

Extreme thinking

By Michael A. Nielsen

School of ITEE and School of Physical Sciences, The University of Queensland,
Queensland 4072, Australia
nielsen@physics.uq.edu.au and www.qinfo.org/people/nielsen/index.html

First draft, completed August 30, 2003

Abstract

The natural sciences have a reputation for posing special challenges to the way we think and learn: they are a form of “extreme thinking”. In this essay physicist Michael Nielsen discusses some of the challenges facing researchers in the natural sciences, and how those challenges shed light on other tough learning situations.

About this essay

This essay is the text for a presentation delivered by the author at the “Tough Learning” conference held in Brisbane, Australia, September 7-10, 2003, organized by Learning Network Australia (www.lna.net.au).

Introduction

I’m a theoretical physicist. For many people, the word “physics” means something like: “the subject I disliked most or was worst at in high school”.

For me, physics is almost completely unrelated to all that stuff in high school - inclined planes, calculus, friction, and so on. When I talk about physics I’m talking about a human endeavour. It’s the endeavour to figure out what the basic rules governing our Universe are. What is the Universe made of? How did it start? How will it end? What’s out there?

When I was in my final year of undergraduate physics, I came across the following quote:

The quote was in an article by Leon Lederman, who won the Nobel Prize in Physics.

Look up and insert quote. Basic point was: Lederman had a very hard time early on in physics, achieved at a mediocre level, and it was a very

long time before he came up to speed, and began believing in his own ability to do physics. He simply decided to persevere, and pursue what he personally thought was interesting.

Reading this turned a light on in my head. You see, I was terribly worried that I didn't have the ability necessary to make a contribution to physics. You feel pretty stupid, pretty often. When trying to understand somebody else's work, it can often take hours of work to understand just a paragraph or two of text, and sometimes even then one does not understand the text.

Whatsmore, reading someone else's work is not doing physics. When I say "doing physics", I mean trying to understand some natural phenomenon out there in the Universe, a phenomenon that quite possibly nobody else in the world understands.

Imagine you're living in the 1850s. How could you convince yourself that matter is made up of atoms? How would you figure out that the sun is made of Hydrogen and Helium?

In the modern day, physicists struggle with similar questions. In my own work I'm trying to figure out what are the ultimate limits to computing power. Not "what can we fit into next year's laptop", but "given the laws of the Universe, what are the fundamental limits on what computations can be done?"

Answering questions like this is tough, and it's tempting, after trying and failing for a long time, to conclude that one is simply not cut out for this.

Well, I showed the Lederman quote to a whole bunch of friends, virtually the entire final year undergraduate crowd in both physics and mathematics. Almost uniformly they had the same response I had. Many of them asked for copies of the article.

You see, physics is extremely difficult *even for physicists*.

So my talk will be from the point of view of a practitioner constantly engaged in tough learning. It will be anecdotal, based on my experience and the experience of others, not based on well-grounded scientific studies. It will only incidentally be about my efforts to help other people learn, because even ten years after reading the Lederman quote, I still find physics plenty tough enough for me.

Alan Kay, the father of much of what's good about the modern personal computer, is fond of saying that "A change of perception is worth 80 IQ points." Well, doing physics is very different, in some ways, from most other human activities. I hope that'll give me an unusual perspective, one that will inform people engaged in other types of tough learning.

In preparing this seminar, I tried to imagine that I was writing a letter to myself 10 years ago, about the same time I was reading the Lederman article. What would I say? I decided that I'd try to identify a few fundamental principles critical to success in tough learning. So today, I'm going to talk about three principles that I believe are critical to success in any tough learning situation, not just physics.

First principle

The first principle is that effective learning requires a strong sense of purpose and meaning.

Let me go back to the example of Lederman. If you've got a bunch of young adults sitting around, trying to learn a subject, but feeling as though it's hopeless, as though they can never do it, how likely are they to be successful in learning that subject?

To give another example, consider Maslow's hierarchy of needs. Learning, intellectual development, comes pretty high up that hierarchy. If one of your more basic needs is not being met, then you probably won't have much luck developing your mind. Your purpose, your meaning, will be all tied up with getting food, or finding shelter, not in learning to spell.

In physics research, meaning and purpose is derived, to a considerable extent, from having *unique abilities* that enable one to contribute in *unique ways*. Research is, by definition, the process of discovering things that no one else has discovered. By developing unique skills and talents, a physicist is able to make contributions to the research community that nobody else could make.

How can we develop unique abilities? It's not likely to come from the production line, from the standard schooling and undergraduate education. The focus of most schooling is to produce people with a standard set of skills. That set of skills gives them a common understanding so they can take their place in a community, like the community of medical doctors, mechanical engineers, primary school teachers, physicists, or whatever.

Developing such skills in common with others is certainly critical for the existence of such communities.

But a sense of individuality, of meaning and purpose, is not going to come from an education in which you are treated as one of a large number of identical automatons, all sopping up the same sets of skills. By institutionalizing learning in this way, we *deprive* people of a sense of having a unique and therefore meaningful contribution to make.

I believe that the key to developing a strong sense of purpose and meaning is to balance three activities.

- The first is development of a common understanding with a large group of people, people with whom one is later able to feel a common sense of community. This is the focus of most educational institutions, from kindergarten through undergraduate degrees.
- The second is development of abilities which are not common to your community, and which eventually give you the ability to make a unique contribution to your community.
- The third is making a contribution to your community. For most people, meaning is derived, at least in part, from a sense that one is making a contribution to something larger than oneself. If one focuses solely on self-development, then one quickly loses any sense of such larger meaning. For this reason, it is important that self-development and practical achievement be balanced. Ideally, one can be used to reinforce the other: we learn a new skill *en route* to some concrete achievement. Or perhaps we can teach others as we learn a new skill, making a concrete contribution to their welfare as we ourselves develop the skill.

Second Principle

Have you ever seen a bunch of seven year old kids playing soccer?

They run around in a herd, with the ball somewhere in the middle, going from place to place.

Every once in a while a kid breaks out who is faster than all the others. He takes off with the ball, running in front of the herd, before scoring a goal.

Is this kid going to grow up to be a great soccer player?

Well, maybe he will, and maybe he won't.

One thing we know for sure is that if he keeps relying on that great pace, and doesn't develop the other skills needed to play soccer well, he's not going to turn into much of a soccer player.

That is, if he concentrates solely on that which brings him short-term success, namely, his great pace, while leaving undeveloped other abilities which are essential, but don't have the same short-term pay-off, then in a few years time he's not going to be much of a soccer player.

The second principle is that effective learning requires long-term vision. A powerful long-term vision can give you the courage and will to do things important things for your self-development, but that don't pay off over the short term, and that may even be discouraged by your organization, by your peers, or by your superiors.

You may laugh at the example of kids playing soccer. Let me give you another.

In Australia, basic research is funded through a body called the Australian Research Council, or ARC. Each year the ARC gives out hundreds of millions of dollars in grants to Australian University academics.

At the end of 2001, the ARC announced four "Priority Areas", areas of particular importance for Australia's future. Those priority areas included things like nanotechnology and biotechnology, for example.

For the 2002 annual funding round, the ARC announced that fully one third of their funding, across all disciplines, including science, the arts, and engineering, would be allocated to the priority areas.

Now, in the previous year, 2001, I've been told that only about ten percent of grant applications were in what would become the priority areas. But after the ARC announced that one third of their funding would go to the priority areas, close to fifty percent of the applications in the next round were in those areas.

That's a five-fold increase in the number of applications.

Wow.

It's possible, of course, that half the people doing research in Australia had secret burning desires to do research in those priority areas. What seems much more likely, at least to me, is that people saw their own short-term interests best being served by changing their so-called research interests in this way.

If we *lack* a long-term vision of what we'd like to achieve, our behaviour will be largely determined by the perceived short-term rewards.

Now, in some instances, that's not a bad thing. In an organization and society that functions well, the short-term rewards will encourage good behaviour. But it is an extremely rare institution which has incentives in place that result in behaviour anywhere *near* as effective as if we each had our own well thought out long-term visions.

This is even true at the everyday level. When someone comes by my office, and says "Michael, would you mind coming to such and such a meeting?" I'm very tempted to say "yes", simply for the warm fuzzy feeling -a short-term reward -

that doing someone such a favour will give me. Not because I've truly weighed up the costs and benefits, to myself and to others, of doing it, but simply because I like that little bit of pleasure one obtains from basking in the approval of others.

The American Philosopher Ralph Waldo Emerson said it well in his essay "Self-Reliance":

What I must do is all that concerns me, not what the people think. This rule, equally arduous in actual and in intellectual life, may serve for the whole distinction between greatness and meanness. It is the harder because you will always find those who think they know what is your duty better than you know it. It is easy in the world to live after the world's opinion; it is easy in solitude to live after our own; but the great man is he who in the midst of the crowd keeps with perfect sweetness the independence of solitude.

The key to doing this is to develop a long-term vision so compelling and well-internalized, that it can override behaviours for which the short-term rewards are significant, but which may be damaging in the long run.

Let me give you an example. In Australian Universities writing papers is seen as such an unalloyed good that there are enormous incentives to write as many as possible. It is extremely tempting, looking at successful young Australian academics to conclude that one should write *lots and lots* of papers. But writing papers takes a lot of time. If you're doing that, then chances are that you're not doing much of anything else, especially if you're teaching, sitting on committees, writing grant applications, and so on. Will you be learning new fields? Will you be developing yourself in other ways? Will you be taking time to explore more generally?

A compelling and well internalized long-term vision will enable you to balance the writing of papers - which is both necessary for survival as an academic, and also a worthwhile contribution to the broader community - with other aspects of your research life, aspects that may be equally important, but which are not necessarily rewarded in the short term.

Comparing the first and second principles

It is interesting to compare the first and second principles.

The first principle was that effective learning requires purpose and meaning.

The second principle was that effective learning requires a long-term vision.

At first glance, these principles appear rather similar. Closer consideration shows that it is possible to follow either principle without obeying the other.

Suppose you have a strong sense of purpose and meaning - something you believe is deeply important, and that you are engaged in on an ongoing basis. It does not automatically follow that you will act in a way best suited to achieve that purpose or actualize that meaning, over the long-term. Unless you think through in detail, and regularly revisit, a long-term vision for actualizing that meaning in your life, the chances are good that you can be misled through short-term rewards, and the influence of other people.

Suppose, conversely, that you have a long-term vision. It is easy to come up with long-term visions that have little purpose or meaning to them. Sometimes people do this in disastrous ways. Someone's long-term vision may be to become a millionaire by 30. They may achieve that long-term vision, and realize when they get there that it *didn't matter to them*. Meanwhile, their other relationships and their health have deteriorated enormously, and they are asking *why* they did what they did.

This problem is also reflected in time management programs. Sure, it's possible to manage one's time in more effective ways. But unless one has a deep sense of purpose about what one is doing, you may simply be doing unimportant things, faster.

This is why it's important to establish meaning and purpose, as well as a long-term vision, in order to learn effectively.

Third Principle

The third principle is that the most effective way to learn something is to teach it.

A striking expression of this idea came in a seminar from management guru, Stephen Covey.

Covey has a large group of people in a room, split up into many small groups, each group seated around a small circular table. At some point in the presentation, Covey tells people that he is going to make five points over the next ten minutes, and that after he's done the person sitting at "6 O'Clock" at the table is going to explain the material to everybody else at the table.

Well, the first of Covey's five points is that the best way to learn is by teaching. He starts off making this point in the abstract. No doubt many in the audience are sitting there, nodding their heads in agreement. Then he looks around the audience, and asks people to compare the behaviour of the people sitting at 6 O'Clock to those not at 6 O'Clock. Invariably, the people at

6 O’Clock are sitting there, assiduously taking notes, paying close attention, and so on. They are, on average, paying much closer attention than those not seated at 6 O’Clock.

By changing the *role* of some of the people in that room - those seated at 6 O’Clock - Covey changed the way they perceived themselves, a change that was reflected in vastly more effective behaviour. They began to see themselves as teachers, and this made them much better learners.

Covey, by saying just a few words, has changed the way some of the people in that room look at themselves. Those people are then learning at a rate that is perhaps two or three times faster than the people around them. He’s achieved this simply by changing people’s role so they see themselves as teachers.

This principle is an expression of a deeper principle, namely, that the most effective way of changing someone’s behaviour is to change their social role.

This is an incredibly powerful principle, which we can use to change our own behaviour. We can *invent roles* for ourselves that reinforce behaviours we want.

One of the greatest exponents of this idea was the American writer, scientist, diplomat and statesman, Benjamin Franklin. Throughout his long life, Franklin was constantly inventing new organizations and institutions that reinforced those aspects of his own behaviour that he thought most desirable.

Perhaps most famously, at the age of 21 Franklin created the “Junto”, a small group of men who formed a society dedicated to their own self-improvement. At each meeting of the Junto members were asked questions like “Have you lately heard of any person coming into wealth, and if so, how was that achieved?” or “Have you lately heard of any offense of the State against liberty?” (Check and revise.) The questions were chosen by Franklin specifically with the goal of promoting his own self-development, and the development of others.

Of course, Franklin could have addressed these questions himself each week, in his own time. Or he could have discussed it among friends on a regular basis. But how much more *powerful* it must have been to create an institution dedicated to addressing these questions on a regular basis!

In scientific research, one expression of this idea is that to learn something, one should teach it to somebody else. This can be done at many levels, from a simple conversation with somebody, through to a formalized classroom setting.

Interestingly, the formalized classroom setting is in some ways the easiest of these, since in that instance the process is institutionalized, which hides much

of the process. What one learns from this setting is also not particularly widely applicable, since most people don't work in classrooms.

Instead, I'll talk about a more informative example, with which I have some experience, the organization of small, informal groups that meet to learn some new skill or subject. The same basic process can be applied to the organization of almost any teaching situation, from a simple conversation on up.

The first step is to identify a suitable person or group of people - people who may have a common interest in learning about the subject you wish to learn, and thus to teach.

The second step is to negotiate with those people, individually. People don't like being lectured to against their will. You need to ask "Are you interested in learning about subject X from and with me?" This has to be done in a no pressure way, and in a manner which gives people escape options if they should decide to leave.

Once common interest is established, you need to negotiate two things.

First is the format the teaching will take.

Second is people's level of personal commitment. This can be discussed individually, but ideally the commitment will eventually be made in a public manner, perhaps through conversation with the entire group. Usually this entails to making some commitment of time *outside the group*, for extra reading, problem-solving, essay-writing, field work, and so on.

The third and final step in the process is to make sure everyone has agreed on some procedure that ensures accountability to their commitment. A very simple procedure I've seen used in discussion groups is for each member to (very briefly) publicly announce how they fulfilled their commitment since the previous meeting. If someone is regularly not meeting their commitment, this can be brought up by the leader, and the issue resolved, either by a renewed commitment, or by a changed involvement in the group, including termination.

Conclusion

To conclude, let me summarize the three basic principles I've described.

The first is that effective learning requires purpose and meaning. This purpose and meaning can, in many instances, be obtained by concentrated self-development focused on developing combinations of abilities that enable us to make a unique contribution.

The second is that effective learning requires a long-term vision. In all organizations the short-term incentives neglect critical aspects of our development. By developing and constantly reinventing a compelling long-term vision we are able to ensure that we develop and learn in the most effective ways possible, even when those ways are in contradiction to the established short-term incentives.

The third principle is that teaching is the most effective way to learn. More generally, if we want to change our own behaviour, one of the most effective ways of doing it is to either create, or to get involved in, social institutions that will promote the behaviour we're trying to encourage.

These three principles are obviously not all there is to tough learning. There are many other important principles and practices. But, if I could write that letter to myself of ten years ago, I think those are the principles I'd try to convey.

Let me conclude by mentioning that I have begun a small discussion group about the practice of doing research, involving several students and staff at the University of Queensland. As part of the discussion group, I maintain an online weblog recording the ideas we discuss, and providing a forum for discussion of those ideas. If you've found any of what I've said today interesting, please visit the weblog at:

<http://www.qinfo.org/people/nielsen/weblog/>