were only one step (I called this 'skipping steps' in logic class and complained bitterly about it. "It's obvious," said the professor. "Whaaaaa?" I responded).

These are very different phenomena that are essentially the result of the same neural process but which instantiate very differently and need to be approached very differently. Kind of like the way the steering used to recover from a spinout may be exactly the same as the steering required to navigate a hairpin curve. Sure, it's the same motion. But you would describe the two events very differently.

- *Individuals use **top-down processing** to reduce overload. This draws automatically on their past experience of the particular context, existing knowledge and intelligence and avoids them having to evaluate all new information from the bottom up. An example would be how people can easily read a sentence in which the letters in each word are jumbled up.*

Yes. But...

This is not 'top down' processing as traditionally understood.

There is a *very large* difference between inferring something on the basis of similarity to a prototype (that is, apttern recognition), and inferring something based on a general principle or rule. By 'top down' we typically mean the latter. But when describing character recognition, as in the example, we are describing the former.

I would also be wary of building the (Darwinian?) intent into the process. People use pattern

recognition. It reduces information overload. But it is not necessarily true that people use pattern recognition *in order to* reduce information overload. People use pattern recognition because that's how neural networks work. Perhaps evolution directed us in this way, perhaps it did not. Either way, our use of pattern recognition in a particular circumstance is not *caused* by some such intent. It occurs naturally, as though by habit.

- *Designers/teachers need to take account of the way in which the information is likely to be encoded and processed - it's not 'what you teach' but 'what is learned'.*

Except... it is very misleading to say that it's 'encoded'. Otherwise, yes, there is a large distinction between what the teacher teaches and what the learner learns (which is why information-theoretic and transmission-theoretic theories of learning are wrong).

- *Different parts of the brain specialise in different tasks. Individuals can engage in more than one task at the same time, as long as each uses a different part of the brain.*

Of course, these parts of the brain form dynamically and according to experience and circumstance, so there's no telling in advance, or in general, what processes occupy the 'same' part of the brain and what processes do not.

That's why I can read and write while listening to loud music, as I'm doing now, while my father couldn't.

- *It's a myth that we only use 5-10% of the brain - we use it all.*

Correct.

- *The brain continues to **change** throughout our lives, even though we stop adding new brain cells in our early 20s. Some parts of the brain are relatively hard-wired (through nature or nurture), some very plastic. It makes sense to concentrate in recruitment on finding those people with hard wiring which suits the job, because no amount of training will sort the problem out later. (Itiel did not go into detail about those capabilities which tend to be hard-wired and those which are more plastic - this is clearly important.)*

The main part - that neural nets are plastic - is true and important.

It is also true that some parts are pretty much hard-wired -- good thing, too, or our hearts wouldn't beat and our eyelids wouldn't blink.

But as to how much this carries over into learning or into life skills - this is very controversial. I can certainly agree that there are people with *currently existing wiring* that may be more or less suited to

the job. That's no more controversial than saying people learn different things. But to say that these capabilities are hard-wired is much more questionable.

- *As you grow older the hard-wired capabilities persist - the most learnable capabilities go first.*

This is demonstrably false. For otherwise there would be no incontinence in old people.

- **Language** *is more than just a means for expressing thought - in many ways it is thought. If a person is not exposed to any language in early years, then by the age of seven they are incapable of learning it.*

I doubt that this is uniquely true of language - it is probably true for any pattern set. Can a person become fluent in mathematics despite never having been exposed to numbers. Can a person become a musician having never been exposed to tone and melody?

People who have not specialized in the nature of language typically take language as a given - some sort of folk-psychological representation of Chomskyan generative grammar. And then suppose that this then must be the nature of thought.

Even if language is thought - which i would not grant for a second - we still know *nothing* about the nature of thought if we do not agree on the nature of language. Which we probably most emphatically do not.

- **The two sides of the brain** *really do have different functions (I thought this was just pop psychology). The left brain concentrates on language and analytical skills; the right has the spacial abilities. The left side of the brain controls the right side of the body and vice versa. The left and right sides of the brain do not interact physically.*

They do interact, through something called the 'corpus callosum'. But yes, the two sides of the brain do specialize, as observed. That said, to my knowledge, this specialization is not hard-wired.

- **The size** *of a person's brain is not an indicator of intelligence.*

Within the normal variation of human brains, that is. My poor cat, with her cat-sized brain, will never reach human intelligence. But she's still very cute.

- *20% of your blood is in the brain.*

Which stresses the importance of nutrition to brain function.

- *You never lose anything from long-term memory, just the ability to retrieve it. Retrieval is a function of how you encode memories / the number of links you provide.*

Well, yeah, in the sense that the connections that constitute the long-term memory are basically permanent, more or less. But you don't 'retrieve' memory the way you retrieve a book from the bookshelf (even though it feels that way).

'Retrieval' (properly-so-called) is a case of pattern recognition - and the less salient a pattern becomes in the mind, the less likely it is to be associated with a current perception.

- **Working memory** consists of 7 ± 2 items (again I thought this was pop psychology).

Yes. And there's that 'grouping' again, in a third guise. We can remember things that are more than 7 words long by recognizing them as coherent patterns. That's why I can remember something like 'Turn right at the third light and then left at the second stop sign, then go four blocks' even though it consists of 18 words (and 70 letters). or 1-800-857-2020 even though it's 11 digits long.

What we put into working memory *first* depends on pattern recognition.

- *To reduce cognitive overload, take out every word or picture that is not necessary or relevant to your learning goals. Even then, don't deliver more than the learner can handle (presumably by modularising the learning).*

This is effective only at the very gross level, and not particularly useful as you get into finer details (and more precise definitions of 'necessary').

I have seen studies, for example, showing that a slide show with bullet points is more easily remembered than a slide show with the same bullet points and animated graphics.

I expect, though, that the placement of the 'NRC' logo in the corner of the same slides would not have an impact either way.

I also expect that the removal of 'unnecessary' letters in the bullet points would actually *hinder* memory. Fr exmpl, rmving mst f th vwls. Thy r nt necsry in the snse that we cn stl rd the sntnce, bt they hlp wth the grping.

So - better advice would be something like - present material that accords (perhaps with some cognition) with patterns that will already be familiar to the learner.

- *Provide the learning when it is needed, not before.*

Sure. But *why*? This isn't determined by whether it is *necessary* or not, but rather, by whether it is *salient* or not.

- Be **consistent** in the manner of your presentation, e.g. the interface.

This can actually be distracting, taken to extremes. That's why documentaries switch from having a person talk to showing some nature scenes with a voiceover to interviewing some other person.

- Be consistent in the level of your presentation, i.e. not too complex, not too simple. Try to work with homogeneous groups; better still personalise the learning.

Yes, but again, *why*? I would argue that this facilitates pattern matching .

- **Engage** the learner by grabbing their attention, allowing them to determine their progress, providing constructive feedback, introducing an element of excitement/surprise.

Again, *this doesn't follow* from the presentation above. What has happened here is that some hackneyed (and vague) pedagogical tips have been attached to some discussion of neural function, without a clear linkage between them.

- Be careful of allowing the learner too much **control** over the learning process if they don't have the metacognitive skills, i.e. they don't know what they know and what they don't know, nor how best to bridge the gap. Ideally help learners to increase their metacognitive skills, i.e. learning how to learn.

This has utterly *nothing* to do with the brain theory discussed above.

If the *content* has more to do with attention than, say, distraction, then taking control away from learners, even in areas where they do not have skills, may cause more harm than good.

And what *are* the metacognitive skills. What is 'learning how to learn', for example?

- Providing the learner with control over pace and allowing them to go back and repeat any step is important.

Same point.

- The learning benefits by being **challenging**. Performance targets, rewards and competition can increase the degree of challenge, perhaps through the use of games.

And again, same observation. This doesn't follow from what has been stated above. Sure, it's good advice (what Seymour Papert and James Paul Gee call 'hard fun'). But why is it good advice. What else can we learn about this piece of advice. What kind of games, for example (see Aldrich on this).

Anyhow, those are my thoughts based on this reading of Shepherd's article. I also read a bunch of Dror's publications online and certainly have no quibble with his neuroscience. I just think that the study of teaching and learning involves more than just neuroscience, and that there are areas of complexity and potential confusion Dror may not have considered in his work.

E-Learning 2.0

First published as E-Learning 2.0 in E-Learning Magazine, October 17, 2005.

<http://www.elearnmag.org/subpage.cfm?article=29-1§ion=articles>

E-learning as we know it has been around for ten years or so. During that time, it has emerged from being a radical idea—the effectiveness of which was yet to be proven—to something that is widely regarded as mainstream. It's the core to numerous business plans and a service offered by most colleges and universities.

And now, e-learning is evolving with the World Wide Web as a whole and it's changing to a degree significant enough to warrant a new name: E-learning 2.0.

Where We Are Now

Before talking about where e-learning is going, it is worth spending a few words to describe here we are now.

When we think of learning content today, we probably think of a learning object. Originating in the world of computer-based delivery (CBT) systems, learning objects were depicted as being like lego blocks or atoms, little bits of content that could be put together or organized. Standards bodies have refined the concept of learning objects into a rigorous form and have provided specifications on how to sequence and organize these bits of content into courses and package them for delivery as though they were books or training manuals.

Today, e-learning mainly takes the form of online courses. From the resources distributed by MIT's [OpenCourseware project](#) to the design of learning materials in [Rice's Connexions project](#) to the offerings found from colleges and universities everywhere, the course is the basic unit of organization.

As a consequence, the dominant learning technology employed today is a type of system that organizes and delivers online courses—the learning management system (LMS). This piece of software has become almost ubiquitous in the learning environment; companies such as WebCT, Blackboard, and Desire2Learn have installed products at thousands of universities and colleges and are used by tens of thousands of instructors and students. The learning management system takes learning content and organizes it in a standard way, as a course divided into modules and lessons, supported with quizzes, tests and discussions, and in many systems today, integrated into the college or university's student information system.

In general, where we are now in the online world is where we were before the beginning of e-learning [\[1\]](#). Traditional theories of distance learning, of (for example) transactional distance, [as described by Michael G. Moore](#), have been adapted for the online world. Content is organized according to this traditional model and delivered either completely online or in conjunction with more traditional seminars, to cohorts of students, led by an instructor, following a specified curriculum to be completed at a predetermined pace.

Trends

As we approach the halfway mark of the new millennium's first decade, the nature of the Internet, and just as importantly, the people using the Internet, has begun to change. These changes are sweeping across entire industries as a whole and are not unique to education; indeed, in many ways education has lagged behind some of these trends and is just beginning to feel their wake.

One trend that has captured the attention of numerous pundits is the changing nature of Internet users themselves. Sometimes called "digital natives" and sometimes called "n-gen," these new users approach work, learning and play in new ways [\[2\]](#).

They absorb information quickly, in images and video as well as text, from multiple sources simultaneously. They operate at "[twitch speed](#)," expecting instant responses and feedback. They prefer random "on-demand" access to media, expect to be in constant communication with their friends (who may be next door or around the world), and they are as likely to create their own media (or download someone else's) as to purchase a book or a CD [\[3\]](#).

The manner in which this new generation of users is changing markets is captured evocatively in a document called [The Cluetrain Manifesto](#). First posted online in April 1999, the document begins with the declaration that "markets are conversations" and continues with a redefinition of the relation between producer and consumer. "Markets are getting smarter, more informed, more organized... People in networked markets have figured out that they get far better information and support from one another than from vendors." Jay Cross, writing in the same vein, talks about the "augmented learner" and the "hyper-organization" [\[4\]](#).

In learning, these trends are manifest in what is sometimes called "learner-centered" or "student-centered" design. This is more than just adapting for different learning styles or allowing the user to change the font size and background color; it is the placing of the control of learning itself into the hands of the learner [\[5\]](#).

"The changing demographics of the student population and the more consumer/client-centered culture in today's society have provided a climate where the use of student-centered learning is thriving" [\[6\]](#). Learning is characterized not only by greater autonomy for the learner, but also a greater emphasis on

active learning, with creation, communication and participation playing key roles, and on changing roles for the teacher, indeed, even a collapse of the distinction between teacher and student altogether [7].

Taking this approach even further is George Siemens's [Connectivism](#). "We derive our competence," writes Siemens, "from forming connections... Chaos is a new reality for knowledge workers... Unlike constructivism, which states that learners attempt to foster understanding by meaning-making tasks, chaos states that the meaning exists— the learner's challenge is to recognize the patterns which appear to be hidden. Meaning-making and forming connections between specialized communities are important activities." Readers of Douglas Rushkoff's *Cyberia* will recognize a similar theme as knowledge-working is no longer thought of as the gathering and accumulation of facts, but rather, the riding of waves in a dynamic environment [8].

The breaking down of barriers has led to many of the movements and issues we see on today's Internet. File-sharing, for example, evolves not of a sudden criminality among today's youth but rather in their pervasive belief that information is something meant to be shared. This belief is manifest in such things as [free and open-source software](#), [Creative Commons licenses](#) for content, and [open access](#) to scholarly and other works. Sharing content is not considered unethical; indeed, the hoarding of content is viewed as antisocial [9]. And open content is viewed not merely as nice to have but essential for the creation of the sort of learning network described by Siemens [10].

Numerous writers, even, have called for what is often referred to as the "open society." Tapscott, for example, writes about ["the transparent burger"](#) and ["the naked corporation."](#) Mougayar tells us that "the future organization is an ["open corporation."](#) And in a widely popular online essay Rob Paterson asked, "Is not the new "big idea" of our time to disintermediate the institutional middleman and to enable direct relationships? Are supermarkets eternal? Do we need factory universities to learn? Is our health dependent on a doctor? Is the news what we see on TV?" [11].

In short, the structures and organization that characterized life prior to the Internet are breaking down. Where intermediaries, such as public relations staff, journalists or professors, are not needed, they are disregarded. Consumers are talking directly to producers, and more often than not, demanding and getting new standards of accountability and transparency. Often, they inform the productive process itself, and in many cases, replace it altogether. Passive has become active. Disinterested has become engaged. The new Internet user may not vote, but that is only because the vote is irrelevant when you govern yourself.

The Web 2.0

The first sign that something was changing on the Web was the underground popularity of a site called [LiveJournal](#) and the very visible surge of interest in a site called [Friendster](#). These sites, which came to be called "social networking sites," were rapidly emulated by such services as [Tribe](#), [LinkedIn](#), [Google's Orkut](#), [Flickr](#), and [Yahoo 360](#). Writers conversant with the works of social network analysts, people like

[Duncan J. Watts](#) and [Mark Buchanan](#), for example, noticed that similar patterns existed in these online networks [\[12\]](#). Something was happening here.

What was happening was that major parts of the World Wide Web were acquiring the properties of communications networks, the sorts of networks found to exist (albeit on a much smaller scale) in the physical world. And that the Web itself was being transformed from what was called "the Read Web" to the "[Read-Write Web](#)," in accordance with [Tim Berners-Lee's original vision](#). Proponents of this new, evolving Web began calling it Web 2.0 and in short order the trend became a movement.

"Enter Web 2.0, a vision of the Web in which information is broken up into "microcontent" units that can be distributed over dozens of domains. The Web of documents has morphed into a Web of data. We are no longer just looking to the same old sources for information. Now we're looking to a new set of tools to aggregate and remix microcontent in new and useful ways" [\[13\]](#).

In a nutshell, what was happening was that the Web was shifting from being a medium, in which information was transmitted and consumed, into being a platform, in which content was created, shared, remixed, repurposed, and passed along. And what people were doing with the Web was not merely reading books, listening to the radio or watching TV, but having a conversation, with a vocabulary consisting not just of words but of images, video, multimedia and whatever they could get their hands on. And this became, and looked like, and behaved like, [a network](#).

Nowhere is this clearer than in the world of blogging. In a few short years the blog went from a few idiosyncratic Web sites to something used by millions of people empowered by content creation tools such as [Blogger](#) and [Wordpress](#). Even more importantly, these blogs were *connected* to each other through the mechanism of [RSS](#), a simple XML format that allows bloggers to send their content to a network of readers (called 'subscribers').

But it wasn't just blogging. Creating an online community became a snap with tools such as [Plone](#) and [Drupal](#). Moreover, using a collaborative writing tool called the [wiki](#) Jimmy Wales and a few thousand of his friends created a site called [Wikipedia](#), rendering Encyclopedia Britannica obsolete in the process. Others, using the free audio-recording tool [Audacity](#), began recording their own talk and music; this, when combined with RSS, became [podcasting](#), a rapidly rising phenomena that is transforming what we think about radio.

For all this technology, what is important to recognize is that the emergence of the Web 2.0 is not a technological revolution, it is a social revolution. "Here's my take on it: Web 2.0 is an attitude not a technology. It's about enabling and encouraging participation through open applications and services. By open I mean technically open with appropriate APIs but also, more importantly, socially open, with rights granted to use the content in new and exciting contexts" [\[14\]](#).

E-Learning 2.0

In the world of e-learning, the closest thing to a social network is a community of practice, articulated and promoted by people such as [Etienne Wenger](#) in the 1990s. According to Wenger, a community of practice is characterized by "a shared domain of interest" where "members interact and learn together" and "develop a shared repertoire of resources."

For the most part, though, what constituted "community" in online learning were artificial and often contrived "discussions" supported by learning management systems [\[15\]](#). These communities were typically limited to a given group of learners, such as a university class, had a fixed start and end-point, and while substantially better than nothing, rarely approached Wenger's theory.

That's not to say no communities of practice were forming. There were some attempts to foster them, as for example [MuniMall](#), directed toward the municipal governance sector, and [PEGGasus](#), directed toward engineers and geophysicists. Moreover, as commentator [Erin Brewer](#) has noted, places on the Internet like Yahoo! Groups have become a locus for community learning activities. But in general, the uptake has been slow, and the support from traditional institutions almost nonexistent.

Educators began to notice something different happening when they began to use tools like wikis and blogs in the classroom a couple of years ago. All of a sudden, instead of discussing pre-assigned topics with their classmates, students found themselves discussing a wide range of topics with peers worldwide. Imagine the astonishment, for example, when, after writing a review of a circus she had viewed, a Grade 5 student received a response from one of the performers [\[16\]](#). In a very short time, blogs were used for a wide variety of purposes in [education](#); an [educational bloggers' network](#) formed and by this year thousands of teachers were encouraging their students to blog.

Blogging is very different from traditionally assigned learning content. It is much less formal. It is written from a personal point of view, in a personal voice. Students' blog posts are often about something from their own range of interests, rather than on a course topic or assigned project. More importantly, what happens when students blog, and read reach others' blogs, is that a network of interactions forms-much like a social network, and much like Wenger's community of practice.

It's not just blogging. Educators have also taken an interest in podcasting. Some have started broadcasting, such as at McMaster, where engineering professors now host an online show [\[17\]](#). "We're talking to the download generation," said Peter Smith, associate dean, Faculty of Engineering. "Why not have the option to download information about education and careers the same way you can download music? It untethers content from the Web and lets students access us at their convenience." Moreover, using an online service such as [Odeo](#), [Blogomatrix Sparks](#), or even simply off-the-shelf software, students can create their own [podcasts](#).

What happens when online learning ceases to be like a medium, and becomes more like a platform? What happens when online learning software ceases to be a type of content-consumption tool, where learning is "delivered," and becomes more like a content-authoring tool, where learning is created? The model of e-learning as being a type of content, produced by publishers, organized and structured into courses, and consumed by students, is turned on its head. Insofar as there is content, it is used rather than read—and is, in any case, more likely to be produced by students than courseware authors. And insofar as there is structure, it is more likely to resemble a language or a conversation rather than a book or a manual.

The e-learning application, therefore, begins to look very much like a blogging tool. It represents one node in a web of content, connected to other nodes and content creation services used by other students. It becomes, not an institutional or corporate application, but a personal learning center, where content is reused and remixed according to the student's own needs and interests. It becomes, indeed, not a single application, but a collection of interoperating applications—an environment rather than a system.

It also begins to look like a personal portfolio tool [\[18\]](#). The idea here is that students will have their own personal place to create and showcase their own work. Some e-portfolio applications, such as ELGG, have already been created. IMS Global has put together an e-portfolio specification [\[19\]](#). "The portfolio can provide an opportunity to demonstrate one's ability to collect, organize, interpret and reflect on documents and sources of information. It is also a tool for continuing professional development, encouraging individuals to take responsibility for and demonstrate the results of their own learning" [\[20\]](#).

This approach to learning means that learning content is created and distributed in a very different manner. Rather than being composed, organized and packaged, e-learning content is syndicated, much like a blog post or podcast. It is aggregated by students, using their own personal [RSS reader](#) or some similar application. From there, it is remixed and repurposed with the student's own individual application in mind, the finished product being fed forward to become fodder for some other student's reading and use.

More formally, instead of using enterprise learning-management systems, educational institutions expect to use an interlocking set of open-source applications. Work on such a set of applications has begun in a number of quarters, with the [E-Learning Framework](#) defining a set of common applications and the newly formed [e-Framework for Education and Research](#) drawing on an international collaboration. While there is still an element of content delivery in these systems, there is also an increasing recognition that learning is becoming a creative activity and that the appropriate venue is a platform rather than an application.

In the future it will be more widely recognized that the learning comes not from the design of learning content but in how it is used. Most e-learning theorists are already there, and are exploring how

learning content—whether professionally authored or created by students— can be used as the basis for learning activities rather than the conduit for learning content.

A great amount of work is being done, for example, in educational gaming and simulations. Theorists such as Seymour Papert, James Paul Gee, Clark Aldrich, and Marc Prensky have all touted the efficacy of games.

Papert writes, "The most important learning skills that I see children getting from games are those that support the empowering sense of taking charge of their own learning. And the learner taking charge of learning is antithetical to the dominant ideology of curriculum design" [21]. This is most evidenced when learners engage not only in playing, but in the design, of games. In the gaming world this practice is widely recognized and encouraged—game "modding" allows players to make the game their own. [22].

Where games encourage learning is through the provision of what a student needs to know in a context where it will be immediately used. As Gee recommends, "Words are only meaningful when they can be related to experiences," said Gee. If I say "I spilled the coffee," this has a different meaning depending on whether I ask for a broom or a mop. You cannot create that context ahead of time— it has to be part of the experience. And in just the same way, the science text doesn't make any sense to someone who has not done any science (though it makes a great deal of sense to an experienced scientist)" [23].

A similar motivation underlies the rapidly rising domain of mobile learning [24]—for after all, were the context in which learning occurs not important, it would not be useful or necessary to make learning mobile. Mobile learning offers not only new opportunities to create but also to connect. As Ellen Wagner and Bryan Alexander note, mobile learning "define(s) new relationships and behaviors among learners, information, personal computing devices, and the world at large" [25].

As this trend progresses, we find ourselves in a world characterized by the phrase "ubiquitous computing." "Where virtual reality puts people inside a computer-generated world, ubiquitous computing forces the computer to live out here in the world with people" [26]. The "Father of ubiquitous computing," Mark Weiser, compares computing of the future to writing. "Today this technology is ubiquitous in industrialized countries. Not only do books, magazines and newspapers convey written information, but so do street signs, billboards, shop signs and even graffiti" [27].

In the world of learning, what this means is having learning available no matter what you are doing. Jay Cross captures this idea in the concept of "workflow learning." Sam Adkins writes, workflow learning is "a deep integration with enterprise applications assembled from Web Services into composite applications" with "task and work support fused into the aggregated business processes that make up the real-time workflow" and supported by "contextual collaboration with people and systems" and "design and modification achieved by modeling and simulation" [28].

Of course, there is no reason to expect that this form of learning would be restricted to the workplace. Learning integrates into every aspect of our lives, from daily household chores to arts and culture. Learning and living, it could be said, will eventually merge. The challenge will not be in how to learn, but in how to use learning to create something more, to communicate.


To The School or Classroom 2.0 Advocates

Christian Long asked, by email, for input. Quoted text is his. Here is my reply:

Christian writes,

In advance of a major keynote presentation called "Designing School 2.0" I'm giving in 2 weeks to a room full of "school design" decision makers -- school architects, their educational clients (board members, superintendents, and administrators), and various stakeholders -- I'd like to ask each of you to consider sharing what you'd say to them if you have 5 minutes and the microphone.

The basis of my talk, if I were giving it, would be the alternative School 2.0 document I distributed in response to the 'official' School 2.0 document.



http://www.flickr.com/photo_zoom.gne?id=268691876&size=l

You can focus your ideas from kindergarten through higher education, although the majority of my audience will focus on K-12 on most days. Likewise, you can think macro answers or burrow down to the library shelves. All feedback is good feedback.

There is no particular focus for this view of 'Scool 2.0'. The main point is that technology allows us to change our approach to education, from one where we segregate learners in specially designed education facilities (classrooms, training rooms, schools, universities) to one where learning is something we do (and what educators provide) in the course of any other activity.

The idea is that 'School 2.0' is the first step toward being non-school, and that our objective should be to use technologies to leverage our ability to personalize learning, and in so doing, facilitate students' learning while taking part as full citizens in the wider community.

Here are the areas I'd love specific feedback from you today (although I will have the blessing of extrapolating many ideas/resources from many of your blogs, wikis, research, writings, etc.): Big Picture trends in the next 5-25 years that will have the biggest impact on what it means to be an engaged learner. This is the firecracker side of things. I'll be sharing Karl Fisch's/Scott McLeod's collaboration to this audience assuming they are in the Media/AV 1.0 world, at least. Never can be sure no matter how nice the hotel/conference center appears to be on the surface!

I have commented in the past, and I reiterate the point here, that from my perspective the predominate use of the term 'School 2.0' has been to promote a view of learning that is traditionalist, rather than oriented to the future, one that seeks to preserve the existing trappings of education, most notably, schools. We hear a lot of language like "the fact is, schools are here to stay," but there is in my mind no fact of the matter, certainly not in the time-frame of 25-30 years.

By way of animation, think back to the time of the one-room schoolhouse as they were rolled out across a rural agrarian nation in the late 1800s and early 1900s. Within a generation society transformed itself from one in which 'book learning' was the exception to one in which it was the norm. Read Jane Eyre, and learn how schooling was conducted - if it was conducted at all - by tutors hired to live in and help raise the children. People my own grandparents' age would routinely not attend school at all. Only in my father's era, the war years, did schooling become the norm for the entire population.

Within a generation or so this was completely transformed. The school was the creation of the industrial age, and will pass into history with it.

YOUR definition of School 2.0 and/or Classroom 2.0...and how to help "school design" decision makers use it to inform their thinking, research, leadership, and solutions. Most will NEVER have heard the phrase...and are still beating the "School of the Future" and "21st Century School" horse over and over without really even understanding what it means besides the marketing pitch. Hoping to shift their semantic lens a bit, and also invest them in co-defining the end game as well.

Most of what is touted as 'School 2.0' or 'Classroom 2.0' is just the whole "School of the Future" and "21st Century School" thing warmed over. There is not much to it over and above marketing (and obligatory mentions of Web 2.0). Those who are a little more adept at marketing will trot out phrases today's business leaders are looking for - things like 'collaboration' and 'teamwork' - without any real understanding that they are describing a generation that 'takes orders well' and 'subsumes their own interest to the common interest' when, really, the opposite is the case.

Insofar as technology enables greater collaboration and group work, it does so only at the behest of those using the technology. You can connect people with computers but you can't make them talk. The predominate trend is not so much collaboration as a much increased sense of empowerment (and, at the more petulant ages, a corresponding sense of entitlement).

Best way to describe how 'kids' (all ages, really, but I'll use the cute version since most still default to it) are transforming as collaborators, creators, project team members, publishers, etc. I'll use Prensky language as a shot across the bow -- i.e. "What is your digital accent?" questions - but I'm looking for more nuanced language/examples from each of you. And also how we MUST respond as educators...and school designers (the entire community of stakeholders, really)...if we are to offer relevant learning environments/programs for our students/communities' future(s).

Honestly, there's no way we 'must' respond.

What is at stake is not so much our children's education - which, thanks to a wealth of freely accessible digital resources, is now more assured than ever. Rather, what is at stake is our own relevance. That if there is anything of ourselves we want to put into what children learn, how we are to go about it.

Educators need to realize that today's students are exposed to much more television, online communication, and other electronic communication, than they are to traditional classroom instruction. School, even as it is, makes up only a small percentage of their learning. It plays virtually no role in values formation, culture and self-identification, language learning and art. If school provides any learning of science, mathematics, geography and history, it is only because the students' cultural environment is almost completely bereft of those subjects. Their performance in those subjects - and especially the latter two - shows just how abject their learning has become.

Instead of bringing students to the learning, as the education system has done for about a century, we must now, if we wish to be relevant at all, bring learning to the students. This means setting students free to pursue their passions, and then being there when they need coaching, mentoring, or a safety net.

Start 'inside' a classroom, studio, lab, or micro learning space that exists TODAY. Offer a set of requests you'd make TODAY that can have a positive impact for learners and teaching guides/mentors without spending a fortune, and set up a mind-shift for bigger 'school design' investments in the future. You can imagine SmartBoards, new ways of teachers 'talking'/asking

questions, or think in really wild ways. No limits. Just start at the scale of a single space for learning...and work up/out.

We need to stop employing students as fast-food servers and sales clerks. They are capable of much better than that, and an exercise in corporate demeaning is probably not the best way to introduce them to society.

We should begin offering students full-time employment in certain fields as alternatives to their formal studies. Such a program should logically begin at the higher grades (grades 11 and 12) as well as being brought on-stream as an alternative to college and university.

Most such employment involves the creation of some sort of content or another. The ranges of possible employment are covered in my diagram:

- students could provide ultra-local news, entertainment and sports reporting
- students could provide up-to-date surveying and inventories of civic property
- students could conduct scientific field-research such as bird-counting, ecosystem sampling, pollution-measuring and the like
- students could help supervise younger children

and more - the possibilities are limited only by our imaginations.

The trick to making this attractive is to present it, not as the dog-eat-dog struggle for survival that characterizes our existing economy, but rather as a large and complex game, played partially on the computer and partially in RL, in which they play an increasingly important role.

The task we should be undertaking is not one of trying to stuff more and more knowledge into students' heads, but rather, finding more and more ways they can make meaningful contributions to society.

(I'll post this to my blog - feel free to re-use however you wish)

Moncton, Monday, April 02, 2007

The Issues in Front of Us

Responding to [Doug Johnson](#), who created this list based on one posted by [Miguel Guhlin](#).

This is one of those lists that looks so good on the surface but is really nothing more than slogans carelessly applied. I don't mean to be so critical (and would rather not be) but really would like to draw out some of the problems of this list.

The way it is: Teachers lecture - students listen

The way it will be: Teachers guide - students do

Does this mean: 'teachers tell students what to do, and students do it?' And if so, is this any advantage over lecturing?

It is always fashionable to complain about lecturing, but this does not automatically make the alternative (whatever it may be) acceptable.

In my view, the issue has nothing to do with the form and content of the educational offering. Lectures have value when used appropriately. Rather, the issue has to do with control. Lecture at me when I'm not interested, and no new information reaches my mind.

The proper approach here is to make learning available, in whatever form is desired and appropriate, to assist students as they do what they choose to do.

The way it is: Students work alone

The way it will be: Students work in groups

I hated groups. Hated them. Despised them.

Groups aren't about some better way of learning, they're about conformity, power and control.

What is so *bad* about working alone? This doesn't mean that I am isolated, without resources or support. It merely means I don't have somebody telling me what to do, taking credit for my work, and excluding me if I don't conform to their rules.

There is nothing sacrosanct about groups. Certainly no particular learning advantage is gained by forcing students into groups.

The proper approach here is to allow students to form groups if they want to (in other words, to

redefine 'cheating') but to allow them to work on their own as well.

The way it is: Subjects are departmentalized

The way it will be: Subjects are integrated

I agree. But what does that mean.

If it just means that the math problems now include examples drawn from the science curriculum, then no real advance has been made.

It is as though only two alternatives are envisioned: subjects taught apart, or subjects taught together. That's certainly what the wording of this item suggests.

The real alternative is, of course, subjects not taught. Rather, students engage in real world activities (which may include the solving of problems, but is not exclusively problem-based). As they engage in these activities, learning (that might correspond to 'subjects') becomes available to them.

The way it is: Curriculum fact centered

The way it will be: Curriculum problem centered

No.

First of all, there are domains of learning that do not involve solving problems. Art and other forms of creativity, for example. There is no known problem that is solved by the Mona Lisa. Or by the Beatles. But the world is the better for them.

More significantly, a problem centred curriculum is still a curriculum. It retains that idea that there is some One Way that will work for all students. But this is no more true of problems than it is true of facts.

Finally, students are still going to need facts, or perhaps more accurately, some way to get facts. This is merely obscured by problem centered curricula, which imposes a layer of obfuscation between students and learning.

The way it is: Teacher primary source

The way it will be: Rich resource environment

Yes. But it should be worded 'resource-rich environment' because it doesn't matter whether or not the resources themselves are rich.

The way it is: Primary print medium

The way it will be: Variety of media

If print is the primary medium, then why is it necessary to harangue against the lecture (as above) which is primarily an **oral** medium?

The **real** distinction here is between language-based and non-language-based learning resources. Multi-media seems to offer a non-linguistic alternative. But of course, a lot of multi-media is intended to deliver text in other ways. That's what podcasts are about.

More significantly, though, is the question of **why** we need to use non-linguistic resources. Normally here we might get some story about the same information being transmitted using multiple modalities.

However (and I write this while listening to some jazz guitar) what is the message delivered by genuinely non-textual modalities? How is it the same as, say, some sentence, or even some concept?

At the very least, if you get this far, you need now to be asking questions like: if knowledge and learning are not textually based, then what are they? If what we learn does not depend essentially on language, then what does it depend on?

You can't just eschew text without knowing where you're going. Which is why people say that they don't want to use print, and continue to use print.

The way it is: Success = tradition

The way it will be: Success = accountability

I am not sure where it has been the case ever that 'Success = tradition' (though I applaud the ambiguity of use with the '=' symbol, a vagueness that ought to be celebrated for looking like it means something while remaining meaningless).

'Success', more accurately, has traditionally been represented as the ability to recite, on demand, relevant facts and information. And in some cases, to solve certain types of (mathematical and linguistic) problems. As, say, defined by the structure of such measurement instruments as the SATs.

So, now, how is 'accountability' different from this?

Turns out, it isn't.

The **real** tension here is in whether the measurement of 'success' is curricular based or not curricular based. Whether you are measured against some sort of academic standard, or something else.

In my thinking, 'success' is measured by 'something else'. Where the 'something else' means leading a good life, whatever you think that to be.

By any curricular definition, by any measure of 'accountability', people like Bill Gates, Steve Jobs, and many others, are failures.

It's actually pretty easy to recognize success when we see it. In a community, we can tell whether the educational system is successful by the lower crime rates, better health of the citizens, inventiveness and creativity, and the like.

Of course, that's a lot harder to measure, and the politics of accountability won't be satisfied by these real measures. Hence the demand for pointless testing of irrelevant 'knowledge' based on 'testing'.

The way it is: Schools are insular

The way it will be: Schools are connected

We all have some sort of intuitive idea of what it means for schools to be connected. But if we push these intuitive ideas, we find that they pretty much fall apart.

They fall apart because it's not really true that schools are insular.

They are connected in myriad ways. They are governed by a school board, share in a state or province-mandated curriculum, have teachers represented by a board-wide teachers' union, have ties to the community through the PTA and other such associations, compete against each other in academic and sports competitions, and, of course, can contact each other through the mail, by telephone, online, and more.

So, in this environment, what could the author possibly mean by saying that schools should be connected?

Probably something like the pairing or twinning of classes, shared classes, co-teaching with teachers from other districts, online classes with students from multiple schools.

Things that, in general, constitute a shift from the class as something that is taught by one teacher to the class as something that is taught by multiple teachers. No?

So now I ask, well why is **that** good?

But - of course - the problem is in the statement.

It should not be 'schools are connected'.

It should be '*students* are connected. And even 'teachers are connected' (though the union already supports that, and should play an even larger role).

Of course, if the principle is that 'students should be connected' then the school is no longer so central as it was. And we have to ask, what was it about the school that made it so central in the first place?

The *school* is what keeps students separate. The less the emphasis on the school - the less, for example, that the school demands of students in a given school day, the less the school blocks and filters content from outside the school - the more students can and will connect.

Magically. Without the help of teachers or schools.

The way it is: 3 R's (Rote, Restraint, Regurgitation)

The way it will be: 5 C's (Children, Computers, Communication, Creativity, Collaboration)

What teachers and what schools are genuinely willing to give up on rote, restraint and regurgitation?

What is 'curriculum' other than rote?

What is a 'school' other than restraint?

What is 'accountability' other than regurgitation?

I think that when most people read these, they will nod in agreement but they will be thinking of the 5 Cs *in addition* to the 3 Rs.

Look at the last four: 'Computers, Communication, Creativity, Collaboration'.

Are children going to be *required* to use computers, to communicate with each other, to be creative, to work in groups?

Will the school be the place where they go to use computers and to talk to other students? Will the school be the place where they find the materials and tools to be creative? Will the school be the place where they are required to work with other students?

Will they need to show what they know about computers, perhaps by answering questions on a test? Will they prove they have communicated with others by answering questions about them? Will 'creativity' be defined as something safe?

If it was *about* children, then they would each get their *own* computer, with which they could

communicate with others - children and otherwise - as they *choose*, with which they could create software or art or literature as they desired, where they could work collaboratively, collectively, or without others at all.

In summary...

My reading of this list is that although it looks like it is on the forefront of advocacy for change, it disguises a sentiment which is at heart fundamentally conservative. It offers the illusion of change, without actually promoting change.

I don't know whether this was the intention of the author, and I won't speculate on motives. It feels to me like a well-intentioned post that was simply unable to move outside of a comfort one.

But the reason why I felt that it was important to comment on this is that it is characteristic of a lot of recent writing that I have seen in the edublogosphere that walks and talks as though it is at the forefront of something new but is in reality an effort at retrenchment, an effort to protect one's own turf while embracing the change swirling around it.

The recent 'School 2.0' movement is a good example. By locking into the concept of 'school' the proponents, while looking for all the world like they are embracing change, are in fact freezing the state of education into an archaic past, where the school is the centre and where everything else - including the students - revolve around that central concept.

The idea of 'school 2.0' by definition eliminates as out of scope any concept that reduces or eliminates the importance of the school (and by extension, the elements that constitute a school, such as classes and curricula, teachers and lessons).

Given that the shift in focus from authority (such as schools) to empowerment (such as for students) is at the very core of the whole concept of '2.0' the idea of 'school 2.0' is inherently self-contradictory. It stands for the very *opposite* of what its public posture presents.

That's why I posted this response. The future is much more difficult to grasp than a mere set of slogans. Fundamental values are shifting under our feet. Pretending it's something superficial, as represented by this list, won't change that. It is important to have an accurate representation of the issues, so people can genuinely understand what they are facing.

The Form of Informal

Part One

Clark Aldrich, as cited by [Tony Karrer](#), writes, "Can one criticize formal learning models in a book? Isn't a book the epitome of what one is suggesting is the wrong model?"

This is a pretty equivocal discussion. The words 'formal' and 'informal' are used in distinct senses.

In *one* sense, something is 'formal' or 'informal' according to whether it forms some part of a recognition infrastructure - the system of classes, programs and institutions that constitute our certificate and degree granting structure. Thus, when one has 'formal' qualifications, it is in reference to this infrastructure.

In *a second* sense, something is 'formal' if it is derived or represented in or as an abstract structure. In this way, systems or methods of doing things are 'formal', as are structures, 'forms' and other abstract entities. Thus, when we study 'formal' logic, we study abstract principles of reason, mechanisms for successful argumentation, truth tables and truth preservation, and the like.

Clearly there is a very large difference between having a method of doing something and having some educational infrastructure regulate and recognize the doing of something.

What complicates the matter is that the criticism - 'how can somebody write a book about informal learning' misses the mark in *both* senses. Because we have to ask - in what sense is a book formal?

The most that we could do is to try to adduce some sort of *third* sense of formal - something based on the length, or perhaps the book being organized into sections and chapters, perhaps, or even something based on the book being written in a *language*, which is at heart sequential. But neither segmentation nor sequence is sufficient to make something 'formal'.

I have over time lost patience with people who criticize informal learning on the ground that, if it is informal, it has no structure whatsoever. Articles like '[Informal Learning is Too Important to Leave to Chance](#),' I might add, fall into that category.

The mere fact that the learner is not being directed by some teacher or educational institution does *not* entail that there can and must be no structure whatsoever. *nobody* equates 'informal learning' with 'structure-free learning'.

The test for this is simple: ask any person who has learned something informally - whether it is how to program computers or how to tune an engine - how they learned. If they have *any* answer, then there was some structure to the learning. If a person says, "well I followed the discussion lists and then I

looked up what I needed on Google," *there's a structure*. If he says, "Well I watched my uncle Fred tune a few engines, then I tried it while he watched," *there's a structure*.

What makes informal learning different from formal learning is *not* that it is formless, but rather, it that it is conducted outside the domain of the formal education infrastructure, with the associated and not trivial implication that it is managed by the learner, and not the professor or institution.

That's why a statement like 'too important to be left to chance' is so misleading. It implies that there is no reason why a person (whether an employee or a student) might choose this or that informal learning method. It implies that nothing can be done to support this person, to suggest some structures or mechanisms, to improve their likely outcome. It assumes that, unless we *control* this person, the outcome is 'by chance'.

But that's ridiculous. It is one thing to say, "I'm personally much more confident if I have a set of performance objectives that I can use to derive learning objectives and skill development opportunities around. I want to put structure in place that guides the learner along the way." It is quite another to say that the learner (a) *must* attend your class, and (b) *must* adhere to your learning outcomes and learning methodologies.

When I have characterized the distinction between formal and informal learning, I have done it this way: by saying, if you can walk out of the room, or change the shape of the discussion, or skip an activity, without (academic or other) sanction or penalty, then it's informal learning.

People can publish a book and still support informal learning. After all, picking up a book, reading it on your own time, and consulting it as necessary, is the *epitome* of informal learning. No classes, no structure, no tests, no grades, no degrees. Reading a book written by James Joyce is not somehow 'more informal' than reading a book by Jay Cross.

From my perspective, I see more of a tension when people refer to books as being authoritative on the internet.

For example, [Dave Warlick did something](#) to one of my comments recently to which I took objection, and when he asked me to justify my complaint he asked, "What book is it written in?" As though *this* would be authoritative.

When I replied, *the internet*, it was taken by some readers to be flippant or even rude. Why would I not provide a citation? But what citation would be better than the internet itself when addressing a question of what is or is not done on the internet? How would what is written in a book actually outweigh what people do online?

In certain circles, publishing a book still confers some authority upon its author, as though cozying up to

some publisher (and signing away your copyright) somehow grants you some sort of status. But I don't agree.

We are at the cusp of a division between the age when people gained authority through what they printed on paper and when people gain authority through what they print online. It will not be a quick transition, because it is easy to recognize when somebody publishes a book, and much harder to recognize when they publish a quality website.

But as time goes by, and when we realize that Paris Hilton has published one too many books for our liking, or that some online writer has published one too few, this transition, from being anointed an 'expert' by the publishing industry, to being anointed an 'expert' by a community (of practice), will become easier.

Part Two

Responding to Tony Karrer, [More on the Form of Informal](#)

What do you think the meaning of the word 'dictating' is?

"How do you aim at performance objectives, provide appropriate support, structure and form without dictating to some level?"

I can 'aim at performance objectives' in a variety of ways, some dictatorial, others not:

- I can suggest that the company (or society) needs more widget experts, for example, or
- I can require that all employees undergo widget training

Similarly, given that for whatever reason a person has decided to learn about widgets:

- I can make available some resources on widgets
- I can provide some just-in-time learning support for people using widgets
- I can give salary bonuses for people who demonstrate they have passed a widget course
- I can assign an online course that all staff must complete
- I can hold mandatory widget training sessions

"No one says its structure free, but how much structure is allowed before it becomes too much to be informal? There's some kind of spectrum here with all sorts of shades."

The presumption here is that 'structure' is some sort of monodimensional property, akin to (say) 'complexity'. But it is not. Something can be very structured in one dimension and very unstructured in another.

For example: viewed a certain way, a forest is very unstructured. Nobody is organizing anything. Trees grow whether they will, deer and other woodland creatures wander about with no guidance whatsoever. But viewed another way, a forest is very structured. It is an instance of a complex ecosystem. Patterns are repeated in the shapes of the trees, the shapes of the leaves.

The point here is: 'structure' is not a 'spectrum' per se. Rather, some types of structure tend toward 'informal', while other types of structure tend toward 'formal'.

I'll come back to that.

"What's interesting is that the moment you begin to understand the form that informal learning took and provide support for that kind of learning in the future you start down the path of dictating solutions."

Again, this simply doesn't follow, unless you have a very unorthodox understanding of the meaning of the word 'dictating'.

The term 'dictating', as normally understood, entails some sense of enforcement. To 'dictate' that something will be done is to require it to be done, to impose an obligation on someone that it be done, and hence, concurrent with the concept of enforcement, is related to the concepts of power and control.

But not all structure is the result of, nor requires, power and control. This is the fallacy I've been trying to get at.

Let's call this the 'teleological fallacy' - the presumption that, where there is order, there must be an organizer, someone who creates and manages, through some process of authority, this order. A 'God' of the training room 'design', if you will.

"Providing job aids was one example that Michael McGinnis cited. The first time a person learned how to do that task/job it was likely through someone showing them how to do it. Once they put it in a job aid - it feels more formal on the spectrum. So maybe it's not a paradox, maybe it's a spectrum."

The mere transfer of some 'showing' into some 'job aid' does not make something more formal. Some very formal instruction - such as, say, by a drill sergeant - can involve 'showing'. And some very informal instruction - such as a 'how-to' manual - can be a job aid.

So why does it 'feel more formal' (in, at least, some cases)? Not because it has been given structure, but because the structure has *changed*. Some types of structure are formal, others are informal.

"But, it feels a lot like what Artificial Intelligence faces - informal is a bit mysterious and putting structure to it makes it feel more formal and much less mysterious."

Perhaps by 'structure' here you mean 'rule-based'?

Again - informal does *not* mean 'having no structure'. Rather, 'informal' means having a *different* kind of structure. One that is, among other things:

- not dictatorial
- not organized or managed by an organizer
- not rule-based

"Of course, I personally am not that concerned with the definition of the term nor really even the paradox as I am in understanding the next level of informal learning: what kinds of guidance, what kinds of support, how can be provided, that ultimately lead us to accomplishing our performance objectives?"

Right. And neither am I, because it is in attempting to define the term that we get right into the sort of fruitless and endless debates that formal semanticism so often entail.

But...

It is still relevant to ask, what sort of structure tends to characterize the informal, and what sort of structure tends to characterize the formal? Structure that is not dictatorial, not organized, and not rule-based, to be sure. But what sort of *structures* look like that?

In a slogan: networks, as opposed to hierarchies.

I have [elsewhere characterized](#) the properties of systems that tend to be networks, as opposed to hierarchies, via their having following structural properties:

- decentralized.
- distributed.
- disintermediated.
- disaggregated.
- dis-integrated.
- democratic.
- dynamic.

Sorry about the cutesy list of Ds.

The sixth condition, 'democratic', constitutes what I call the semantic condition, and is constituted of four major elements:

- autonomy

- diversity
- openness
- connectedness

These (and conditions similar to these - I am not wedded to any *particular* characterization of these conditions) constitute a metric that distinguishes formal from informal. Things that have more of these properties are more informal, and vice versa (this will not be a strictly linear progression, which is why it is not merely a spectrum; take away even one characteristic - 'autonomy', say, and you instantly convert a system that was very informal to something very formal. The way a party becomes something very different entirely the moment someone says you have to attend.

But the main point here is that, *within* these parameters, there is a great deal of room (and *expectation*) of organization. The concept of 'distributed', for example, by its very nature assumes the existence of communications protocols between the parts. And the same for the other conditions.

When we look at this sort of model, we can see how unsatisfactory discussion akin to "Providing job aids was one example..." can be. We ask, immediately, *what sort of job aid?* Is it something employees were required to use (like, say, project template forms)? Is it something that needs to be scheduled? Managed? Produced from a central depot by a team of experts? All of these effect whether it can be characterized as 'formal' or 'informal'.

Different types of technology, different types of learning materials and different types of pedagogical practice may be characterized as more or less informal according to these criteria. Not based on the simplistic assessment of whether they are 'structured', but rather, on the much more insightful determination of *how* they are structured.

Uniqueness and Conformity

'Short Input' to [International Monitoring Conference](#), September 30, 2010

It is interesting that Fritz Bohle immediately characterized the dilemma between 'stability and flexibilization' as 'management of uncertainty', and focused on the idea of science having as an enterprise the reduction of uncertainty. The reality is that, as he said, the uncertainties resist elimination. I will consider why this is.

Johannes Sauer writes, "In order to protect and extend Germany's capacity for innovation and competitiveness, the extension and organisation of learning cultures are of major significance within the process of transforming the industrial society into a knowledge society."

Unstated in this assertion, and in assertions like it, is that the nature of 'knowledge' itself is changing as society changes. So we should not interpret the phrase 'knowledge society' from our comfortable definitions of knowledge.

The fourth dilemma outlined in the International Monitoring discussion paper "describes the demand of individuals, organizations, networks and societies for safety of current and planability of future processes." It is possible that the depiction of society as a 'knowledge society' offers for some this safety and stability.

If knowledge is derived to any significant degree from experience, however, then as new technologies, social structures, and innovation are generating an increasing number of novel and unexpected experiences, the continuous state of knowledge itself is one of change, as what we know adapts to what we have experienced.

Consider the concept of 'knowledge processing', from the opening keynote – this treats knowledge as though it is some kind of resource or raw material, like iron or coal, that will be transferred, reformed, processed. This is a traditionalist perspective of knowledge, which is no longer appropriate today.

Where there is structural complexity and process complexity, there is also epistemic complexity. Fully realized, a state of total knowledge is indistinguishable from total complexity, or chaos. That which is 'static' or even 'dynamic' is nothing more than an interpretation, a pattern recognized and indeed imposed on the world.

They say knowledge is power. But in fact, power is knowledge. The only order in the world is that which is imposed, by those in power. In order to understand the changing nature and role of knowledge, we need to understand the changing nature of power. As we have evolved historically, from the power of the monarchy, to the power of the corporation, to something (which lies still in the future) a more

decentralized power, so also knowledge evolves from a single, centralist concept, to the pluralism of corporatism, to the chaos of individualism.

In a chaotic environment, knowledge is nothing more than pattern recognition.

The challenge of commonality where there is no static underlying essence to unite us.

Mike Bullard, Canadian comedian, on the secret to stand-up comedy

- first, you establish something in common with the audience
- then you bring them around to your point of view
- then you get them to laugh at themselves

The point is – the joke doesn't first exist in the teller, and then appear in the listener. The joke exists entirely in the listener. The teller possesses only the mechanics of joke production, but not the actual humour. The comedian laughing at the audience is completely different from the audience laughing at themselves.

Knowledge works the same way.

The proposition from the keynote was, only companies that are unique are competitive. But knowledge is found in the recipient, not the company, which contains only the mechanism for knowledge production. If the company must be unique to achieve value, it must at the same time find a point of commonality in order to realize that value. That point of intersection is the critical point of innovation.

Management-union, social partnership, shared values – are artifacts of the older perspective.

Thinking of industry – manufacturing – the factory vs the artisan vs the individual... the industrial age created tools that could be wielded only by masses of individuals working in concert. But the post-industrial age has resized tools again. "The value of a tool in a man's hand has to be re-valuated."

The 'tools' of knowledge are the same.

The success of, say, electricity was based not on uniqueness but on commonality. The current that was sent was accessible, via a point of interaction, to every person in the world (the challenge to consumer power lies in this same point). But the semantics of electricity – the use to which it was put – was unique and determined by the individual.

Moncton, September 30, 2010

New Technology Supporting Informal Learning

Abstract

We often talk about games, simulations and other events in learning, but these technologies support only episodic learning. Equally important are those technologies that provide a context for these learning episodes, an environment where students and interact and converse among themselves. This paper described experimentation in the development of distributed online courses and in software - particularly, the personal learning environment - that support the formation of connections between the far-flung pieces of such courses. This work, in turn, is suggesting and supporting the model of learning described in the first section, that of a course network supporting and informing an ever-shifting set of course episodes. This in turn suggests a pedagogy of participation rather than retention, and even suggests distributed and locally-based forms of evaluation and assessment. Future developments will focus on realizing these concepts as software or at least software prototypes. The intent of such systems is to to facilitate the conversation and interaction around episodic learning events in a distributed environment, transforming them from elements in a linear flow-based design to free-floating objects in an environment.

Context

Online learning today is beginning to be dominated by developments in games, simulations and related technologies. (Akili, 2007) And there is no doubt that this is a positive development for the field. Such applications are almost unambiguously beneficial for the student. In addition to providing an engaging and immersive environment for student learning, substantially improving motivation and interaction with the learning material, games and simulations are able to support learning in complex environments, offering a subtlety simple instruction-based or lecture-based learning cannot offer. (Squire, 2005)

What most characterizes games and simulations is that they are not merely forms of instruction, they are environments, into which students must immerse themselves in order to participate. (Foreman, 2004) As environments, they model complex relationships between variables, resulting in an experience that is unpredictable and unique each time played. (Aldrich, 2005) It is this feature, and not simply the action and the graphics, that motivates learners and draws them in. With the addition of interaction with other participants, as seen in massive online gaming environments, the experience can be almost addictive. (Ng & Wiemer-Hastings, 2005)

That said, these environments, by their very nature, require intense preparation on the part of the designer. In addition to graphics and game play, there is the content of 'storyline' to consider. (Rollings

& Adams, 2003) In the case of learning environments, the planned learning objectives or outcomes need to be programmed into the game design, involving a further layer of preparation. Consequently, games and simulations fall into a category similar to lectures and presentations in that they involve statically designed learning objectives and strategies. (Amory & Seagram, 2003)

As a consequence, interaction with such environments, even the most immersive and addictive game or simulation, must have a start point and an end point. Such systems are by their very nature episodic. Because they must be designed in advance, they are inherently static, at least at the level of overall design and framework. Consequently, they represent a separation between the learner's in-environment experience and his or her wider life of leisure and work. Consequently, in order to place games, simulations and other episodic learning events such as classes and lectures, into the context of the student's wider life, a wider frame of reference is necessary. In this wider frame we would expect to find a wider environment of conversation and interaction with friends and associates. This wider frame situates, and plays a significant role in the selection, of episodic learning events.

Why is this necessary? In short, it is simply impossible for simulation, game and learning designers to design unerringly for the learning needs of the student. First, and significantly, we often do not know what it is we want to teach the student. (Caine & Caine, 1997) Today's environment is variable, which means situations - and hence, fact - change fluidly. One day Pluto is a planet, the next day it is not. One day Czechoslovakia is a country, the next day it is not. One day capitalism is the unassailable foundation for our economic system, the next day, following a market collapse, it is not. Moreover, today's environment is complex. The relations between variables cannot be described or even predicted. An understanding of such things as the financial system or global climate change requires a subtle and ever-changing perspective on the discipline.

Second, learners themselves are changing. There has been much discussion in recent years about the rise of the 'digital native' or of the 'net generation'. It has even been suggested that our interactions with modern communication technologies change the way we think. Even if we reject such descriptions as students as overly broad and inaccurate generalizations - and there is good reason for doing so - it is nonetheless the case that the needs, capabilities and interests of the target audience is rapidly shifting and changing. As much as it is tempting to say that human nature is unchanging, it appears nonetheless the same that human experience is endlessly varied, resulting in any number of approaches to media in general and learning in particular. A child raised on text alone will think and learn differently from a child raised on cartoons or a child raised on Facebook. (Peters, 2006)

In part, our best response to the variability and complexity of the subject matter along with the changing nature of the learner is to design systems that are decentralized, to push learning decisions down the hierarchy or out to the edges of the network. (Wiley & Edwards, 2002) This logic, which is characterizing not only new learning but also new approaches to business and management (Malone, 2004), is based on the idea that those who are closest to the situation are in the best position to make decisions about it. In the military, this means that company commanders, and not generals, must make

tactical decisions. In business, this means that salespeople and customer service representatives must determine marketing policy. And in learning, this means students must be empowered to make their own learning decisions. This is the basis for the models and strategies that characterize what has come to be called informal learning. (Cross, 2006)

But there is in addition a second and critical aspect to this wider environment of conversation and interaction. It is not merely to create a network into which to situate episodic learning, but rather, to create a network that learns and thus adapts and reshapes itself based on those conversations and interactions. (Downes, 2007) We need to consider learners not only as the subjects of learning, entities to whom we deliver learning content, but also the sources of learning, functioning as the perceptual input for the wider network. (Marotzki & Specht) The things we say, the things we choose to read or view, the things we link to, the people we send messages to - all of these constitute input to the learning network, causing it to reform, causing it to present, say, one learning episode rather than another, one game rather than another, one simulation rather than another. And, moreover, our reflections and commentary on various games, simulations and learning events constitute feedback for those systems, modifying them internally as well, either directly, or through a series of design iterations, just as we see in (for example) agile programming. (McCall, 2005)

Learning networks capture an essential element in learning today, the simple fact that we don't know what we want to teach. Indeed, it is often suggested that the best we can manage is to teach students how to learn, and to encourage them to manage their own learning thereafter. But even this principle is subject to changing affordances of technology and changing capabilities of students; how we learn itself is something that changes, and cannot be precisely taught. The way musicians learn, for example, changes as they grow from novice to expert. (Münste, Altenmüller, & Jäncke, 2002) For this reason, we need to see the educational system itself as adaptive rather than merely prescriptive.

We are seeing the development of specific instances of this approach to learning today. For example, a learning system called Company Command, designed by officers in the U.S. military starting in 2000, is essentially a learning network composed of company commanders. (Dixon, Allen, Burgess, Kilner, & Schweitzer, 2005) While most traditional military training is conducted from trainer to learner, Company Command starts with a significantly different proposition: that knowledge exists in the minds of the members or participants, and this knowledge is derived from their direct (and recent) experience in the field.

In addition, the need for content and support emerges from conversations among the participants. These interactions are able to reveal not only what company commanders know, but also what they don't know (and need to know). The interaction, in other words, meets and addresses an objection often put of self-directed learners, that they don't know what they need to know. (Clayson, 2005) While this may be true, through participation and interaction in this wider environment they are able to identify these needs (as expectations, for example), and hence to select and conduct appropriate learning episodes. (Chickering & Ehrmann, 1996)

The model of Company Command is one that has been repeated many times on the internet. Company Command itself began as one of thousands of Drupal applications. The core purpose of Drupal is to facilitate the creation and management of communities online (Drupal, 2009), including communities of practice of the sort that typify such cases as Company Command. Drupal, an open source content management system, enables the creation of individual accounts, the creation of discussion posts and pages and other content, and the sharing of this content online with other community members. Many other systems provide similar functionality (CMSWire, 2009), and in learning, the learning management system (LMS) provides the wider conversational context for in-person or online learning episodes. (Garrote, 2007)

More recently, social networking technologies have come to be applied to content and learning management systems. (Dignan, 2008) The core of a social networking technology is the capacity to create links between members in a community - to create, in other words, social networks. These links are usually created explicitly, through the declaration of each of the as 'Friends'. Often, the creation of links is associated with the creation of content, as in content management systems. The last few years have seen the development of social network services online such as Facebook, Friendster, LinkedIn, and MySpace to name a few) as well as a service for creating social networks, Ning. (boyd & Ellison, 2007)

Social networks represent a gradual decentralization of content and contact online. Content management systems (and before them, email lists and usenet groups) organized people and content by hierarchy, by topic and content thread. In social networks, such associations are created by the users themselves. Topics, for example, are not assigned centrally, but are instead created by individuals 'tagging' certain content with terms or categories they choose themselves. (Barsky & Purdon, 2006) Each person's social network on a social networking site, moreover, is unique; there is no definitive grouping of people, only a clustering of people with more or less similar interests.

Software to create social networks is the logical successor to content management systems such as Drupal, and in the field of learning, the most prominent such system is Elgg. In addition to supporting content creation and the creation of networks of friends among members, Elgg allowed people to import content from remote sites and to syndicate this content through the use of RSS feeds. (Elgg, 2009) Systems such as Elgg therefore combine the functions of content management and social networking.

In summary, then: we often talk about games, simulations and other events in learning, but these technologies support only episodic learning. Equally important are those technologies that provide a context for these learning episodes, an environment where students and interact and converse among themselves. Historically, this role has been played by the content management system, while more recently systems supporting social networks have also come into prominence. Such systems represent a partial decentralization of the management of learning, pushing some decisions (such as association

with other learners or clustering of material into categories) from central decision-makers to the learners themselves.

Current

At the turn of the century, the dominant model of online community proposed by pundits was one that could be characterized as a 'destination resort'. It would be, suggested writers like Hegel and Armstrong, a complete interest-based community revolving around travel, homeowners or personal finance. (Ill & Armstrong, 1997) Members would be attracted through marketing and content, would be encouraged to contribute content, would become loyal to the community through friendships and interaction, and would be monetized through value-added services and sales. (p. 59) What developed instead resembles barrios more than resorts: a complex interweaving of online services, sites, interactions and applications developed ad hoc rather than at the behest of some community planner.

Also at the turn of the century, it was thought that online services would interact with each other in an organized and managed way - they would, in other words, be "choreographed". The classic example involved a travel service where one central provider - the agent - would send requests via web services to hotels, car rental agencies, airlines and even caterers in order to seamlessly manage the experience. (Yendluri, 2003) What developed instead were random, individualized and often ad hoc assemblages known as 'mash-ups', these based on lightweight communications technologies such as REST, AJAX and APIs. (Rushgrove, 2007)

Far from being neat and organized, the internet has become complex. Far from settling into one web community, users jump from service to service, creating (and discarding) new identities as needed. A typical web user may have multiple 'home pages' - their personal blog, their photo page on Flickr or Picassa, their Google Reader account, shared documents through Zoho, their video page on YouTube, their Twitter account, their profiles, on Facebook, MySpace and LinkedIn, their Wikipedia login, their email accounts, and (often least) their university LMS login. While real friendships and communities develop through this *mélange*, loyalty to online sites and services is limited and fleeting. (O'Brien, 2007)

It is in this environment that the University of Manitoba's Connectivism course was developed in the fall of 2008. (Siemens & Downes, 2008) Developed by George Siemens and myself, at least one intent of the Connectivism course was to facilitate the transition from a neat, constrained and centralized learning management system to a distributed environment in which students and instructors employ multiple online services and applications. (Downes, Tools, 2008) Consequently, the Connectivism course must be viewed as one of the first distributed courses to be created - not simply distributed in terms of time or place, but distributed in terms of website or application.

Much has been or will be written about the course elsewhere, but it is sufficient and relevant to say that roughly 2200 people signed up to participate in the course. While the course was offered as a tuition-based for-credit course, in order to foster the network dynamics we also chose to open the course to all

participants. (Siemens, On Finding Inspiration, 2008) In this we were following not so much the model offered by OpenCourseWare and others, which made learning materials freely accessible online, but rather David Wiley and Alec Couros in offering actual course instruction available online. We wanted students not merely to consume learning materials, but rather, in the manner of the wider environment discussed above, to contribute to the learning through conversation and interaction. (Siemens, Narratives of Coherence, 2008)

To this the students contributed in droves. The central course aggregator listed 170 separate weblogs or similar RSS feeds contributed by students, each of whom used their own blog or website to participate in discussion. (Downes, Feeds, 2008) Additionally, thousands of comments were contributed to the central Moodle forum, three separate areas in Second Life were contributed, Google Groups were created, a Ning was created, and more. In fact, student contributions to the course continue to this day even though the course was completed in December, 2008.

As no viable mechanism for connecting the disparate and distributed course contributions exists, we adapted my newsletter software, gRSShopper, for the course. (Downes, gRSShopper, 2008) This software was developed out of a need for a personal online web space to do more than was possible in Drupal (in fact, I document my trial with Drupal in a series of posts). gRSShopper is available as free and open source software for public download.

gRSShopper is a prototype personal learning Environment (PLE). The PLE is a concept developed in 2005 in conversations with and among members of JISC CETIS and their friends and associates. (Wikipedia, 2009) The idea of the personal learning environment is that it performs many of the functions of a content management system and of a social network system but from the perspective of the individual rather than the community or the institution. (Attwell, 2006) Hence, the PLE may be understood as the intersection of the multiple home pages employed by any given individual. In the first instance, the PLE is a concept, rather than an application - it is the idea that a person's web presence can be distributed. (Attwell, Graham Attwell: "Knowledge is best shared and developed through communities of practice", 2007) And in applications such as gRSShopper it is the instantiation of that idea in a personal application.

In the context of the Connectivism course, the prototype PLE proved to be an admirable teaching application as well. While online course content was provided to students using the more traditional mechanisms of a Moodle learning management system (to host discussions and conversations) and a course Wiki (to host the course outline and links to learning resources) the use of gRSShopper allowed us to send, by email and RSS, a daily newsletter to students' own email or RSS readers. (Downes, The Daily, 2008) In this way, we were linking course material out, to students. Student registrations to the email newsletter remained at a constant level of 1870 subscriptions through the full duration of the course.

Additionally, however, gRSShopper has a built-in RSS aggregator. Hence, we collected the feeds from the 170 separate blogs and websites created by participants and stored the student contributions in the

gRSShopper database. This allowed us to filter content by tags and to include this content into the daily course newsletter mailouts. We selected and distributed material containing the 'CCK08' tag (thus not diluting the newsletter with unrelated material) . We also created (manually) links to online events such as Elluminate and Skype discussions, sessions in Second Life, occasional videos, diagrams and other resources. (Siemens & Downes, Wiki, 2008)

Because there were so many people contributing to the course, and because the content of the course actually shifted and varied according to participation and input into the course, it was necessary to emphasize to students that their role in the course was not to attempt to assimilate all course content. This was neither possible nor desirable. Rather, students were told that their role was to select and sample course content, pursuing areas of interest, reading related material from both within and outside the course, and then to contribute their unique perspective based on this reading. (Siemens, *Where does the learning occur??*, 2008) Students would be evaluated, we said, not based on their retention of course material, but rather, on the basis of their contribution to the discussion, their interaction and sometimes collaboration with other participants, and their evolving capacity to work within a network to produce new knowledge in the field.

Indeed, the world-wide and distributed nature of the course suggested an alternative method of evaluation entirely, one that separated course content from evaluation. Students from other countries and other institutions could register into the course as students and participate in the course, and use their work in the course as material submitted for evaluation in their own home institution. To that end, we made all assignments and evaluation metrics available to all participants, to share with their home institutions. At least one person requested, and was evaluated, in this fashion.

To summarize: we are currently seeing experimentation in the development of distributed online courses and in software - particularly, the personal learning environment - that support the formation of connections between the far-flung pieces of such courses. This work, in turn, is suggesting and supporting the model of learning described in the first section, that of a course network supporting and informing an ever-shifting set of course episodes. This in turn suggests a pedagogy of participation rather than retention, and even suggests distributed and locally-based forms of evaluation and assessment.

Future

Future developments around the concept of the conversational and interactive environment begin with preparations for a second offering of the Connectivism course in 2009. In particular, work to date has revolved around the idea of simplifying the production of course newsletters. Even with content aggregation, these were taking the author (me) about an hour every day, as course content (such as planned online events, readings, etc.) needed to be input into the newsletter body. To this end, a system to develop serialized feeds (Hirst, 2009) was created, in order to automate the distribution of scheduled course content. (Downes, *Serialized Feeds*, 2009)

The idea of a serialized feed is to create elements and to store them into a data base. Each element of course content corresponds roughly to a blog post - that is, it is dated, has its own page, and may link to external resources or services. Each post is then given an offset value which stipulates, in number of days, how long after the onset of the course a material should be delivered. When a course is initiated (by the registration of students into the course) the timer is started. The system automatically delivers a newsletter each day. Student contributions, filtered for the CCK08 tag as before, are harvested and inserted into the newsletter. Then any content from the database with an offset matching the current course day is also added. The completed newsletters are distributed by email or RSS.

Serialized feeds are one aspect of a more general development program being undertaken around the idea of the personal learning environment. As noted above, the PLE merges the function of the content management system with the social network service, and does so from the perspective of individual students. Hence the PLE could be depicted as being a node at the centre of a network, connected (using standards such as RSS) to content and other services across the web. Examples of such services in Scott Wilson's paradigm document include Flickr, 43things, LiveJournal (a blogging service), an academic publisher, and more. (Wilson, 2005)

In the PLE project being undertaken by the National Research Council, the functionality of the PLE is depicted in four major stages: to aggregate, that is, to collect content from the individual's and other online content service providers, where aggregation includes elements of recommendation, data mining and automated metadata extraction ; to remix, or to organize content from several different sources in different ways, including through automated clustering; to repurpose, or edit, localize, or otherwise modify or create new content; and to feed forward, or send the content to subscribers and other web services, either via RSS syndication, email, Twitter, or other relevant services. (Downes, Theory of Learning Networks, 2004)

When viewed from the perspective of a collection of students taking a class (such as, for example, the Connectivism class), what is created using the PLE is not a recreation of the capabilities of the learning management system, but rather, a learning network. Though through the use of serialized feeds and similar mechanisms educational institutions and instructors can feed content, services and resources into the network, actual structures of the network, along with many of the resources exchanged in the network, are created by the students themselves. These structures are reflective of the students' interactions with each other and with the wider community (surrounding a particular content domain) and hence the structure of the network varies as student experience varies. (Levin, 1995) A network of PLEs is a learning network.

As suggested above, the pedagogy of such a network is strikingly different from what we might find in a content-based (content-management based) course. The order and structure of the content is dissolved; while episodic content (such as books, simulations or lectures) maintain an internal logic and structure, the linear or hierarchal structure that previously defined courses is entirely absent. This does not mean

that the relation between course, participants and content is completely unstructured, only that the nature of the structure has changed. It makes more sense to think of learning episodes as objects that inhabit the wider environment, the conversational and interactive environment that constitutes the course. (Downes, *Learning Objects: Resources for distance education worldwide*, 2001) The entities in such an environment - individual students, as instantiated as PLEs, along with course episodes, as instantiated as readings or services or games - interact with each other much in the way physical objects interact in an environment: not according to any central plan, but via the internal motives and affordances of each object.

The computational structure for such a model exists in the field of object-oriented programming, where computer programs consist not (simply) as lists of instructions to be followed, one after another, in a linear or branching manner, but rather in an open-ended consideration of the properties and states of each object. (Meyer, 1997) On the internet, the best example of such systems are the Multi-User Dungeons (MUDs) that were developed in the late 80s and early 90s, where objects could have properties and methods (functions) that influence those properties, and could send messages to other objects invoking methods in those objects as well. (Cuciz, 2004)

This model informs the design of learning experiences as well. In traditional learning design, as instantiated by (say) EML or IMS-LD, learning design consists (essentially) of a flow of learning experiences, choreographed (or directed) just as web services manage access to different travel resources, where there is a script, planned learning outcomes, and localized, content-based evaluation or assessment. (IMS Global Learning Consortium, Inc., 2003) The traditional model suits a world of content management systems where the delivery of learning episodes as well as the content of those episodes can be anticipated and planned in advance. (Vries, Tattersall, & Koper, 2005)

An alternative to traditional learning design, state-based learning design, will be proposed, such that the presentation (and even the content) of a learning episode will vary depending on the relative states of the objects in a given environment - that is to say, the set of values and methods present in the set of objects in a given interactive space (defined by linkages between individuals and content). In state-based learning design, learning resources are not arranged as sentences in a paragraph or chapters in a book, but rather, are used as a form and means of communication, more in the manner of words in a vocabulary. Their use is suggested by content rather than mandated by learning imperatives.

Such a change is essentially a migration of IMS Learning Design into a Rule-Based design more characteristic of object oriented systems. (Martínez-Ortiz, Sierra, & Fernández-Manjón, 2009) Rules may resemble simple functions, such as "show an activity" or "hide an activity" or may represent more complex interactions. While traditional IMS learning designs could be mapped into such a system, the reverse would not always be true, as such a system would be capable of more open-ended interactions not describable in a flow-based format. The interaction between user and content would resemble the dynamics and interplay of a simulation or a game. Indeed, these latter learning episodes would take their place as objects within this larger learning environment.

To summarize: future developments have focused on realizing the concepts displayed in the Connectivism course as software or at least software prototypes. The intent of such systems is to realize the objective of the design of the Connectivism course, to facilitate the conversation and interaction around episodic learning events in a distributed environment. This realization is essentially that or re-orienting learning objects, transforming them from elements in a linear flow-based design, such as described in IMS-LD, to free-floating objects in an environment, activated by the triggering of rules in an object oriented environment.

Akili, G. K. (2007). Games and Simulations: A New Approach in Education? In D. Gibson, C. Aldrich, & M. Prensky, Games and simulations in online learning: research and development frameworks (pp. 1-21). Hershey, PA: Idea Group, Inc.

Aldrich, C. (2005). Introduction to Systems Thinking. In C. Aldrich, Learning by Doing: A Comprehensive Guide to Simulations, Computer Games, and Pedagogy in E-learning and Other Educational Experiences (pp. 18-32). San Francisco: John Wiley and Sons.

Amory, A., & Seagram, R. (2003). Educational Game Models: Conceptualization and Evaluation. South African Journal of Higher Education , 17 (2), 206-217.

Attwell, G. (2007, April 16). Graham Attwell: "Knowledge is best shared and developed through communities of practice". Retrieved from eLearning Europa:
http://www.elearningeuropa.info/directory/index.php?page=doc&doc_id=9098&doclng=6

Attwell, G. (2006, June 1). Personal Learning Environments. Retrieved from Wales Wide Web:
http://www.knownet.com/writing/weblogs/Graham_Attwell/entries/6521819364

Barsky, E., & Purdon, M. (2006). Introducing Web 2.0: social networking and social bookmarking for health librarians. JCHLA / JABSC , 27, 65-67.

boyd, d. m., & Ellison, N. B. (2007). Social Network Sites: Definition, History, and Scholarship . Journal of Computer-Mediated Communication , 13 (1), 11.

Caine, R. N., & Caine, G. (1997). Understanding Why Education Must Change. Retrieved from New Horizons for Learning: http://www.newhorizons.org/trans/caine_change.htm

Chickering, A. W., & Ehrmann, S. C. (1996). Implementing the Seven Principles: Technology as Lever. AAHE Bulletin , 3-6.

Clayson, D. E. (2005). Performance Overconfidence: Metacognitive Effects or Misplaced Student Expectations? Journal of Marketing Education , 27 (2), 122-129.

CMSWire. (2009). Content Management Systems (CMS) List. Retrieved from CMSWire3:

<http://www.cmswire.com/cms/products/>

Cross, J. (2006). *Informal Learning: Rediscovering the Natural Pathways that Inspire Innovation and Performance*. John Wiley and Sons.

Cuciz, D. (2004). The History of MUDs. Retrieved from GamneSpy:
<http://archive.gamespy.com/articles/january01/muds1/>

Dignan, L. (2008, June 30). Content management software vendors eye social networking. Retrieved from ZD Net: <http://blogs.zdnet.com/BTL/?p=9219>

Dixon, N. M., Allen, N., Burgess, T., Kilner, P., & Schweitzer, S. (2005). *CompanyCommand: Unleashing the Power of the Army Profession*. West Point, N.Y.: The center for the Advancement of Leader Development and Organizational Learning.

Downes, S. (2008). Feeds. Retrieved from Connectivism & Connective Knowledge - The Daily:
<http://connect.downes.ca/feeds.htm>

Downes, S. (2008). gRSShopper. Retrieved from <http://grsshopper.downes.ca>

Downes, S. (2007). Learning Networks in Practice. In e. David Ley, *Emerging Technologies for Learning*. London: BECTA.

Downes, S. (2001, July). Learning Objects: Resources for distance education worldwide. *The International Review of Research in Open and Distance* .

Downes, S. (2009). Serialized Feeds. Retrieved from course.downes.ca: <http://course.downes.ca>

Downes, S. (2008). The Daily. Retrieved from Connectivism & Connective Knowledge:
<http://connect.downes.ca/>

Downes, S. (2004, November 10). Theory of Learning Networks. Retrieved from Slideshare:
<http://www.slideshare.net/Downes/theory-of-learning-networks>

Downes, S. (2008, August 1). Tools. Retrieved from Connectivism & Connective Knowledge Blog:
<http://lrc.umanitoba.ca/connectivism/?p=59>

Drupal. (2009). About Drupal. Retrieved from Drupal: <http://drupal.org/about>

Elgg. (2009). About. Retrieved from Elgg: <http://elgg.org/about.php>

Foreman, J. (2004, September/October). Game-Based Learning: How to Delight and Instruct in the 21st Century. *EDUCAUSE Review*, 39 (5), pp. 51-66.

Garrote, R. (2007). The use of a Learning Management System to promote group interaction and socialization in a trainee project. *högskolor och samhälle i samverkan. Högskolan.*

Hirst, T. (2009, January 6). Serialised OpenLearn Daily RSS Feeds via WordPress. Retrieved from OUseful Info: <http://ouseful.wordpress.com/2009/01/06/serialised-openlearn-daily-rss-feeds-via-wordpress/>

Hegel III, J., & Armstrong, A. G. (1997). *Net Gain: Expanding Markets Through Virtual Communities.* McKinsey & Company, Inc.

IMS Global Learning Consortium, Inc. (2003, January 20). *IMS Learning Design Information Model.* Retrieved from IMS Global Learning Consortium, Inc.: http://www.imsglobal.org/learningdesign/ldv1p0/imsld_infov1p0.html

Levin, J. A. (1995). Organizing educational network interactions: Steps towards a theory of network-based learning environments. *American Educational Research Association Annual Meeting.* San Francisco.

Malone, T. W. (2004). *The Future of Work: How the New Order of Business Will Shape Your Organization, Your Management Style and Your Life.* Cambridge, Mass: Harvard Business School Press.

Marotzki, W., & Specht, M. *Matching Learners and Technology I - Community-based adaptive systems: Learning from learners, a smart idea? Interdisciplinary Approaches to Technology-Enhanced Learning.* Darmstadt.

Martínez-Ortiz, I., Sierra, J. L., & Fernández-Manjón, B. (2009). Translating e-learning Flow-Oriented Activity Sequencing Descriptions. *6th International Conference on Information Technology: New Generations ITNG.* Las Vegas.

McCall, J. (2005). *Simulations Games as Historical Interpretations: Critiquing Rome: Total War in the High School History Classroom.* Retrieved from *Historical Simulations in the Classroom:* <http://www.historicalsolutions.net/theory.htm>

Meyer, B. (1997). *Object-Oriented Software Construction.* Prentice-Hall.

Münste, T. F., Altenmüller, E., & Jäncke, L. (2002). The musician's brain as a model of neuroplasticity. *Nature Reviews Neuroscience*, 473-478.

Ng, B. D., & Wiemer-Hastings, P. (2005). Addiction to the Internet and Online Gaming. *CyberPsychology*

& Behavior , 8 (2), 110-113.

O'Brien, K. (2007, May 18). Web 2.0 Makes You ReThink the Basics. Retrieved from Marketing Thought: <http://marketingthought.blogspot.com/2007/05/web-20-makes-you-re-think-basics.html>

Peters, S. G. (2006). Do You Know Enough About Me to Teach me?: A Student's Perspective. King Lindsay Printing Corp.

Rollings, A., & Adams, E. (2003). <http://www.gamegrene.com/node/820>. In A. Rollings, & E. Adams, Andrew Rollings and Ernest Adams on Game Design (pp. 89-119).

Rushgrove, G. (2007, June 4). APIs and Mashups For The Rest Of Us. Digital Web Magazine .

Siemens, G. (2008, August 6). Narratives of Coherence. Retrieved from Connectivism & Connective Knowledge Blog: <http://ltc.umanitoba.ca/connectivism/?p=61>

Siemens, G. (2008, June 18). On Finding Inspiration. Retrieved from Connectivism & Connective Knowledge Blog: <http://ltc.umanitoba.ca/connectivism/?p=25>

Siemens, G. (2008, September 16). Where does the learning occur?? Retrieved from Connectivism & Connective Knowledg Blog: <http://ltc.umanitoba.ca/connectivism/?p=123>

Siemens, G., & Downes, S. (2008). Connectivism & Connectiove Knowledge. Retrieved from University of Manitoba: <http://ltc.umanitoba.ca/connectivism/>

Siemens, G., & Downes, S. (2008). Wiki. Retrieved from Connectivism & Connective Knowledge: <http://ltc.umanitoba.ca/wiki/Connectivism>

Squire, K. (2005). Game-Based Learning: Present and Future State of the Field. Retrieved from Masie Center: cecs5580.pbwiki.com

Vries, F. d., Tattersall, C., & Koper, R. (2005). Pre-Discussion Paper Future developments of IMS Learning Design tooling. Retrieved from Open University of the Netherlands: <http://dspace.ou.nl/bitstream/1820/553/1/discussion%20paper%20IMSLDtooling%20FreddeVries%20preprint.pdf>

Wikipedia. (2009). History of personal learning environments. Retrieved from Wikipedia: http://en.wikipedia.org/wiki/History_of_personal_learning_environments

Wiley, d., & Edwards, E. K. (2002). Online Self-Organizing Social Systems: The Decentralized Future of Online Learning. Quarterly Review of Distance Education , 3 (1), 45.

Wilson, S. (2005, January 25). Future VLE - The Visual Version. Retrieved from Scott's Workblog:
<http://zope.cetis.ac.uk/members/scott/blogview?entry=20050125170206>

Yendluri, P. (2003). Web services choreography. SOA News.

Moncton, April 25, 2009

How I Would Organize A Conference

The structure would be more like a market or a fair. Mostly, there would not be a schedule. Participants would certainly not get a schedule; organizers would have a bit of one, in order to choreograph the event.

The conference area itself would consist of a largish central area with various side areas with more or less privacy (the presumption is that while people will want to go to the side to chat, etc., they won't want to cut themselves off completely from the main event).

The main area would itself have various types of things, including:

- presentation booths - these are not lined up in rows, but are rather islands, round, surrounded by the attendees - exhibitors don't get a booth for the whole conference, rather, just for certain time slots (lots of rotation in the booths) - at any given time there will be many of these booths in operation - some of these (especially those in strategic areas) are no more than soapboxes, whereon a person stands and makes his/her pitch
- bearpits - where from time to time famous people are surrounded by an audience, where - instead of presenting a lecture - they answer questions tossed in by the audience
- demo labs - again, like the booths, nobody owns these, they are used for a certain amount of time - they consist of a large screen and about a dozen workstations - need lots of these, in various configurations - in the open, so people can stand around and watch
- a great big wall where people can put up any sort of notice or advertisement they please (people who spam the board will have their messages removed)
- electronic games and activities (could even make it possible to win 'tokens' by playing and use the tokens to auction things - check your local laws) -- there would absolutely have to be a Wii area - but also, there are many video-cam games (eg. the video cam boxing game), have those set up as well, whatever games people want to play
- Big screens everywhere - some of them are showing the games (especially the hockey games, especially the finals from the conference-long tournament), others are showing the conference 'backchannel' where participants (you need a conference login) post their thoughts
- The announcers, of course, on the speaker system, letting people know when an event is about to take place - a '10 minute keynote', a 'Flash video demo', the 'George Siemens bearpit', the 'EA Hockey semi-finals', 'the Wii-learn SIG in the alcove...
- Entertainment, including mainstage shows at noon and in the evening, side stage workshops (esp. with electronic music tools), wandering minstrals, jugglers
- The Vendor's Parade
- plenty of tables and chairs throughout, where people can sit and work or chat
- wireless and numerous ethernet ports with the bandwidth to back them up

There would be numerous side areas, including (pay for meals and coffee, etc., with conference

vouchers)

- coffee shops (how much would Starbucks pay to have a coffee shop there?) and cafes (meals tend to be 'ad hoc', not huge 'everybody eats at once' factory-style conference dinners)
- computer stores (Future shop? Apple)
- bookstores - a proper bookstore, not just a few titles on a table
- various types of pubs, some open and lively, others more like lounges, others with quiet out-of-the-way nooks
- very quiet areas, with couches for sleeping
- Art galleries / events (local artists are given a space and told to 'create')

You get the idea...

Now then, what we want participants to do is to *add* to this, in any way they can - we would want *lots* of ways participants can contribute...

- code jams - where coders create a new applications
- the participants' art gallery - any art, any way
- the conference radio station / podcast (which plays on speakers in various locations, including some of the cafes, as well as online (of course))
- 99-second presentations

Again - from the participant's point of view, none of this is scheduled ahead of time - what they are intended to do is to arrive and follow what interests them

What about papers? After all, that's how many people get funding to travel to conferences...

- the conference book and DVD - participants will be asked to come and, at some point during the event (probably have to sign up for a slot) go to a studio and record a presentation of their work - they can also bring in audio or video clips to add - these will be facilitated by program directors and hosts - the idea here is to make the presentations less formal, not just a person reading, but rather a person showing, being interviewed, interacting with the audience, etc -- a lot of this would also be the material used for the conference radio and shown on some of the video screens throughout - after the show, you can but the book/dvd from Lulu

- Typically conferences take place in convention halls and the like -- and there's nothing wrong with that -- but it's worth nothing that a conference as described here can fit into pretty much any (large) location
- school, college building, small town, whatever...

Moncton, January 30, 2007

What I Learned in High School

Following on a [thread from Clay Burrell](#) and [Harold Jarche](#)...

I actually learned a great deal in high school. Most of it wasn't the approved curriculum.

English:

In grade ten my English teacher Jamie Bell - a young idealistic educator full of new ideas - had us all do writing journals. It could be anything we wanted and - as I've mentioned before - I filled mine with stories, crosswords, drawings and more. From that point on I kept writing for myself, the way that project taught me, filling numerous notebooks before finding a web space in which to express myself.

I was also taught public speaking in English class. Technically this began in grade five. But it continued throughout high school. I won the school championships in grades five, eight, nine, ten and eleven (it was a small school).

I discovered science fiction in high school (specifically, John Christopher's 'The White Mountains' and Arthur C. Clark's 'A Fall of Moondust') and read that during English class instead of the official texts. I was also reading the classics (from a series that my mother bought) - Twain, Crane, Stevenson, Swift, Weiss, London and many more. These gave me a vista far more sweeping than the school texts, and let me see myself as (potentially) a hero.

In grade 12 I was supposed to read Dickens (we were finally done with years of Shakespeare) which I hated. I was tested on content (what colour was so-and-so's shirt) which I thought was degrading, so I boycotted then. I did take the final exam, though, and so managed to finish the year with a respectable (but still failing) 44 percent.

French:

I took 12 years of French and found I was not qualified to work for the government after I graduated because I was still unilingual.

Social Studies:

I enjoyed World Politics and I supposed I learned the basics of political systems. I enjoyed the model stuff that we did:

- model parliament (I managed to win 30 percent of the vote in a school-wide election running as the leader - and only member - of the Fascism Reform Party. This made me leader of the opposition in a

minority government. With the socialists I toppled the government and then made their party illegal. Now possessing a clear majority in parliament, I made the socialists illegal as well, thus becoming the only member of parliament. The governor general intervened so I shot him (no I didn't - I had a screaming fit and swore at Mr. Greenfield and his "bloody class" - I was very passionate when I was 16))

- model Premiers' conference (naturally, I was the prime minister - I wrote to the government for advice and got pages and pages of policy papers and procedural notes, which I basically committed to memory
- something that has served me well in chairing meetings ever since)

- model commonwealth conference (subbing for Jane Cooper as the representative from 'England' - and therefore having to crib overnight to prepapre) city-wide, in which I learned that the high class kids from Ashbury may be dressed to the nines, but they weren't any smarter than I was)

- model revolution - definitely an unsanctioned action, the 'Movement for Autocratic Organization' overthrew social sciences - I learned that it's easy to plan a revolution and to write a manifesto, but the reality is very very difficult to pull off - because you have to win the support of the people (aka the students) which is not such a simple matter.

Drama

I tried out and lost a male part to a girl. I learned some play. I learned I liked Randy-Lee Gbert (not a typo, it was a very odd last name). Nothing ever came out of either thing though.

Mathematics

Pretty much nothing. By the time I hit high school I was competent with basic mathematics. I could do geometry and measure areas and stuff like that. I zoned out right around the day they decided to teach me quadratic equations (which, somehow, I *knew* was specialized knowledge).

Science

I'm not sure when I learned the basic laws of motion, friction, force, acceleration, and all that, but I learned them.

In my closed ecosystem project I learned that nature needs sunlight. I also learned that going way way overboard on a project (I kept my ecosystem long after the rest of the class gave up, and then submitted a detailed notebook with graphs of months of measurements, drawings, references, theories, the rest) is sometimes rewarded (with a 20/10 woo hoo!).

Um, what else? I never did made anything explode (not for lack of trying). I learned I didn't want to cut animals open (and hence skipped most of the biology classes). I learned that the experimental; method

was a fraud (because we were doing 'experiments' but were penalized if we didn't get the 'right' results - which, of course, is a contradiction).

Geography

Oh, I loved geography. I learned the shapes of every country in the world, their location, their capital cities, their flags, their forms of government, their populations (roughly), their major exports, and some of their history. Mind you, I learned this from reading world almanacs, but I digress...

I learned everything about Ecuador. I did a major project on the Danube River (so it was a special thrill to finally see it when I went to Vienna, even if it has been rerouted far away from the city center). I wrote to the embassy of every government along the river (I picked the river because it had lots of governments) asking for information. They all responded - the communist countries right away, Austria next, and West Germany dead last and well after the end of the school year.

I did a project about the Northwest Territories and learned all the islands.

I learned urban geography which led to a lifetime habit of creating complex city maps in the margins (and sometimes on whole pages) of my notebooks. I remain to this day an inveterate critic of transit systems, highway intersections, left turn lanes, parkland planning, city profiles, and more. I was seriously tempted in later life by a career in geography (I didn't have the eyes for cartography, sadly, because it really was my first passion - I still love maps of all kinds).

I probably learned more, but that's the main stuff.

History

I learned all the explorers (specifically: the Vikings, Columbus, Cabot, Tasman, Hudson, Magellin, Drake, Frobisher, Mackenzie, Franklin, Thompson, Livingston, Stanley) and where they went (I made maps, of course).

I also learned enough ancient history in Ancient History class (it was an experimental class with only eight people - I sat next to Janet McGee (the girl who took my role in drama) and she liked me. I learned just enough about the Greeks and the Romans to become fascinated by them (my *actual* knowledge of the Greeks and Romans is based much later readings of Herodotus and Gibbon).

Economics

I took grade 11 and grade 12 economics in grade 11, discovered I was very good at it, and lost interest.

Art

I was pretty good at art, but never really received any instruction or technique. So, to this day, my only real artistic ability is to copy. I like taking photographs, though, and with my father set up a black and white photo lab at home.

Industrial Arts

I learned I was very good at drafting. I learned to print (that is, to print properly, with proper form, beautiful writing), which became my 'handwriting' thereafter (including even my signature, which is today (sadly) a bad scrawl - maybe time to fix that). I learned basic drafting techniques, including the three-sides diagrams and the exploding diagrams.

I learned how to weld. I could probably still run a bead - I really liked that. I learned basic carpentry (which was, essentially, how to use power tools without injuring myself).

Phys Ed

I hated phys ed.

I learned I cannot kick a 30 yard field goal, even if my grade depends on it.

I learned I was the 4th fastest kid in school (over distances greater than a mile). I learned that this does not translate into a starting spot on the soccer team. Apparently you need to be able to kick.

I learned the rules of curling and how to curl. As a skip, I am undefeated (1-0, lifetime).

I learned the proper way to hold a golf club. Sadly, this does not translate into 200 yard drives. It does, however, help a lot with putting.

I learned that if you show even the lightest weakness, people will exploit it; that if your clothes are even slightly ripped, people will rip them to shreds; that if you feel pain, people will inflict it.

25 years Later

I went to my high school reunion. My presence at the high school had been completely obliterated. there was no trace.

Moncton, February 4, 2008

My Personal Passion Trajectory

[John Hagel hits on an interesting thing](#), a way of describing how his interests have changed through time, merging to create his unique perspective on the world. Mine, too...

Construction

Perhaps because we lived in a subdivision when I was young, and watched the houses on the next street, then the next block, and so on, being built from scratch, construction was an early interest of mine. Perhaps all young boys want to build forts, but we had an entire subdivision's worth of scrap lumber and nails to play with. We built a fort in the back yard, one down the road, several tree-houses -- and then, when we moved to Metcalfe, the process started anew, using the wood from a woodpile that was all that was left of a demolished barn.

Exploration

Moving the Metcalfe put us in the country, and greatly increased my scope of activities. I had always explored to some extent, but living in Metcalfe meant following the creek one way to its source, and the other way down to the river, and the river down as far as we could go. It meant visiting one village after another by bicycle - Vernon, Kenmore, Russell, Embrun, Edwards, Osgoode, and in one challengingly long trip, Cassleman. It meant exploring the north woods two fields deep from the house, tracking through the swamp to highway 31. It meant creating local maps, and then buying World Almanacs and creating maps of the world. I cannot count the hours I have spent creating maps and diagrams of everything from fictional cities to plots to relations between people and more. It meant being in the outdoors, camping and canoeing, and above all, exploring.

Writing

I'm not sure what first got me on to writing but I do remember writing children's books as a child (which seems appropriate) - "Voyage Under Water" was my first, followed by several others. By grade 5 I was producing my own newspaper, "The Eagle Report", which went hand-in-glove with my after-school job delivering newspapers. In Grade 10 this interest took off when my English Teacher, Jamie Bell, had everyone keep a journal - this could be anything, he said, so long as it was something. So I wrote a series of stories - "The Adventures of Homer Higgens" - created crosswords, drew pictures, wrote poetry, and more. I would write hundreds of (bad) poems over the next five years.

Computers

My father worked for Bell Canada and always encouraged my interest in information and communications technology, which, by the time I was old enough for college, meant computers. My first

foray into post secondary education was a three-year certificate program at Algonquin College, majoring in computer science. I ran out of money and employment, though, and moved to Calgary, where I found a job with a computer company. I was right at home in the machine room filled with desk-sized Texas Instruments 'TIMAP' processors crunching geophysical data. I mapped out the sequence of programs we would run, got pretty good at debugging the card decks submitted by 'users', learned 'Multiple Virtual Systems' (MVS) and 'Job Entry Subsystems' (JES) and how to program some of the first TI computers - including one summer spent filling a cassette tape with 'the ultimate Star Trek program' running on Rex Hayes's TI-99. I also took more computer courses at SAIT, learning Fortran and - yes - doing more work on the card-punch machine.

Journalism

Passed over for promotion, I quit my job (and it would be 15 years before I would earn the same salary - \$1300 a month - again) and went to university. On my first day on campus I sought out the student newspaper and signed up. I started on news, covering education issues - my first story was a front page headline '\$1.5 billion program cut' on reductions to Established Programs Financing transfers for education and health care. I spent more time at the Gauntlet than I did in classes, writing regularly, learning layout and graphic design, taking news photos, and arguing politics with the other 'Gauntleteers' in the office. I was sports editor for a year, the first year our football team won a national championship, and was editor for two terms. I took my journalism seriously; I learned how to write well and quickly, I learned about propaganda and fallacious reasoning (the film 'Not a Love Story' and Eleanor Mclean's book 'Between the Lines' were major influences), and, well, much more. In the end, I wrote hundreds of articles, editorials and features.

Philosophy

I enrolled in university as a physics major, even though I needed to take remedial math, because I wanted to be a scientist. I aced the remedial math but bogged down on integral calculus and was overall a struggling B student. Except in philosophy - a subject I had taken only because English was full. It turned out that I was *very* good at philosophy (I rarely got less than an A) and, moreover, found it interesting as well. My interests and inclinations leaned toward the philosophy of knowledge (epistemology) and philosophy of science. I learned logic and analysis from some unrepentant logical positivists and got a *very* good grounding in scientific method, paradigm shift, models and representation. My interest in the philosophy of mind would follow. meanwhile, I came within a course or so of minoring in religious studies, examining things like religious experience, redactive criticism and Biblical analysis, world religions, and more. I found no contradiction in the different areas of study (or perhaps, more accurately, learned to live happily with contradiction). More and more I looked into the nature of mind and knowledge, cumulating with a connectionist theory of mind and similarity-based theory of knowledge, summarized in 'The Network Phenomenon'.

Politics

No journalist is far from politics, but after an unsuccessful run at student office in my second year I left it alone until graduate school. In my Master's program I became our department representative on the Graduate Representative Council (GRC) and was thence drawn (willingly) into the debates and demonstrations over tuition fees. When I went to Edmonton to study for a PhD I went on my very first day to the Graduate Students' Association and within a few days found myself the Communications Officer, a low-grade VP position. Before the end of the year we had overturned the president and embarked on a more radical agenda; I was VP Communications my second year and President for the two years following. I also ended up sitting on the university Board of Directors for a couple of years, as well as the Faculty council, some very eye-opening experiences. We were a very activist and very successful Council; I reformed the organization of the GSA based on democratic principles, made student support and lobbying our primary function (the GSA had been basically a well-funded social club before that), launched a (successful) lawsuit against the university, and for more than two years was in the newspaper at least once a week. I ran unsuccessfully for a nomination as a federal candidate in Edmonton Strathcona, worked on campaigns when I lived up north, and when I moved to Brandon, ran for mayor. The suicide of my close friend (and campaign manager) pretty much ended my interest in a political career, but I still have an (angry) interest in politics.

Distance Education

While studying for my Masters I earned money as a teaching assistant; this, combined with summer jobs as a programs coordinator for a local development education centre (the Arusha Centre) led to my being interviewed for a job with Athabasca University (I had wanted to focus on development education much more, having been very influenced by Francis Moore Lappe, but was told I would need to travel internationally before I would be considered, something that put it financially out of reach for me). There was nothing open in Calgary, but when I moved to Edmonton an opening was available almost right away, and so I spent the next seven years as a tutor for Athabasca University. I learned a lot about education and education theory and got a lot of practical experience teaching by telephone, in person (in remote northern communities) and even over the computer. I did some course development and a pilot program in computer audiographics. I also got involved with the Tutor's Association (a branch of the Canadian Union of Educational Workers) and ended up sitting on Athabasca University's governing council. In retrospect, my focus on distance education was pretty narrow, but it was very deep, and I learned a lot.

The Internet

While studying at the University of Alberta, Jeff and Istevan told me about this thing called a 'MUD' I could access using my University computer account. I had been using computers for years by this point, as they were invaluable for writing articles and philosophy papers (as well as for creating games, which I had never stopped doing). The modem thing wasn't too difficult; I had set up a Bulletin Board Service

(BBS) to support my distance education work. 128.227.96.54 2000 and I was in to Muddog Mud. Jeff, Ishy and I would keep building MUDs for the next few years; I figured out how to compile and run my own MudLib (most notably the Nightmare Mudlib) while all three of us were proficient in 'LPC', the standard code library. I spent half the time I live in the north connected to the internet, spending every cent I had on long distance and connect charges.

The Web

My interest in the internet - and in distance education - got me a job with Assiniboine Community College and I finally started making the same kind of money I was making with Texas Instruments. At the same time the World Wide Web arrived on the scene, so when I joined Assiniboine the first day of 1995 I almost immediately began working on a web site. The web was still a hard sell in the early days, so I had to do everything myself, installing my own server, creating my own pages, then a College web site, then demonstration courses, then our own learning management system. It was relatively straightforward to learn Perl, the web programming language, and web servers and CGI programs installed and compiled pretty much like Mudlibs. I was also drawn into the whole ethos of the web, partially influenced by the first few years of Wired and partially influenced by those first garish gushing web sites. I created course sites, Star Trek sites, a City of Brandon sites, some association sites, and more. I went to a bunch of NAWeb conferences and won a few NAWeb awards. The first, good, simple days of the web.

Online Learning

As I worked at Assiniboine I found myself connected through mailing lists like WWW_DEV and DEOS to distance and online learning practitioners worldwide. It was not a large community, and mostly an offshoot of the already large and well-established distance education community (so I knew the jargon and was familiar with the concepts). This brought me into the realm of learning theory, which was mostly new to me, but which fairly consistently was derived from the philosophical concepts I *had* studied. I did a few papers describing how we were building online courses and such, but basically things came together when I was asked by our own management to explain exactly what I was trying to do. This led me to write 'The Future of Online Learning', which almost overnight made me into a futurist and internet theorist. Another paper, written while I was working a year later at the University of Alberta, Learning Objects, established my tech credentials. I've spent rather longer on online learning than I ever intended. I certainly didn't expect to make it a career.

Blogging and Social Networks

I come into the field of social networks as an outsider. Blogging was a natural for me, and from the day I created my first website in 1995 I started posting my writing online. To this day, I don't follow the 'standard form' of blogging because blogging didn't have a form when I started, and by the time it was reified by services like LiveJournal and Blogger I was too set in my own way of doing things (with my own

software) to change. As my term at the University was ending I had a chance to go to Australia to do some web development for a philosophy course (which gave me that magical and elusive *international experience*) and it was while there I decided to start my newsletters. I had been, since 1998, involved with a group of friends on a news and politics website called NewsTrolls, which was a big part of my life, and I transferred that experience over to create OLDaily, a newsletter I have been running to this day. I wrote 'Educational Blogging' and then 'E-Learning 2.0' and began to be asked to attend some education technology conferences. I won a few blogging awards and began to think more seriously about the link between connectionism, networks and pedagogy.

Photography

Ever since my days bicycling to neighbouring communities I have had an interest in photography. As a child I had an enlarger and developed some of my own film. I still have some prized aerial photos taken from the day as a 14-year old, determined to at least experience my first flight, I cycled 10 kilometers to Embrun and hired a small plane for the 15 minutes I could afford with my 25 dollars. The 15 minutes became an hour looping back around Metcalfe and I emptied the roll. Working with the Gauntlet gave me the opportunity to work with free film and developing, but only in black and white. The arrival of digital photography - and the opportunity to travel, which finally came with my first *international experience* - created a renaissance in my photographic career, an interest I pursue to this day. While I've done some videos, I think, honestly, that I'll stay with photography; there's a beauty of a still image that can't be matched with a video. I surround myself these days with my own photos, and I have to admit, I had never known the world could be filled with such stunning beauty.

So that's the summation of my passions to this point in my life. the major ones; I haven't touched on some of the less major ones, like science fiction, darts, hiking, camping, cycling, old time radio, birds, gardening, flying, architecture, and more. I'm overdue for a new passion, though honestly, there's still a lot there in those other passions to occupy me for a lifetime. These passions have always driven me to distraction, the need for things like a stable income or home life taking very much a back seat. And I have some unfinished business - the novels I want to write, the philosophy I want to complete, the art and media I want to create, the places I want to see, and the house I want to build out there somewhere in my forest.

Moncton, December 21, 2010

My Academic Upbringing

Responding to Steve Hargadon, who wrote,

Do we believe in rigor and passion in our own educations? It's a hard message, but if our free time is filled with unchallenging and mindless entertainment, and if when we talk about our school days we speak of something that is behind us that "we got through," then our children will not know any better.

"When our major method for accomplishing something is enforcement (which is really what the culture of school is now), we give the implicit message that it is not something that is going to be enjoyed, no matter how much we say otherwise. Want to help your child become a better learner? Let them see you studying math or reading a classic..."

This comment is exactly right.

Neither of my parents were academics. Neither attended university (except some night classes my father took at Sir George Williams). But in our household, academic virtues were celebrated and practiced:

- the radio was tuned to CBC (Canadian public broadcasting) and so we would hear world news, scientific programs, 'Ideas', and more...
- there were always newspapers in the house - we all ended up delivering newspapers - and articles of importance, such as the current membership of the Cabinet (in the Canadian government) were posed on the wall. There was always a big map of the world on the wall.
- my mother bought a complete set of the classic works of literature for the house (these were very specifically my mother's, and I had to ask to read them), very cheap Pelican's (low-cost Penguins) that fell apart when you read them. Everything from Shakespeare to Butler to Thoreau to Twain. I read about half the 120 book collection before the middle years of high school (talk about an advantage!)
- I joined the Book of the Month Club with my father, through which I learned a lot of history - Pierre Berton, William L. Shirer, and Albert Speer all stand out, as do my Complete Sherlock Holmes
- there was also technology, and an evident interest in technology, in the house (we weren't just about academics). Our house radio was built from a kit. Bits and pieces of telephones were always about, as my father worked for Bell. We had telescopes and microscopes (much to the distress of the local bug and amphibian population). My younger brothers benefited from my father's interest in computers, but by then (1980s) I had left home. Still, I got my first model, a big 300 baud box, from my father.

- somehow I came into possession of an old Underwood typewriter (the reason I can't type to this day, because the keys took too much force to push) and a limitless supply of paper.

- I also somehow had access to tools - hammers, saws, screwdrivers, the works - to build things (and we built numerous things, including clubhouses, tree houses, go-karts and even a stage coach).

- we had a (large) garden and learned how to grow food. We were involved in preserving and canning the food (I can still remember piles of beans, a supply that would last the entire winter). We could cook basically whenever we wanted, so I took the opportunity to bake some cakes and pies.

Things like this - which, really, began with my first set of blocks, which had letters stamped on the sides, characterized my childhood. Knowledge and learning were always valued and supported.

At the same time, though, none of it was forced on me. These things were always in my environment, but I wasn't required to read the books (though the garden work was not voluntary - everybody helped because everybody ate). It was all about the environment, and not some rigorous academic regime.

Moncton, February 25, 2007

Social Media and Me

I was [asked](#):

- * How did you get started with social media?
- * What was your introduction, and how did the journey unfold?
- * What difference has it made in your professional practice?

Here's my answer.

My first experience with anything that could be called Social Media was in my Masters level 'Philosophy of Mind' course, offered by John A. Baker. I had to go to the computer lab in the department of Philosophy at the University of Calgary to gain access, but I participated in a lengthy discussion on the topic with the rest of the class (I still have the complete archive, on paper and bound).

My introduction to networking comes in the late 1980s with Bulletin Board Services (BBSs). It took me a lot to figure out how to configure the modem and connect to a BBS, but when I did, a world opened up to me. I took part in online chats, I posted to discussion boards, and I downloaded software.

I thought this was a great way to connect and I made some friends on the BBS systems in and around Edmonton. One of my favorite sites was run by the Alberta Alcohol and Drug Abuse Commission ([AADAC](#)), believe it or not. Another was run by some guy from south of the city in a small town; I remember visiting him once and thinking that he was like my father, an older guy with an office filled with electronics and other junk (it's a certain character type).

I was so enamoured of this online life that I eventually built my own BBS, using the Maximus BBS system. I had to create my own batch files to run the server, something that took literally weeks of labour (I still have the code listing at home somewhere). My BBS, Athabaska BBS, was less successful - people had trouble connecting, just as I had. And as my interest was province-wide, long distance was an issue. As was the fact that it tied up my telephone line.

My next encounter with social media was the Multiple User Dungeon (MUD). This was also my first introduction to Internet (ie., TCP/IP) social media. I recall it clearly. We - the philosophy department - had finished losing another [slo-pitch game](#) and were recovering in Tom Daly's Chop House (and pub).

Two friends, [Jeff McLaughlin](#) (Kane) and [Istvan Berkeley](#) (Nomad) had already joined the MUD, Muddog MUD (based at the University of Florida). (Istvan and I shared a house in Edmonton for a while and [explored connectionism](#) together). They described how to join, and I figured I would. We discussed my character name for a bit, and I settled on Labatt - it was short, easy to type, and was suggested by the beer posters on the wall (Labatt is a brand of beer in Canada).

The next several years spent on Muddog were transformative. I worked my way up the ranks and eventually became a wizard, responsible for contributing to the MUD programming. *En route* I made a number of friends (one of whom I eventually married) and went through many shared experiences. It's the sort of thing that has been [written about elsewhere](#), about other people ([Sherry Turkle](#) is the [danah boyd](#) of the MUD set). Muddog closed down in 1994.

Jeff, Ishy and I worked on some other projects together, most notably the '[Painted Porch MAUD](#)' (Multiple Academic User Domain, which I coded from a Nightmare MudLib). We caught the attention of some other people, most notably Wes Cooper and Terry Anderson. While I was living in a cabin in Northern Alberta Jeff sent me an email describing the job at [Assiniboine](#) Community College in Brandon, a job that my MUD experience uniquely qualified me for.

I made the jump from telnet to the web when I moved to Brandon in 1995. In Brandon, my first experience with social media was a site called [The Spot](#). I was one of the original 'Spotniks' back in the days when it wasn't certain whether or not The Spot was real (it wasn't - it was launched and run by [Fattal and Collins](#), an advertising agency. The Spot sort of drifted away before being eventually closed in 1997.

I remember visiting my family in Ottawa over the Christmas break of 1996. I spent most of my visit exploring a site called [Firefly](#). Basically, it was a prototype social network, with friends and recommendations and many of the features we see today. When I returned back to Brandon, Conrad Albertson and I had long discussions about the site - and the sorts of things we could do online. Firefly was eventually bought by Microsoft, which killed it.

Around this time I picked up [ICQ](#) - as I recall, at the instigation of my father - which is probably the first genuinely social software I tried (in the sense that you have a list of 'friends' or contacts). ICQ is the original instant messaging software. It was [launched in 1996](#) by Mirabilis, an Israeli company. My ICQ number, 1287181, was given out in 1997. My list of friends included mostly friends and family - work acquaintances never did take to it. Significantly, ICQ allowed my father and I to reconnect after many years, and we had a good dialogue in the last years of his life (he died in 1998) and with Andrea, who I married in 1998. ICQ was eventually bought by AOL, which cannibalized it to promote AOL Instant Messenger.

After I left the Spot I landed at the [HotWired](#) website and in particular the Media Rant discussion area. Here I came to know quite a number of people. I also started saving my online postings (having seen too many communities close - you can see them collected [here](#), way down at the bottom of the page). This was prescient - the Threads part of the site closed in 1998, not long after I posted my [lament](#) for the once-great magazine. A bunch of us left the site in the fall of 1998 to form [NewsTrolls](#), which continues to this day, but which, over time, has lost all its members.

Also around this time (starting in 1995, with my job at Assiniboine) I had my first real experience with email discussion lists, specifically [WWWDEV](#), run by Rik Hall at the University of New Brunswick, and [DEOS-L](#), run by Mauri Collins (among others) at the University of Pennsylvania. This connected me with the more academic community, and especially the Canadian community through WWWDEV. I attended a number of WWWDEV conferences and had many too many late-night beers with people like Rik, Terry Anderson, Rory McGreal, and the various other people the conference attracted.

In 1999 I was attracted from Assiniboine to the University courtesy of a job offer from Terry Anderson, who hired me to build an online community for the municipal sector in Alberta. My work in online community as well as my experience building websites and courses for Assiniboine qualified me for this job. I learned a lot creating [MuniMall](#), not only about content syndication and community building, but about politics and communication. And I must say, I am very gratified to see MuniMall still alive and still using my original design, eight years after I left the University of Alberta. It's gratifying, and it tells me I got some things right.

While at the University of Alberta I started my [email newsletter](#). I had launched my website in 1995, and in 1998 built the engine that I used to post news on my website and also on NewsTrolls (and a city of Brandon website, [The Brandon Pages](#), which was beautifully designed, but attracted no community). But in 2001, while on a [three-month visit to Australia](#), I decided I needed to expand my web presence - my job was coming to an end and I thought a newsletter might help me get paying work.

In the summer of 2001, Rory McGreal contacted me, telling me I should apply for this NRC job in New Brunswick. It was, again, another position for which my community-building experience perfectly suited me. And since taking that position in November, 2001, I have expanded my personal social network in two ways: first, by an ever increasing list of subscribers to my newsletters and RSS feeds, and second, by an ever increasing list of feeds in my RSS aggregator, the people I read on a regular basis. I now read about 500 feeds regularly, and I distribute my newsletter to 3000 email subscribers and roughly a similar number of RSS subscribers.

From there has been a proliferation of social networking tools. I have used [Skype](#) a lot and have a long list of Skype friends. My list of [Orkut](#) connections, by contrast, is small and silent. I have a long list of connections on [Facebook](#), but my only real use of the site is to play scrabble (though I was invited to speak at a couple of conferences via Facebook messages). People follow me on [Twitter](#), because my Facebook status update outputs to my Twitter account, but I do not read Twitter.

My main social network today is my own website. Mostly because I know it's the one site that nobody is going to close on me.

You work with these sites long enough, and you come to realize, if it belongs to someone else, it's going to die, eventually. No matter what kind of site - whether a MUD, a bulletin board, a discussion list, an online community, a communications app, whatever. Either it will shut down outright because it was no

longer profitable (or no longer fit the 'funding mission') or whatever. Or it becomes so commercialized that it becomes useless, a cesspool of advertising, spam, and other detritus.

That said - if you follow this history, you will see that pretty much everything I have done in the last fifteen years has been the result of social networks one way or another. *All* of my jobs have come from the network. Most of my ideas, my work, and my thoughts and opinions have been shaped through the network. Even my personal relationships have been established through the network.

Moncton, May 28, 2008

Seven Habits of Highly Connected People

With apologies – and all due credit – to [Stephen Covey](#).

1. Be Reactive

There's a lot of talk about user-generated content on the web. That's great. But if publishing your own stuff comes at the expense of reading, and commenting on, other people's stuff, that's not so great.

The first thing any connected person should be is receptive. Whether on a discussion forum, mailing list, or in a blogging community or gaming site, it is important to spend some time listening and getting the lay of the land.

Then, your forays into creating content should be as reactions to other people's points of view. This will ensure, first of all, that they read your comment, and second, that your post is relevant to the discussion at hand.

Posting, after all, isn't about airing your own views. It's about connecting, and the best way to connect is to clearly draw the link between their content, and yours.

2. Go With The Flow

We all know those people in our online community who are out to Prove Something, to Get Things Done, or to Market Themselves.

These are people we tend to avoid. Because no matter what the topic of discussion, they'll weigh in with their pet project, peeve or talking point.

When connecting online, it is more important to find the places you can add value rather than to pursue a particular goal or objective. The web is a fast-changing medium, and you need to adapt to fit the needs of the moment, rather than to be driving it forward along a specific agenda.

This doesn't mean you shouldn't have any goals or principles for yourself. You should; that's what will inform your participation. It's just a reminder that your goals are not the same as other people's goals, and therefore that your online participation needs to respect that fact.

3. Connection Comes First

People talk about not having time for email, of not having time for blogs. Sometimes they even talk about working without an internet connection.

It's good to take a break and go out camping, or to the club, or whatever. But the idea of replacing your online connecting with busy-work is mistaken.

In almost all fields, connecting with others online is the work. The papers you write, the memos you read and toss – all these have to do with connecting with people. Even if you work with your hands, making cabinets or rebuilding engines, all your contacts with customers and suppliers are about connecting with people.

If you don't have enough time for reading email, writing blog posts, or posting discussion lists, ask yourself what other activities you are doing that are cutting in to your time. These are the things that are often less efficient uses of your time.

If you are spending time in meetings, spending time traveling or commuting to work, spending time reading books and magazines, spending time telephoning people (or worse, on hold, or playing phone tag) then you are wasting time that you could be spending connecting to people online.

If you make connecting a priority, you can take that walk in the forest of vacation in Cadiz without feeling you are not caught up.

4. Share

We're all heard the advice to "think win-win". Forget that advice. If you follow that advice, you will always be looking at things and saying, "what's in it for me?" That's exactly the wrong attitude to have in a connected world.

The way to function in a connected world is to share without thinking about what you will get in return. It is to share without worrying about so-called "free-riders" or people taking advantage of your work. In a connected world, you want to be needed and wanted. This will, over time, cause resources to be sent to you, not as a reward for some piece of work, but because people will want to send you stuff to help you to be even more valuable to them.

When you share, people are more willing to share with you. In a networked world, this gives you access to more than you could ever produce or buy by yourself. By sharing, you increase your own capacity, which increases your marketability.

5. RTFM

RTFM stands for 'Read The Fine Manual' (or some variant thereof) and is one of the primary rules of conduct on the internet.

What it means, basically, is that people should make the effort to learn for themselves before seeking instruction from others.

Almost everything a person could need to know has been recorded somewhere online (by people who are sharing their knowledge freely). Taking the time and effort to look at this work is not merely respectful, it demonstrates a certain degree of competence and self-reliance.

For example, if your software fails to install, instead of calling customer service or posting a note on a bulletin board, copy the error message into the Google search field and look for answers. Almost every software error has been encountered (and documented) by someone before you.

Finally, when you do ask for help, you can state what you've read and tried, and why it didn't work. This saves people from giving you advice you don't need, and helps them focus on what's unique about your problem.

6. Cooperate

Offline people collaborate. They join teams, share goals, and work together. Everybody works in the same place, they use the same tools, and have the same underlying vision of the project or organization. Online, people cooperate. They network. Each has his or her own goals and objectives, but what joins the whole is a web of protocols and communications. People contribute their own parts, created (as they say in open source programming) to 'satisfy their own itch'.

This is probably the consequence of distance. Online, it is not possible to enforce your will or (beyond a limited extent) to get your way by shouting and intimidation. This means that online communications are much more voluntary than offline communications. And successful online connectors recognize this.

To cooperate, it is necessary to know the protocols. These are not rules – anybody can break them. But they establish the basis for communication. Protocols exist in all facets of online communications, from the technologies that connect software (like TCP/IP and HTML) to the ways people talk with each other (like netiquette and emoticons).

7. Be Yourself

What makes online communication work is the realization that, at the other end of that lifeless terminal, is a living and breathing human being.

The only way to enable people to understand you is to allow them to sympathize with you, to get to know you, to feel empathy for you. Comprehension has as much to do with feeling as it does with cognition.

People who use online communications ‘only for business’ – or worse, feel that other people shouldn’t be posting cat photos or playing Scrabble on Facebook – are employing only a small part of the communications capacity of the internet.

Learning and communicating are not merely acts of sending content over a wire. They are about engaging in (what Wittgenstein called) a ‘Way of Life’. Having a cat is as important for a physicist as having an advanced research lab. These common everyday things form the mental structure on which we hang the highly theoretical structure.

The idea behind ‘being yourself’ is not that you have some sort of offline life (though you may). Rather, it’s a recognition that your online life encompasses the many different facets of your life, and that it is important that these facets all be represented and work together.

Moncton, April 8, 2008

The Reality of Virtual Learning

Transcript of a keynote address delivered to the Canadian Defense Academy, January 30, 2008.

There's always a danger when you come in and you do a talk like this: the idea that you're presenting something, it's the facts, the truth this is the way it is, I'm going to lay it out on the line, I'm the expert, the guru, you're the mob, etc.

That's not how I view this material at all. So I want to preface this. I'm looking at a particular slice of online learning, not the whole field of online learning, obviously.

There are many things, including a wonderful helicopter simulation in Gagetown that I got to fly once and crashed very quickly –

Voice: seven seconds.

Is it still a record?

Voice: yes. [Laughter]

All the way across to the Second Life stuff that you saw, to the Company Command software that you heard about yesterday, from Nate Allen, and the rest.

This is just one small part of it and it's one person's perspective on it. It's not all a seamless whole. It's not a beautiful step-by-step narrative that wraps up nicely and neatly in a bow like a Sherlock Holmes movie or short story.

It's something that's a little bit loose and scrambly and something that you should take, look at critically, analyze it in your own perspectives, take what you need and discard the rest and take it from there.

So I want to begin this talk by talking about reality. I know that seems like an odd sort of place to begin a talk. But perhaps it's appropriate because finish a demonstration of Second Life.

A lot of the times we would be here among our own discipline, our own domain being talked about as a virtual reality. But my background is in philosophy and as a philosopher I've learned over time to look at reality in different perspectives and in different lights.

We often hear the phrase 'the reality is' as a rhetorical device and I'm sure you've all heard it, right? You have to have this great idea, "I think we should do such and such and so on," but someone comes along

and says, “Yeah, but the reality is...” Sounds familiar?

‘The reality is’ is the enemy of innovation.

When we look at reality, when we analyze what reality is, almost everything that we think is real (properly-so-called) is construction. It’s an artifact. It’s something that we’ve created. It’s a way that we understand the world. It’s a device that we use in order to understand the world.

Kant talked about ‘space’ and ‘time’ as the necessary prior conditions of perceptions. The idea here is that we don’t know space and time exists absolutely certainly for all time, but as perceiving, thinking humans the only way for us to make sense of our perceptions whatsoever is to come up with the concept of ‘space’, come up with the concept of ‘time’, to create a framework in order to place our many different perceptions and come up with some understanding of them, an explanation of them, a way of dealing with them. Hume called them ‘useful fictions’.

There are many ways of looking at ‘real’ in today’s world.

We think of ‘real’ on one hand, we contrast it with the ‘artificial’. So the real, it’s like the natural, the artificial is like the fake.

Now we can think of ‘real’ as ‘genuine’, as the actual thing or real as the fake thing that is not the original. It’s just a pale imitation. We have – “it’s the real thing,” and by implication, Pepsi is the fake thing.

(Pepsi should have done that. They should have come out with the advertising. I bet you it would have worked. Pepsi: fake Coke. [Laughter] That would have worked I tell you.)

These days of course we also have the contrast between the ‘real’ and the ‘virtual’. Analyze that with what we mean. Because the ‘virtual’ is every bit as real as the ‘real’ except that it’s not physical. But I mean it’s not like the things that are going on in Second Life are artificial or fake. Well, now, maybe they are.

How about ‘real’ versus ‘illusory’? We have the idea of real, something that exists. A mirage is not real. Not because we don’t perceive it. We do perceive it, but it’s not real because it does not in fact exist.

Now does this sense of real apply to something like Second Life? When we look at Second Life would you say, “Hm, that doesn’t exist?” Not really. That’s not the sort of thought that goes through our mind, is it?

We have ‘real’ versus ‘delusion’. The idea here is that having a grasp of reality is having a grasp of sanity, being able to make sense of your perceptions without all the static or interruptions that somebody who

is deluded might think. Is there a great conspiracy? Do we need to wear tin foil hats? Well that's not reality, is it? And hence we have the admonition to people, "Get real," as "Become sane again."

Now there are many ways to find reality, many ways of sensing, perceiving, touching, measuring, many different points of view. We have different perspectives, different models. The models tell us what things exist and what things don't exist.

It's an interesting thing. Here we have – well, look at this right in front of you. We have here a podium, right? We say this podium exists. This light exists. This light is part of this podium. It's almost artificial the way we give this particular podium an identity.

It depends on a certain perspective, a certain world view. And so, as I talk about this think about what the "realities" are in real world. I think about all of the challenges that National Defense is facing these days and there are many different realities, shifting realities, changing realities that you're dealing with.

Think about this. The 'realities', the constructs that we will accept as real, that characterize your institution, that characterize your students or your fellow staff, that characterize your values. What of your values is real? What of your values is artificial? What is the construct? What is faith? What is deluded? Your finances. My finances are very, very heartbreakingly real.

So, against that background I want to draw out the traditional conception of knowledge. The traditional conception of knowledge is exactly that conception that I tried to warn you against as I took to this podium today.

The traditional conception of knowledge is where things like knowledge, facts, values, and institutions are 'real' in all the full senses of the word and unchanging. They're presented. Here they are. You will consume.

Even 'change', which is kind of ironic, is viewed as inevitable, is viewed as real, is viewed as something that is out of our control. People talk about change management. Change management is getting people to accept change. That is, something that they cannot control and have to deal with. That's one way of perceiving change management.

The reality is, especially in today's world - and I'll have different ways of talking about that as I go through this talk - the reality is we define what counts as real.

We define it by our theories, our world views, our underlying values, our moralities, our religions, our perceptions, our beliefs, our intuitions, our different systems of logics, mathematics, all of these things combined to create a reality and this reality is different for each person and is very different for each person.

Look at all of you right now. The reality is based partially on perception. You're all looking at me; true (it's enough to give me a complex). But if reality is what you perceive, each of you is perceiving me from a slightly different angle which means that we've got about 100 different versions of Stephen Downes sitting in this room, which is really more disturbing. [Laughter]

Here's another exercise and this is in a particularly relevant exercise for this organization. What is a 'student'? What is a natural student? What is a genuine student? What is a physical versus a non-physical student? An existing or non-existing student. An actual or non-delusional or non-deluded student?

We have this way of dividing the world. We have teacher and students and never the twain shall meet, except in a classroom. But these categories, even the idea of real and unreal, existing or non-existing students, are beginning to change. We use the term fuzzified. Fuzzified, yes, is a word.

Here's another challenge for you about real. What is learning? Particularly with the emphasis on testing and repetition and all of that, we may get the impression that learning is something that we have a very good grasp of, that we know what it is. It's 'remembering', say or something like that.

But what is learning? What is natural learning? What is genuine learning? What is physical learning? Learning from experience as opposed to learning from virtual reality. Are those different kinds of learning, or is it the same kind of learning created in a different way?

What is existing learning? How can you have existing learning if learning isn't a thing? You have a thing that isn't really a thing; you can't say that it exists. You see what I mean?

This is the problem with studying philosophy. You look at these words. You come into a domain like education and people say, "well there is learning and there were facts and there were textbooks." You look in that and say look, "yeah, but the world isn't like that."

The reality is, learning has changed. Learning has changed from being about reality, from being about facts and objects and things and dates and names and ways of building things and ways of taking things apart, to being about ways of creating understanding, coping with, managing reality, to verifying reality, to being able to determine what is true in our perspective, what is not true in our perspective, to creating reality, to making reality. The skill that it takes to build something in Second Life is a type of learning, but it's a learning how to produce, how to create.

It's learning in an age that has changed. We used to live in a world that was very certain, that had a cause and had an effect. You knew what created a positive effect, if you wanted to cause it. These days, with multiple, complex, interrelated variables, even one thing means we sometimes don't know what's going to happen.

It's learning in an age of obscurity where the reality is simply not known. The reality is beyond our perception, like Kant's space and time.

It's learning in an age of chaos with multiple, independent variables, where from the same cause you can get one of a range of effects. It's like predicting the weather, which is particularly rough these days.

It's learning in an age of change, where what was true today ceases to be true tomorrow.

There's change that's very much about the way we have to go about learning. The old model is – well, this diagram (the SCORM diagram) is just a few years old - the old model is less than a decade old. The old transmission model.

The idea of learning as content, as facts and you take the facts and you assemble them in some way and then you run through the gears of a learning management system and then - I love the delivery device shoots them into people's brains, you know, as though if you shoot something hard enough it'll lodge in their brain and become learning. That's the old transmission model.

The problem with that model isn't that it didn't work. The problem is that it did work. The problem is that, that kind of learning shoots facts into people's brains in a world where there weren't facts.

Learning is not remembering. Learning isn't simply acquiring facts because that will not leave you the capacity to deal with the world. It will not allow you to create 'space'. It will not allow you to create 'time'.

Learning - you can remember things without meaning. This is from Lewis Carroll – "T'was brillig and the slithy toth..." I remember those words. Haven't a clue what they mean. So, would I be said to have learned those words if I have no idea what they mean?

Or, another example – mathematics. Everybody took mathematics in school and most people passed that. Yet when they go into the work place – I used to work in concession stands and 7/11 stores and things like that - every time an employee comes to the store we teach them a skill. We'd teach them what's 'counting change'.

The reason why we teach them to count change is that even though they've passed grade 12 mathematics, they do not understand that if somebody gives you a 20 for a 79 cent item, you don't give them a 50 change. Because they don't associate the mathematics with the value.

The syntax of the mathematic doesn't associate the semantics of being a 7/11 clerk. So 'making change' is a process that ensures that the change corresponds to what was actually given to them. So if somebody gives you a 20 you count the money back and add it up till you get 20 again.

So they can learn math, oddly, interestingly without understanding what math is supposed to do. We actually see that all the time.

Learning is not 'content'. I know, it used to be, "content is king, content is the web." Learning is not content. It is not shooting those facts into your head.

Rather learning, as I characterize in this slide, is a process of 'becoming' rather than a process of 'acquiring'. Now when I say something like that I sound like one of those German philosophers and I'll just start making up words.

That's not really what I want to do, but rather, what I want to say is, learning is a process of creating a particular mental configuration, particular set of connections between the neurons in your mind.

Learning is shaping yourself rather than acquiring something. To learn, as the slide says here, is to instantiate patterns of connectivity. So what you're doing is like exercising. You don't make someone strong by putting muscles into your arm (that'd be kind of neat, like Schwarzenegger, "I'm Ah-nold...").

But you don't put muscles into your arm. You have to grow muscles, you have to develop them and you do that through certain processes. I imagine you guys understand those processes much better than I do.

What learning really is, on this model therefore, is 'not propositional'. It's not a bunch of sentences. It's not a bunch of facts. Very often, it is tacit, to use the words of Michael Polanyi, Personal Knowledge, 'learning is like riding a bicycle'.

You could not describe what it is to know how to ride a bicycle. You'd sit there and you could write it out. Write sentence after sentence after sentence and the person could read all those sentences and still not know what it's like to ride a bicycle.

Knowing how to ride a bicycle is, again following Polanyi, in an important sense ineffable - it is not expressible in words, and as an aside, this means that the efforts to (as they say) "capture" tacit knowledge are misguided. What they are attempting to do is take something that is ineffable - not expressible in words - and express it in words. Obviously when you do that you change something important. You've taken something that is knowledge and changed it into something that is not knowledge.

Learning and knowledge are also, very importantly, personal. What you learn, what you know, depends on context. That means in a very important sense you can't generalize it. You can't say because I know this fact about the world, this fact about the world should be equally known by every other person in the world. Because what I know about the world very much depends on my own perception, my own background, my own culture, my own experiences.

What knowing is – on this picture - changes as well. Knowing to know something used to be: to know the rules, to know the categories, to put things in their place, to understand the laws of nature. On the new picture knowing now is much more about patterns. It's much more about similarities. It's not knowing that a tiger is a kind of cat; rather it's being able to recognize a tiger when you see it.

We can see how this works from the perspective of a network. We see both the aspects of recognition here and we see the aspects of ineffability, or tacitness, here. This is a network. It's a very stylized network. A network of this type is properly known as a neural network (this is a term from connectionism more than a term from neurophysiology).

The idea here is that our perceptions, such as that perception of a tree, corresponds to a pattern of connectivity in that network. It's kind of hard to see here, but you see the darker red lines that are intended to represent that.

So you see the tree and as you see trees over and over and over and over and over through your life, that creates a pattern of connections between various neurons in your mind. Obviously this is a gross simplification. That's about 20 neurons and the mind has about ten billion or whatever neurons. So the actual patterns of connectivity are very different.

Now what's important here is that no single neuron corresponds to the perception of a tree and there is no propositional representation of that perception of a tree. It is what we would call 'sub-symbolic'.

Another aspect of a network that operates in this way is that the same network manages multiple perceptions. Here we have a network that has three separate perceptions. It can recognize three different types of objects. It can recognize a tree, which again is strong connections in the red. It can recognize this cute little puppy dog with the – oh, I can't really see the colors - the different colored connections. Then it could recognize a couch. Again, different colored perceptions.

What's really important here is to understand that the same network carries the representations of all three of those subjects. Those representations aren't – they're not like pictures. They're not like words. They're like patterns of connectivity.

What's interesting about this is: change my perception of a couch and even though it is not in any way logically connected to a tree, it changes my perception of a tree because they (the perceptions) are operating in the same (neural) environment.

The mechanisms by which networks learn, the mechanisms by which networks form these connections, tell us about the mechanisms through which learning happens in humans and, for that matter, through which learning happens in society or networks of people in the world.

I've identified three major types of mechanisms by which these connections are formed in this diagram here.

One is Hebbian associationism, which is based on, as this slide says, concurrency. You have two neurons are active in it at the same time and not activated at the same time, a connection is going on between them.

(Second) based on that propagation, two neurons draw connected between them and then send the signal forward. From that signal forward a signal is sent back that may correct that original creation of the connection.

Then finally (third), the Boltzmann mechanism, which is a 'settling out' mechanism. We have a set of neurons and a set of connections connecting those neurons and the Boltzmann mechanism will try to find the most stable configuration, a configuration of lowest energy. Think about it by analogy throwing a stone into a puddle of water. The puddle of water will ripple for a while, but eventually it will settle out into the most stable configuration, which is, for water, flat. The mind works the same way. The connections go back and forth to create perception in a person's mind. You can visualize the connections going back and forth, back and forth, and eventually the mind just sort of settles and it settles into the most stable configuration.

Obviously as I've talked about those principles of association I've glossed over quite a bit of detail there. But that gives us a picture or a model of what learning looks like in this network environment. I've represented it in this diagram. This diagram merges or combines two major functions.

On the one hand the function of teaching, which could be said to be 'to model' or 'to demonstrate', to present perceptions for other people to perceive, things for them to emulate, things for them to follow.

Then on the other hand, the process of learning itself, which is 'to practice' and 'to reflect'. To practice is a lot like exercising. Then to reflect, which is similar to the Boltzmann mechanism, to try to comprehend, to understand putting the experience into context, into line with the previous beliefs, the experiences that you've had.

These four things combined together, modeling, demonstrating, practice and reflection, form the four cornerstones of the learning experience.

And then through that learning experience is the personalization or the context dependency of the learning experience beginning up at the top with the exercise of choice by the learner, which manifests itself in a definition of an identity for a learner, which in the end expresses itself as creativity.

Now, very important: this is a model. This is not the way the world is. This is a framework that I'm using in order to try to grasp or apprehend this process. I'm not coming here and saying, "oh you know, the

whole world is defined by these seven elements.” That’s not what I mean at all. I’m trying to give you a perspective or a point of view on the process.

So we have, in summary, E-Learning 2.0, the idea that learning is not based in objects and contents and books and course objects and classes and lessons that are stored as though they were in a library, but rather the idea that learning is like a utility, learning is part of our experience, learning is something that flows, learning is something that’s dynamic, that changes. It is this network, both externally and externally reacting and adapting to a wide variety of perceptions.

So that leads us to the underlying concept of E-Learning 2.0. The first aspect of that concept is that learning is learner centered. That is to say it is centered around the interest of the learner.

I would also say – it’s not on here – but I would also say it’s centered around the perspectives of the learner, the situation or the context that the learner find themselves in, the job that they’re doing and all of that.

Learning is also important when it’s owned by the learner, rather, and again I’m just simply being the -- reception of content is something that is managed, created and deployed – bad word, I don’t like that word – by the learner. I don’t want to say ‘used’ because that would imply that it exists and I don’t want to say it like that either.

But you get the idea what I mean. It’s something that the learner does. The learner is not a passive receiver. The learner is an active creator of their own learning.

The second major aspect: we’ve seen this already in presentations today. It’s immersive learning. It’s learning by doing.

The third major aspect is, it’s connective learning. It’s connective learning in the sense that we’re creating connections in the mind, but it’s also connective learning in that the learning occurs through the process of creating connections with other people in the world: teachers, other learners, colleagues, whatever. And working with those connections, receiving information, sending information.

Learning is, as this slide says, based on conversation and interaction. Some examples - and again we’ve seen this already with Second Life and the simulations that you all use - game based learning.

[Clark Aldrich](#) has identified four major types of learning: the branching, spreadsheet, quiz game simulations kinds of learning.

Just as an aside, for those of you who worked with SCORM and Learning Design: SCORM and Learning Design really amount to only the branching type of learning and not the other types of learning. I would not think of a Storm application that was a spreadsheet type of a game like Sim City or something like

that.

Another example of this learning is what Jay Cross calls workflow learning or informal learning. This is learning where the learning and process of doing your job or doing your work happen at the same time and in the same place using the same device.

So, electronic performance support systems, for example, which are applications that actually provide learning inside the application that you're working on. I'm sure you're familiar with those.

Where we see this most of all interestingly is in games where, if a person is playing a game, the learning for the game has to be inside the game because the game player will not read the manual; absolutely won't read the manual. There's no point in giving them a manual. So, whatever learning the player is going to need during the playing of the game has to be inside the game.

Also related to this is the model of the community in practice that's described by Etienne Wenger. Again, the community in practice, the idea of the community in practice, is that this community exists right inside the work place.

You might wonder about that. We did some work with municipal officials in northern Alberta and we asked them what their primary means of learning was on their job. Their primary means, if they needed to learn something, what they did was they picked up the telephone and they called somebody. They called a town manager in another town.

Basically what their learning model was was to have access to their community of practice right on their desktop. They picked up the phone and called somebody. That was learning for them. That model can be applied to the online world as well. It can be applied to the much more complex learning and working environments we find ourselves in today.

The environment, visualization technology in games: the helicopter that I so famously crashed is an example of learning inside an environment being able to provide visualization for me to understand what it feels like to fly in a helicopter.

Another aspect is mobile learning. One of the things that I've always said from the very beginning is that E-Learning is an online learning, learning not about being tied to a computer. They're not about sitting at a desk looking at a computer screen.

Concordant with the idea of immersive learning is the idea of learning in the place where learning is most appropriate. A slogan I've used for many years is "the best place to learn about a forest is in a forest; the best place to learn about law is in a courtroom." What these technologies now provide us is access to learning materials in the physical situation where learning is most needed.

One aspect of that, and that's just as an aside, is again learning – I want to emphasize – learning isn't simply the consumption of information. Learning is also the production of information. So, learning and multi-learning isn't just about getting content like flash cards or drills or instant messages, it's also about resource capture.

Here I am on the road – I have a microphone there. I'm recording this talk; somebody else is recording this talk (Hi, all of you out there in video land - I always like to put a personal message to people who are...) . It's capturing the learning as well.

One of the major things that we see in E-Learning and especially E-Learning 2.0 is the idea that the process of learning is inherent with the process of capturing and creating new learning content.

Online learning, on the one hand, we all look at our own environments, we develop tools and systems intended to support traditional classroom learning: the learning management system, learning objects, SCORM, learning design, that whole infrastructure. I'm sure you're all familiar with it.

On the other hand, what we should be doing, what we need to be doing and what we could be doing is developing tools and systems to support immersive learning, personal learning, dynamic learning - as I say they're living systems.

The first iteration of this is user-produced media. As I say blogs and blogging is a very simple example, podcasting and vodcasting.

Again, Nate talked about Company Command yesterday. One of the major things that makes this site like Company Command work is that the members of that site are actually creating the learning materials. They're drawing their own perspectives, their own experiences and they are contributing to that site.

There are many ways of producing content, and particularly with the new technologies we are able to more and more produce very good multi-media content. Not just text content, like blogs and blogging, although that's very important.

I guess I should mention Twitter because people were talking about Twitter, which is very short textual content. But also podcasting, as I'm doing with my audio recording here. Vodcasting, which she's doing with her video recording there. And game mods, game modifications and other multimedia are now very easily in the hands of people. People are able to create very simply complex multimedia.

I'm playing right now if you look at my web site; I'm playing right now with something called Kaltura. What Kaltura is is a system that allows me to create a video I've created either recorded off my own video – I have a little video camera in my computer - or I can upload some clips and grab them from YouTube or whatever. So, I create my video and I can add voiceovers, whatever to it. I put it on my web site

and then the next person who comes along can edit my video.

And so, there's a web site out there called Wiki Educator, which is using these Kaltura videos. What they're doing is for the different subjects they're having user generated, user created videos, educational videos, under different topics where people who come along after the video was created, add their own segments to the video. Amazing. Who would have thought that ten years ago? Oh, and of course, it's all free. Of course.

So we have Web 2.0, the learning network. The idea of this place, this weird place, this virtual place, that is an intersection between education and work and home, that allows us access to easy-to-use tools supported by hosting services, like Kaltura or Flickr or YouTube, and allows us to create, to create types of learning.

For example, the E-Portfolio is blogged kind of learning. The idea here is that you're learning is creating and presenting materials online.

I used to use the – still use it – slogan about “aggregate, remix, repurpose, feed-forward” as characteristics of the learning process. The aggregate is to bring in content, information from multiple sources. In other words, to reach out to all your connections online.

To remix is to bring different things from different sources together. To repurpose is to shape it to your own needs, to your own learning context, and then very importantly to feed-forward, to distribute this new material either as a video or a multimedia or a blog or whatever to other people in your network.

This learning is unorganized, it's unmanaged. It doesn't have an 'outcome' or a 'purpose'. It is based on the flow, the communication, the content of the moment. It doesn't have presenters and receivers.

It's characterized by the “unconference”. I'm not sure if you've seen a whole lot about the unconference movement that has sprung up recently. The unconference movement is a bunch of people – like you, say – get together and there is no pre-defined agenda as to what will be talked about. Different people volunteer to talk about different things. They write what they're going to talk about on a notice board at the side of the room, and people decide to create their agenda at the time they're having their conference.

Sometimes people talk, sometimes people don't talk. The content and the – well the entire structure of the conference is decided by the participants importantly not by votes or anything like that, but by each participant doing what they feel is most helpful to them at that point in time.

Now what typically happens is they move into different clusters, one cluster talking about one thing, another cluster talking about another thing. The idea here is that they are producing the best possible learning for themselves that they could have produced at that point in time with those people.

It's messy. It's messy, as opposed to neatly pre-defined structure. Look at this conference at this time (nothing personal to the organizers). Here you are. You're all sitting here listening to me. Is this the absolute best use of your time that you could possibly have at this point in time? Be honest with yourself because you're not going to say that out loud are you?

Voice: No! [Laughter]

The situation here has been pre-defined. There was a conference organizer who, for reasons unknown, picked me, brought me here through freezing rain [Laughter] and now here I am talking to you, right. And they had the best intent, and of course they made the best decision that they could given the resources that they had at the time and they may have been pretty good, but is it the best. Could you collectively have done a better job with the resources that you have?

The proposition by the unconference movement is "yes," you working together to organize yourselves could come up with a better structure, better use of your time right now than a single or small group of people organizing this conference for you. That's the proposition.

Again, it's the idea of user-generated content, user-generated conferences, user-generated learning as a whole. Again it's based on flow. I just mentioned Twitter in passing. Twitter is flow defined. Twitter is 140 character messages. You're not going to have anything static and permanent with Twitter. You're going to get messages to the effect, "what are you doing right now?"

The idea here is that learning is to immerse yourself in this flow. Not to try to capture this flow. Not to try to control or hold this flow, but rather to learn how to adapt oneself to this flow.

Douglas Rushkoff even talked about the internet being like surfing and learning on the internet as being like learning how to surf. A person who surfs doesn't try to come to one single, concrete understanding of the wave because there is no such thing. A person who learns how to surf learns how to adapt him or herself to the wave. Drawing on all their experience, all their perceptions to make minute adjustments, to be able to react to the wave as it changes, as it forms.

So we have Web 2.0. That's the concept and I just want to look briefly at some of the core technologies that underlie this. So, I ask how much time do we have. Oh, there we go.

Voice: Fifteen minutes.

Fifteen; good. See, I have a watch here, but it's ten minutes fast. So I don't want to trust myself. Time is – there's McTaggart, , he wrote a paper called the 'The Unreality of Time', I believe him.

The first underlying tools of Web 2.0 technology are social networking tools, tools to create these

learning networks. Social networking tools are very simply defined as tools that allow you to create lists of your friends, connect to them, and then use those lists to interact with them.

The earliest social networking tool, like instant messaging, ICQ, where you have a list of buddies, Skype you have a list of contacts. Now we have sites like Friendster and Orkut and Facebook and all of the rest where you very explicitly draw out a list of friends.

These social networking technologies, I'll just point out, are very much in flux right now. The future of learning is not in Facebook, no matter what you've read. The future of learning is not going to be any social networking site, properly so-called, but rather something called – something that is much more distributed – Tim Berners-Lee called it and I quote, “the giant global graph” – GGG.

This is the idea that each person's contribution to this web wide social network is a stand-alone thing, not connected to any particular web site like Friendster or Orkut, and in connecting to other people anywhere they are on the web rather than other people who are in Orkut or in Friendster or anything like that.

That's what is being evolved, there's a whole movement on social data portability, and as well (and it'll come up a little bit later) OpenID, which makes that possible (I'll mention that a little bit later).

Another underlying technology - and again this speaks to the messiness of the uncontrolled nature of Web 2.0 - is tagging. Tagging is a very simple phenomenon. A person presented with a resource like a photograph or a video or a web page or even an object, takes a word off the top of their head that they believe describes that object rather than selecting a word or a classification from a controlled vocabulary or a taxonomy. The idea of tagging is that people create their own vocabulary through use.

This is interesting because this creates organizations and structures of concepts that might not have been, perhaps could not have been, anticipated by taxonomists or librarians.

Another major technology underlying Web 2.0 is something called asynchronous Javascript and XML. The name sounds complex. It looks complex, but it's actually very simple.

When you submit information to a web site, the old Web 1.0 way, I know you've all done this, the typical way you do it, you type information into a form. Fill out your name, address, blood type, other maiden name, bank account number, the rest. Then you hit submit and the page reloads.

What AJAX does is it allows you to enter information in the same way, but instead of reloading the entire page, a Javascript using something called HTTP_REQUEST sends a message to a web server. The web server sends a message back to the Javascript and then the Javascript updates just that little bit of the page that is relevant.

So, for example, if you're logging in, here's your page. You have a little log in form you will out your name, password, hit submit. Instead of the whole page reloading, the little java script sends a message, gets a result back and then just changes the little box where you log in.

This is important because it allows web pages to become interactive. It allows you to create a single web page that can have multiple interactions with the web server.

This allows for the creation of online applications, such as Google Documents, for example, or the Zoho suite of applications or things like Gliffy.com, which is a web page that allows you to draw diagrams, flow charts and things like that. And there are perhaps hundreds and hundreds of these applications in the web page (linking to Web 2.0 apps).

One of the things that we're seeing is more and more people are using online applications like this instead of applications that would be loaded on their computer, like Microsoft Office Word document and PowerPoint slides.

One of the reasons for this is that it doesn't matter what computer you're on. You can always access your application and therefore you can always access your data. So if you're at a cyber café in Kuala Lumpur and you logged on, you can go to Google Docs and work on your document just as though you were on your own home computer.

The other thing that's important is, if the document is on a web site like that, you can work on that document at the same time that other people are working on that document. A classic example of course is the wiki. But right now I'm writing a paper with a friend of mine. We're working on Google Docs. I actually - was just last night in fact - I was sitting there typing on the paper and he's typing on the paper at the same time. We're both working on the same document at the same time and it's live. Brilliant; we love it. Got half the paper written like that. Unfortunately that was the easy half.

Another major aspect of Web 2.0 technologies is representational state transfer (REST), probably best described in contrast to something that you've probably seen a lot of buzz about, web services. Web services, which are supported by the simple object access protocol (SOAP) as diagramed at the bottom, are very structured, I consider them very top heavy. There's a lot of infrastructure. I think that they're slow. REST does the same sort of thing as web services, but it does it in a very lightweight sort of way. Data on the web and services on the web are accessed simply by sending a call to a web address, a URI.

The importance of this is that it gives online web sites, online services a very simple, low overhead way of rapidly sending data back and forth. This allows us to combine, to merge and combine data from multiple web sites.

That leads us to the concept of the mashup. The idea of the mashup is you're working with data from one application and data from another application, mashing it together to create a new application.

One of my current favorites is working with Flickr, which is a place where I store my photographs, because I have nowhere nearly enough disc space on my computer to store my photographs, especially now that I've got an eight megapixel camera. So I upload my photos to Flickr and then I access – I just click on 'organize my photos' and then 'map'. Map opens up a Google map, which Flickr has accessed from the Google site. Then I drag and drop my photo on to the Google map and it creates location data for my photo (note: in fact it's a Yahoo! Map. –ed).

So this match up here is allowing me to use two separate applications on separate web sites, Google Maps and Flickr, in order to create data that I couldn't have created otherwise. Very precise latitude and longitude locations for my photographs.

Now this is an aside: imagine that this is currently happening. Hundreds of thousands, maybe millions of people geo-tagging all of their Flickr photos. There are billions of Flickr photos. For any location on the planet you'll have photographs that have been taken by dozens and dozens and dozens of people.

Another major aspect is JSON. The only thing I'll say about JSON right now for those of you who worked on SCORM and sharable objects and faced the cross domain scripting problem, JSON solves the cross domain scripting problems.

JSON allows what is known as the 'tag hack'. Basically what it is is a way of importing into the head or the body of an HTML document, structured data in the form of Javascript arrays that could be used by Javascript on the current page.

So basically I've created the data on one place on my own web server, imported to JSON and use it on another web server thus solving the scripting problem. Brilliant, I love it; it's so simple.

Just as an aside again, JSON eventually is going to be a serious, serious challenger to XML.

Finally, OpenID, which I mentioned earlier: OpenID is a technology, a very simple technology that allows people to have a single identity to cross all of these different web sites.

You can see how having the same identity across all of these different web sites is going to leverage a lot of the interactivity, a lot of the mashups that the connections between these web sites enable.

This treats - and this is a bit of an aside - but this treats your identity as personal rather than institutional. I have no idea how that's going to play out in your context. I think your context is a bit special.

But for learning institutions, like universities and colleges, this is going to have a big impact because, historically the person, the individual, the student, has always been defined by the institution. You have

your university log in, your university number, your university ID. That's who you are.

But as learning happens more and more as a consequence of an interaction between multiple web sites, your identity has to persist across those web sites, which means it cannot be defined by any one of those web sites.

In practice what that means is it takes the definition of an identity out of the hands of the institution and puts it into the hands of the individual. You define your own identity, which you then project into the different web services.

So it means basically - again, your context might be different for the rest of the world - An end to walled gardens, a way of sharing social networks and content networks across institutions, across web sites.

Now I just want to – and I have like what, one minute? I just want to talk very briefly about networks and I'm going to zip through a slides here and not deliver them, but what this has to do with is how to design these structures, these networks, so that they're effective.

One of the things about networks is, just because it's a network doesn't mean it's good. Just because it's a network phenomenon doesn't mean it's good. Networks can have bad results as well. Sometimes they're known as 'cascade phenomena'.

If the disease spreads through a network of people, for example, that's a bad thing. Well, unless you're the disease, in which case it's a good thing. But I mean, for people, it's a bad thing.

So you need to set up your networks so that you're resistant to this. You need to set up your networks so that connectivity is possible, so that you can create this web and interactivity, but also in such a way that you minimize the risk of cascade phenomena and other things that will inhibit the function of the network.

So – I'm just going to skip through why networks and skip through distinction between groups and networks, which is a way of approaching this - and talk about networks as ecosystems and think about what makes a successful ecosystem.

Out of this - again, this is just my tack on this, right? - I've come up with what I call the 'semantic principle'. The semantic principle is the set of conditions that allows you to create reliable networks. There are four major principles to the semantic principle: diversity, autonomy, openness and interaction.

Diversity is the idea that the elements or the members of the networks represent the widest possible spectrum of points of view. You think about your brain, right? You don't want all of your neurons doing the same thing all the time in your brain. That would be really silly. Your brain would cease to function.

Your brain would be inert, like a lump of iron.

What you want are different neurons doing different things at different times. You want each neuron to have its own individual, unique set of connections, its own different perspective or point of view on the world. The idea of the network is you're collecting these many different perspectives and combining them to create a new, overall view of reality.

Autonomy - and this is tied very closely to diversity - in order to get this diversity, each of these individuals in the network needs to work autonomously.

By 'autonomously' I don't mean simply 'making their own decisions', but also making these decisions against their own background, against their own perspective, their own culture, their own world view.

So, we get a very genuine set of – how do I want to say this – very genuine set of distinct points of view or perspectives on any given state of affairs or entity.

Openness – we need openness in order to ensure that all of these different perspectives are heard and that no perspective is omitted from this overall collection.

Openness allows all perspectives from all points of view to contribute to the network, to create what Rudolph Carnap used to call the 'requirement for total possible evidence'. The Carnap requirement was the requirement for the total evidence.

Then finally, connectivity. This is kind of a tricky principle, but the idea here is that the knowledge in the networks is not the propagation of some knowledge from one individual to another, but rather the knowledge that is creating the network is created by the interaction of the networks as a whole.

It is not simply a reflection of the knowledge contained in any particular individual. That's a very tricky concept, especially when explained in less than five seconds.

But if you think about it you're looking at a television set and you see a picture of Richard Nixon, what makes the picture of a picture of Richard Nixon has nothing to do with any individual pixel. There's no pixel that's a little tiny representation of Richard Nixon that is blown up big. Each pixel is just a little black and white dot. Richard Nixon exists in that set of pixels only as a result of all of those pixels working together.

Similarly another example, flying an airplane from London to Canada. No single individual can do this. No single individual can build the airplane, make the tires, pump the gas, navigate the airplane, take off the airplane, transport the airplane across the ocean, land the airplane without crashing. It's too much work for any given individual to do. That knowledge is distributed across a set of individuals.

Those four principles, those four semantic principles, are intended to provide, if you will, a framework or a metric for the evaluation to design and select Web 2.0 and Key Learning 2.0 technologies. With that more or less on time, I conclude my thoughts and I thank you for your patience and your interest.

Cornwall, February 18, 2008

Nine Rules for Good Technology

Published as Nine Rules for Good Technology in The Technology Source July/August 2000 online July 1, 2000. Also published as Mmm, Good technology in University Business September 1, 2000.

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.

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.

Edmonton, March 6, 2000.

What Not To Build

I get to play a government scientist on the internet. As a result, a large part of my work involves being exposed to new and interesting technologies, whether they are the latest military simulators, academic papers delivered at scientific conferences, or product proposals being promoted by aspiring developers.

My sort of environmental scan is a bit different from what you'll get from consultants and venture capitalists. Don't ask me what companies are developing what products, how industry stocks are performing, or where all the 'smart money' is going. I don't know and I don't care.

What I can tell you, though, is what technologies are working, what technologies are flopping, and what technologies are fads. It's practical, down-to-earth advice. For example, if you are a technology developer, you already know that you should not try to build a new operating system, a new word processor, an online store or an auction site, for example. These have been built and have established a mainstream presence. You would need thousands of engineers and billions of dollars to compete with them.

The rest of my advice is like that, only more nuanced. It's obvious to everyone that they should stay out of the operating system market, yet much less obvious that they should avoid building a new content management system. It's less obvious, because these things are harder to see, but it is none the less certain.

So, here is my advice on what not to build. Actually, it's a bit more than that: it's a list of what not to build, a list of some things that people are working on now, some fads to avoid, and some indication of what's out there for the taking, if you can get your act together in a hurry. And what lies beyond that? The domain of real innovation and progress.

What Not To Build

Don't build a destination website

People are still building destination websites. They expect to build *the* location to find such and such. It's somewhat surprising to see in 2009, given that every company, every school, every library, every museum, and every other organization, product, service and even many pets have websites. Even if you have an original website idea, your site, unless it is *very* special (like, say, the Dead Sea Scrolls) will quickly be swamped by the noise and verbiage that is the web, your only traffic search engines and spammers. Even if you have original content and original ideas, don't just build a website.

Don't build a CMS

Unless you have an established market of community and content web sites, you have no business building a content management system (or, for that matter, a learning management system). There is a wide range of choices for people out there, everything from Drupal to Blackboard to SiteScape. And people looking for hosted content can use Blogger, WordPress or LiveJournal. And even more to the point, basically every large scale operation that is going to want a content management system already has one. You will be facing tremendous competition as every new and existing client will be choosing from a range of well-funded commercial and open source products.

Don't build a platform-specific app

2009 is likely to be a year in which everyone is building Facebook apps, Flickr apps, Twitter apps, iPhone apps and Second Life apps. But this is a market you want to avoid. For one thing, it is already saturated. Indeed, any time something becomes popular these days, it is designated as 'a platform' (in homage to web 2.0) and a horde of app-builders descend upon it. The platform remains popular for a while, but as it declines (as it inevitably does) it takes the entire set of platform-specific applications with it. And you risk, at any moment, the platform proprietor building competition for your app and putting you out of business (this applies to Android too, in case you're wondering).

Don't build a Java application

This is a bit of a special case of the preceding recommendation. Java is the original 'platform' application. What that meant was that you had to have Java installed (and, as time went by, the *right* version of Java installed). Java has been around for ages now, and yet most computer users could count on one hand (or, in many cases, zero hands) the number of Java applications they use. The situation is a bit better now that Java is built into some operating systems. But java's day has come and gone - everybody is into the platform-building game now, and most have learned from Java's mistakes.

Don't build a framework

This is one of those bits of advice that would not apply if you could actually do it - that is to say, if you *could* build a framework, then it may be worth doing, but for the most part, you probably can't. It's advice the Perl Parrot and Radaku projects should probably have heeded ages ago, advice the Ruby on Rails people should keep in mind today. People involved in those perpetually running framework projects are tossing good money after bad. Basically, if you are working in any well-known computer system - Microsoft, Java, Javascript, Python, whatever - a number of frameworks already exist. Javascript, for example, supports a number of frameworks for doing web 2.0 stuff - JQuery, etc. Now - you may say, the framework doesn't do everything we want. Maybe not. But that's not the lesson. The lesson is, if a framework already exists in your domain, your domain has been commoditized. Get out, get out now.

Don't build an educational game

This bit of advice is pretty specific and probably does not apply to most people (since most people would not dream of doing this in the first place). But the question to ask yourself is, what is a game doing for you that a straight-forward presentation of the information is not? If it is specifically an *educational* game, the answer is, "nothing." You're not getting new users, you're not presenting material in any way that's easier to understand, you're not adding to motivation. You're simply disguising the old 'teach and test' methodology as a game. Nobody will be fooled well, except maybe purchasers of fad educational products.

Don't build a new standard

People are still proposing to develop, or work on, new standards, be they metadata languages, vocabularies, application profiles, and the like. back in the days when no standards existed, this may have been a good idea. But today, the standards landscape is full. There are standards for every domain under the sun. Things that probably should not have standards - like carrier pigeon messages - have standards. What's worse, few of these standards projects made any effort to work with or cooperate with existing standards. So the standards landscape is a mish-mash of convoluted over-engineered and competing standards. Unless you absolutely have to, don't add to this landscape. Work with what's there and extend it (even if the rules say you can't).

Don't build a new social network

First we had several dozen social networking sites, like Friendster and Orkut and MySpace and Facebook. These became platforms (see above) and then we had social network multiplier sites, like Ning. And now (so-called) social network websites are multiplying like, well, websites. These social network sites are nothing more than reworked mailing list sites (like Yahoo Groups and Google Groups) and content management sites. And the blog-based social networking sites, like MyBlogLog, have already been commoditized. The irony is, as the number of these social network sites increases, their usefulness decreases. How many people are now refusing invitations from new social networks? Right - that would be everybody.

Don't build a wiki

This is a special case of the social network site. A wiki requires a community of people to work together to provide a common base of content or services. In order for a wiki to work, the contributors have to massively outnumber the spammers and the griefers. This works well if (a) the site is sufficiently massive, like Wikipedia, or (b) the site is sufficiently obscure. The Wikipedia project could be duplicated a few times before the pool of potential contributors is sufficiently diluted. That time has long since passed. Your wiki will be either (a) obscure, or (b) filled with spam.

Don't build a travel site

This is another special case. What it refers to is not the travel site specifically - though this market is saturated with the likes of Expedia - but the web services sites generally. The 'travel site' was always the paradigm example used to promote web services. But, just as it would be foolish to try to build another travel site, so also it would be foolish to try to build most web services applications. The point is, when you choreograph multiple applications, the market fills up very quickly. One travel site, for example, basically has a lock on hotels, airlines and car rentals. Web services sites are category killers, and most categories have already been killed.

What is there? Stuff everybody is working on

These are not things I would say you should avoid outright. The market is not saturated, there is room for innovation, and new products will be appearing over the next few months. But beware - a lot of people are already working on these things. If you have to start from scratch, you will have a lot of difficulty catching up. Your best bet right now is a niche play somewhere at the margins.

Alternative interfaces

Nintendo scored a huge hit last year with the alternative Wii interface and their success is drawing a lot of attention. People are now looking at all sorts of ways to control a computer game or computer interface. Webcam interfaces appeared a few years ago. Motion-sensitive and orientation interfaces are featured on things like the iPhone. I've seen gesture-based interfaces (with and without data gloves). I played a game with a heart-beat monitor last summer. And I've even seen a game based on a brain-wave detector.

Portability / cloud / smart cards

Cloud computing has attracted so much attention recently that it's a candidate for fad status. But behind the fad is a set of concepts that have legs - the idea of computational portability. By this I don't mean mobile devices (though obviously they play a role) but rather computers that can plug into other computers to allow you to move your data, software, authentication, and whatever else you want. We have smart cards in our credit cards now, but why can we have our web browser, email application, and social network in our smart cards? The answer is: we will.

Calendaring / coordination / events

There is a range of applications we might call Kantian applications - they depend on time and place. Historically, Apple and Microsoft have kept calendaring to their proprietary little selves, but this logjam is

breaking, and calendar-based applications are becoming available in our personal lives as well as our business lives. Which is good, because everything from social events to concerts to television listing to anniversaries depend on time. Finding new uses for time - that's an opportunity that will not go away any time soon.

Location-specific applications

The second group of Kantian applications are those that are taking advantage of publicly available GPS to create location-based services. These should become widely mainstream as mini-GPS systems are built into cameras, phones, PDAs, laptops, cars, belt buckles, keychains, and more. I personally could have used a GPS based keychain locator this month - I don't even know what city my keys are in. And keeping track of children, vehicles and pets will pass from quirk to mainstream over the next few years.

Intelligent apps / recommenders

We want in our everyday lives what is already available in some aspects of our professional lives - the ability to pick the best product or service in a given environment. Expedia, for example, allows me to pick hotels quite efficiently, and while it can be fooled by unscrupulous proprietors, the service is getting better over time. No such system exists - reliably - for consumer electronics, for rental accommodations, for cars, for food. Imagine, for example, a system that created my grocery shopping list for me, so I simply didn't need to figure out what I needed and wanted. Or that reliably recommended (and delivered, for free) books and music. Moreover, there is not only room for an extended range of recommenders, but there is also scope for increasingly reliable recommenders.

Connected applications (walls, desks, fridges, toasters)

OK, maybe not toasters, unless you really value weather maps on your morning toast. But with ambient wireless in an increasing number of homes it has become feasible to connect appliances to the internet. This creates a whole range of possible products - paper-thin displays that hang on walls, desks with smart, interactive surfaces, fridges that keep track of your food, automatic light switches that switch off when the room is empty, health monitors, and more. And it's not just that these applications are connected to the internet, it's that these applications can access *your* data, remember choices you've made, and interpreted and project your needs. What needs? That is where the room is for innovation.

Sensor networks and sensor data processing

A lot of work is being done in the field of sensor networks these days. There is a number of obvious emergency-related applications: fire sensing, flood sensing, intruder detection. Weather reporting should evolve in short order from a small number of central weather sites to a dense grid of home and business weather stations - and these, in turn, by making their data public will allow businesses to better manage staff, stock supplies and anticipate markets. Sensors already manage the flow of traffic in cities,

and will increasingly manage the flow of goods and people. Room for innovation here includes coming up with new things to measure and developing algorithms that analyze and understand large grids of related data.

Summarizing, data extraction, decision support / workflow support

Business intelligence services already monitor and analyze web and internet traffic in support of corporate and military intelligence. But there is room for personal intelligence services - wouldn't it be nice to know about that patch for your Zune, for example, *before* it suddenly freezes? And there is a need for people to be able to make sense of an increasingly diverse information space - especially as the traditional media can no longer be trusted (if it ever could) to describe events fairly and faithfully, or to report on obscure or unpopular disciplines.

Predictive data visualization

Data visualization will become even more useful when it becomes predictive. We already have a sense of this: we use predictive visualization every day in order to understand what the weather will be like (and despite widespread criticism our weather predictions are surprisingly accurate). Being able to predict crowds, shopping trends, stock prices, fads and fashions, and more, will become an increasingly lucrative industry. Imagine how Air Canada could respond if it had a reliable way to visualize the mess that would develop as holiday travel merged with a series of blizzards.

Fads

Green computing

Green appliances have been identified for a number of years with an "energy star" designation, a system that worked mostly because there was little business advantage to proclaiming oneself green. That has all changed. Consequently, titles and labels will be of little value as being "green" becomes a marketing ploy rather than an indicator of energy conservation.

iPhone

The key sign that the iPhone is a fad is the fact that most of the attention being paid to it has to do with applications and games, not telephony. In addition, the market for iPhones saturated itself within a few months of its initial release: pretty much everyone who wants an iPhone has one. Finally, other vendors - and in particular, Research In Motion, which has survived American patent protectionism - are matching (and sometimes exceeding) Apple point for point with product and service.

Cloud Computing

The interesting thing about cloud computing is that almost nobody in the public knows what it is. This makes it ripe for fad status. But to survive, cloud computing will have to actually be used - and people who don't know what it is won't be using it.

Maybe

Online instruction system

It has always been the holy grail of the e-learning industry: a totally automated system that manages instruction for you. There will be no end to the number of people who say teachers are indispensable, but if the social function of teachers can be replaced by community, and the informational function by software, then a stand-alone online instruction system is possible. And that's what we're seeing people try to build, step by step, with learning objects, competences, and the rest.

Distributed systems

The idea here is to have a thing - a concept, an idea - that rests on, and floats above, a non-specific computing environment. This was the thinking behind the connectivism course we ran last fall. The idea is that the 'course', via its constituent teachers and students, simply grasps whatever computing environment is convenient and available, creating communications channels between those environments, and hence establishing a virtual presence above those environments. Most human organizations can exist in this way, and become much more robust and flexible when not tied to a specific system.

Out there for the taking

By its very nature, most genuine innovation can't be predicted. But there are some obvious targets out there for the taking - extant problems which, if solved, would revolutionize the marketplace.

Marketing that works

To be clear: marketing works already. That's why vendors pay millions of dollars to television channels and radio stations. But as these media shrink, and as marketing money becomes more scarce (the demand from R&D is ever increasing) vendors are looking for a new definition of 'works': marketing that is welcomed, even requested, by potential customers, marketing that is not wasted on people who will not buy, marketing that is viewed as positive and helpful, not vile, crass and commercial. Product placement (been thinking of Cherry Chapstick lately?) is the new nirvana, but is still hit and miss.

Intelligent radio/television (live conversation / events)

Television and radio had show a surprising resilience in the face of the internet onslaught, and the reason for this is that they're easy. Turn it on and it will entertain you without pause until you turn it off (or until it shuts down for the night, an oddly archaic practice that still exists). Shows featuring online content - such as CNN's replaying of YouTube videos - have been, well, awful. But there's so much out there, and so much we could do for ourselves. If we had internet-enabled television or radio that programmed itself, that was personalized, that let us interact with it (in a meaningful way), imagine the future. Talk radio, for example, that is a conversation with people around the world that *you* find interesting.

Personal presence / personal health / personal learning

Personal health records, personal learning environments, personal publishing and printing, personal presence: all of these are ways of imposing the personal on the technical, about making these tools about **you** instead of about them. This is essentially a combination of technologies - of smart cards and their mobile ilk, of content analysis and presentation, of connected applications, of distributed systems. From the point of view of the internet, 'you' are a concept - the one thing in the whole system that isn't actually a part of the system. How to leverage that will be the stuff of genius and innovation.

Simulation / immersion

It's easy to do simulation and immersion if you have a lot of money. I have been in flight simulators that are absolutely convincing. But they cost millions to build; nobody is likely to have one in their living room any time soon. But in the field of affordable immersive simulation is a wealth of opportunity - imagine being able to experience 3D environments from the **inside** (and not just viewed through a screen) without leaving the house. We had reading rooms and TV rooms in the past: the device that creates the Sim room of the future will make somebody rich.

E-government

The problem with representative democracy is that your representative is often looking out for someone else's interests. Often his own. Internet technology creates the possibility for direct democracy, but this in turn requires a way of rethinking how we manage society. We want to connect people to government - but only those parts that affect them directly. We want to create mechanisms that allow people to govern themselves - but not to govern others. Collaborative and community-oriented systems for resource management and decision-making will be a fertile field in the future.

Energy nets

As energy becomes increasingly scarce, we will look not only to alternatives in power generation, but also better ways to manage transmission. We already have an energy grid, but as events have shown, the grid is unstable and liable to cascade failures. It also depends unreasonably on a small number of very large power sources, such as coal or nuclear powered generating plants, hydro dams, and the like. We want to be able to manage energy nets of the future using distributed sources - wind and solar powered, for example. Such systems would be tasked to minimize transmission load, insulate against cascades, and promote diversity of sources.

Dead tech

Telephone

By the 'telephone' I don't mean voice-to-voice communications generally - people will always want to talk to each other - but about telephony in particular (and in particular, dedicated lines and switched services). The reason is simple: it is simply too much overhead to maintain an entire infrastructure premised on the possibility that any given person may require a direct audio link to any other given person. We want to use the wires for other things. And the overhead required to support switching is immense and expensive. The closest thing we'll see in the future is something like bandwidth guarantees for specific services .

Television

Television is so close to being over we can almost taste it. Once digital television comes into households (2009 in the U.S., 2012 in Canada), the previous monopoly owned by the cable companies will be broken. Televisions will no longer be tubes, they will no longer have channels (increasingly, we'll just program numeric selections), and we will no longer watch networks (increasingly, we will watch providers - Fox, Gawker, Google, CNN). Yes, we will continue to have background audio and video displays in our room - often more than one - but we will no longer be 'glued to the tube'.

Radio

Radio is rapidly ending its life as an electronic transmission medium. Today, it is almost as common to listen to 'radio' stations online (through iTunes, for example) as it is to pick up signals from the air. As broadband becomes ubiquitous, we will more and more frequently simply pick up ambient internet and stream audio - whatever that entails. Satellite radio was the last harraugh of a medium that depended on mass broadcasts.

Print / paper

Paper is a resource-intensive industry and will become more and more expensive over time. Already, we are seeing the shut-down of pulp mills in remote regions of Canada, ostensibly because they are "inefficient" but in reality because the market simply isn't there for their product. Bookstores are filling their shelves with trinkets, DVDs, toys and games. newspapers are losing subscribers in droves (mostly as they die off). The affordability of electronic combined with the wastefulness of paper makes this an easy prediction.

Transportation

It has for the last few decades been cheaper to transport goods around the planet than to manufacture them where resources and labour are more expensive. Similarly, it has been cheaper to transport people from their pleasant homes in the country (or pseudo-country) than to live and work in the same location. This all changes as transportation becomes increasingly expensive. We will live and work in closer communities again, which means that systems that support local self-sustainability will be in demand. Can we build apartments that are as comfortable as homes? Can we grow our own grapefruits and coffees? Can we design specialized production systems that do not depend on cheap labour?

Moncton, January 4, 2009

Ten Futures

Drawing on Richard MacManus's [10 Future Web Trends](#), this is a bit linear, but has the virtue of identifying *future* trends, not things that are around today.

1. The Pragmatic Web

Forget about the Semantic Web. Whether or not it ever gets built, you can be sure that we will be complaining about it. Because while the Semantic Web gives us meaning, it doesn't give us context. It will give us what we can get from an encyclopedia, but not what we can get from phoning up our best buddy.

The pragmatic web, by contrast, is all about context. Your tools know who you are, what you're doing, who you've been talking to, what you know, where you want to go, where you are now, and what the weather is like outside. You don't query them; they carry on an ongoing conversation with you. The pragmatic web is chock-full of information, but none of it is off-topic and none of it is beyond your understanding (and if you need to know more, it will teach you). The pragmatic web isn't just a web you access, read to and write to, it's a web that you *use* every day.

2. Global Intelligence

While from time to time our computers are going to appear pretty smart, some of them even smarter than we are, they will be dwarfed by the emerging global intelligence or world mind. This won't merely be the 'invisible hand' of the marketplace, this will be the whole body. And it won't be based on the mere one-dimensional system of valuations of things in terms of capital, it will be composed of multi-dimensional interactions of wide varieties of media, including all of what we call 'media' along with money, votes, population movements (aka traffic), utilities (power, water, gas, oil) and resources (minerals, food) and more.

The global mind will to a large degree be inscrutable. We won't know what it is trying to do, what it wants, what it thinks are 'good' and 'bad', or whether it is even sane and balanced. That won't stop a slew of populists from claiming to 'know' where the global mind is headed (a la evangelists or Marxists) - though of course, except at a very macro level, the destiny of an individual is independent of the destiny of the global mind. The global mind is the sort of thing that raises questions about the meaning of life, the value of ethics, and the nature of knowledge. Our answers to these questions over the next few decades - even as global climate change and wars and natural disasters ravage our populations - will shape the course of society through the next centuries.

3. Extended Reality

We think of 'reality' as being constituted of the physical world and then of 'virtual reality' as being the digital world, or as we sometimes say, 'virtual worlds'. The two worlds are very different in that, well, one world is real and the other is not.

'Extended reality' is a digital version of the real world such that the digital version is as real as the real version. What that means, pragmatically speaking, is that if it hurts in the extended world, it hurts. We will have full sensory coupling with the virtual world, making the virtual world every bit as 'real' to us as the real world.

This reality will not just be a simulation of 'reality'. Rather, what will emerge as the combination of the two is a kind of 'hyper-reality', where objects exist both in the physical world and the digital world (think 'Spinoza' rather than 'Descartes'). The physical world and the virtual world will act as one; eat in the 'virtual' world and your body (such as it is) in the 'real' world will be nourished.

How could this ever happen? Well, take something like, say, 'money'. Is it real, or is it virtual? If you spend money, do you give the other person something real or something virtual? 'Money' is a perfect example of something that can exist in both realms. That's what makes it such a powerful force in today's society! But if money - which, when you think about it, was tangible, solid *gold* and therefore the *last* thing you would think could become virtual - then what else? Food, say?

We will live in an age of biochemical manipulation. Yesterday, we could create synthetic virtual worlds biochemically with drugs - 'take a trip and never leave the farm'. With sufficient computational power, we can create these worlds directly through interaction with computer systems. But we can also - by manipulation of matter electronically - create the 'fuel' that makes continued presence of the body possible. Doing something in the 'virtual world' has real-time direct biochemical consequences, some of which are constituents of energy inputs, which are converted to 'food' - or at least, the biochemical consequences of food.

4. Mobility

We will again in the future become a species of nomads, moving in tribes and herds through society, grazing on energy and information inputs as they become available.

This will happen as a result of a convergence of two factors. First, we will no longer be in want. At a certain point in time, sooner than we think, the technologies we have put in place to ensure the continued uneven distribution of resources (which we then use to extort labour out of deprived populations) will become moot. It will not be possible to maintain wealth technologically; there will be no 'means of production' unique to a certain privileged class of people. Hence, we will not need to hoard food and other possessions; we can simply take what we need from the ambient environment.

Secondly, we will by then be in the habit of needing much less. Consumer goods - ubiquitous today - will become expensive and impractical in the future. Owning a library of books, for example, will be a "wealthy man's folly" - a lot like keeping a Spanish Galleon in the back yard to support your own personal trade link to China. We will have few possessions, and those mostly as keepsakes or mementos. 'Rooted' people will be thought of in the future the way we think of 'nomadic' people today - unable, for some social-cultural reason to mesh with the rest of society.

5. The Human Grid

Human minds will continue to be efficient and effective processing systems, able to assimilate megabytes of information in seconds, intuitively recognize patterns, make decisions, and communicate ideas. Consequently, human contributions to the 'economy' (ie., the system of production of material goods for the sustainment of life) will consist primarily of providing mental 'inputs' to the production engines that actually do the work (much the way we 'drive' tractors today, but at a much more complex level).

Consequently, organizations will be able to derive value by enabling human minds to cooperate in the coordination or operation of elements of production. By contributing our thoughts and opinions on everything from celebrities to the weather to tomorrow's sports scores, computational systems will be able to derive the algorithms that will process iron ore, grow grain crops, and harvest energy from the wind and the Sun. It will be understood by these programmers that pop culture is a *metaphor* for the instruments of production, and that therefore human cognitive capacity can be mined directly by tracking thoughts and opinions about popular phenomena.

The collection of these thoughts and opinions from a network of people, all interacting with each other in an environment that includes entertainment, sports and other pastimes that engage the mind will be called the 'human grid'.

6. Smart Objects

This is discussed in Bruce Sterling's *Distraction* a bit, where he describes a hotel that instructs its owners on how it should be built. Objects - even everyday objects - will have a built in capacity for at least a primitive level of intelligence.

More importantly, these objects will be *connected* with other objects. We don't expect a lot of intelligence from strawberry jam, for example, but we expect it to at least know about what types of bread and peanut butter there are in the house (ie., your current mobile dwelling), to be able to monitor its compliance with your physical systems, to be able to suggest itself as a solution to current needs, to be able to offer relevant instruction, or to at least provide some input to the overall ambient room's conversation with you.

Your use of a product - whether it be strawberry jam, a fishing rod, or an auto-gyro, will have an impact on a whole network of other human and non-human systems. Taking the vehicle out for a spin, for example, will prompt a host of services to prepare themselves for your eventual arrival (and, indeed, you might not be going back). When you land - wherever that happens to be - your personal needs will already be in place (including any artifacts that you may have left behind). Consume a bit of strawberry jam and the global production system will conspire to manufacture that much more (assuming, of course, that it believes you will live to consume it and will have the inclination to do so).

7. Holoselves

No person can be in two places at once, of course, but one's avatar can travel one place while you travel to another, so when it comes time for that meeting in Colorado, you just shift your sensory input matrix to the holoself sitting down at the desk in Denver. Time for a lunch-time walk, so you transfer to the next holoself, which has been waiting patiently (like a book on the shelf) for you to pick it up in the Amazon eco-reserve. In the evening (after a holoself meeting in Zurich) you settle in with your 'real' self in Cairo for a nice evening meal and a show at the Pyramids.

Holoselves are, for all intents and purposes, artificial humans - you'd be hard pressed to tell the difference, and when they're legally occupied by a human, have all the rights and responsibilities of a human. People will naturally prefer to own their own dedicated holoselves, but it will be possible to share holoselves (the physical structure adapts to suit the host intelligence). Actual cognition (sensation, reflection, and the like) takes place partially in the 'real' brain, partially in the 'holo' brain (after a certain point the distinction between 'real' and 'holo' brains becomes more philosophical than practical - asking "Am I the same person in Cairo as I am in Denver" is pretty much the same as asking "Am I the same person tomorrow as I am today?"

The neat thing about holoselves is that they need not be human; they need just enough resident intelligence to input and process (coherently) perceptions and to communicate with other (holo and non-holo) instances of the controlling intelligences. This will lead to numerous holo-fads, like holo-birds, holo-fish, and more.

8. Living Art

When sentential utterances (words and sentences) are abandoned as a means of communication, it will become more natural to convey thoughts and information in multi-modal multi-sensory artifacts. We are beginning to see these even today with things like lolcats and YouTube videos. As our powers of expression (and the tools that help us) become more sophisticated, we will create complex multi-faced forms of expression, the most advanced of which will (almost?) qualify as 'life' and will most certainly qualify as 'art'.

Consider, just to gain an idea of this, how one wizard might express a thought to another in *Harry Potter*. Certainly the wizard would not write a note. Rather, the wizard would conjure an object of some sort - like a message owl, say. But the artifact will not 'carry' the message; the artifact will *embody* the message. On receipt of the 'message owl' the person would not merely read or be told, but rather, would interact with the owl - have a conversation with it - such that the subtleties and nuances of the message are expressed in a way that the recipient can understand them.

We think of communications today as means of carrying 'information'. This function will not cease in the future - we'll still need to say "My name is Johnny" or "I have an apple" to people in the future. But we'll say it in such a way that everything the recipient could want to know - the type of apple, the genetic history of the apple, the provenance of this particular apple, my preferences and opinions, stated and implied about apples, the current market value of apples - will also be contained in the message, not necessarily (and not typically) in sentences, but through a range of conventional multimedia iconology (kind of like giving somebody a white rose to say "let's be friends" and a red rose to say "let's be more").

We will, of course, also have 'living graffiti' - bits of badly created living art that clutter city streets and cling to walls - they'll have to be flushed with high-powered steam hoses into the organic recycling facility. And we'll still have spam - but at least when the message is delivered, you'll be able to eat it.

9. Global (Non-)Government

This is kind of an obvious one, but it should be clear that we will not have 'nations' in any geography-based sense of the term in the future.

This will become necessary due to the clamour of refugees trying to get to the highly developed regions of central Asia and Africa from their economically backward homes in North America and Europe. Many of these will be brought over by formerly American and European corporations, which will relocate to the centre of their major markets in India, the Congo and China.

In any case, the concept of government will have been radically redefined by that time in any case. Government will be no longer of geographical region but rather of sectors. We see a good example of this in its infancy in standards development. Standards are not managed by national governments; they are rather managed by councils with (interested) representatives from around the world.

More and more, sector councils will govern affairs. Fisheries, for example, having recovered from the panic of the early 2000s, will have been removed forever from national control. Energy production on the global grid will have followed. Many other industries - aviation, telecommunication, food production, finance - are already being governed in this way.

The big change will happen during the mass-democratization events that (I expect) will take place in the

middle of the 21st century. The sector councils will be badly managed by the corporate oligarchy that created them - they will act against the best interests of people (though it will take a disaster greater than Bhopal to demonstrate that to people) and will serve to preserve the privilege and wealth of a few. This, combined with the world wide 'free movement' movement - arguing that people, as well as capital and trade goods, should be able to move freely - will cause a crisis and an economic collapse. Governments will move militarily against corporations, which will agree to a power-sharing structure.

For the most part, after that, government will disappear from the lives of people. There won't be elections or anything like that; rather, people will participate directly in the management of sectors in which they are involved. Because people will have (what we today call) guaranteed incomes (but which amounts to free necessities of life) it will not be possible to coerce people in managerial hierarchies, and so corporat governance will be by networked decisions - each person will create creatively and 'pseudo-entities' composed of temporary collections of simultaneous inputs will achieve corporate outputs. That's how the first mission to Mars will be managed.

10. Cyborgs

This is a pretty easy one. The only thing preventing us from merging humans and machines today is that we cannot yet build machines at the scale and complexity required for human-machine interaction. Human inputs operate at the microscopic level, and require complex interactions. Even something so clumsy as replacing an organ requires that we grow - rather than make (though there are some few exceptions, like the artificial heart) - the organ, and then deal with interactions we couldn't design for with anti-rejection drugs.

But it should be evident that with biocomputing and nanotechnology we will be able to build, say, neural nets that can be installed alongside our existing cerebellum and can take over functionality as the original equipment wears out.

Most likely, the initial successes of cyborg technology will be in artificial perception. Replacing eyes, ears and other sense organs will succeed because base mechanical devices will be able to interface (much like a computer peripherals) with sensory input layers. Parts of these will also be created; we already have an artificial hippocampus.

There will, of course, be a large-scale industry in the psychology of cyborgs. Can a person be a ship and not become insane? How do we keep such a person occupied? Several of the technologies outlined above - like holoselves, for example, will be crucial. Metaphor will become reality - and it will become a major ethical issue - and a human right - to know one's actual situation.

Moncton, September 06, 2007