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

September 8, 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 A. Nielsen discusses some of the challenges facing researchers in the natural sciences, and how those challenges shed light on other tough learning situations. *(4800 words, viewgraphs and text online at www.qinfo.org/talks/)*

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

Good morning.

Thank you all for coming along to my talk today.

To some extent, I'm moonlighting today. Professionally, you see, I'm a theoretical physicist, a scientist. My interest in tough learning, though, is very personal, and goes back a long way.

Since my undergraduate days, right up through the present, I've occasionally considered quitting my job, and taking up something else. Not because I don't enjoy my job – quite the contrary. The reason is simply that I find the subject matter of science extremely challenging. I know that the same is true of many of my colleagues and students.

My talk today is therefore from the point of view of a practitioner constantly engaged in a battle with his subject matter.

My job is, in some sense, a type of tough learning. Of course, most jobs have special challenges that make them difficult. I'd like to begin by briefly conveying to you some of the flavour of the special challenges that make the natural sciences difficult.

I'm going to describe a game to you. I'll explain the rules of the game, and then ask you to respond to a few questions with some shows of hands. I'll explain the point of the game after we're done. This game is by far the most informative, by the way, if you just give your initial gut response when I ask for a reaction, and don't worry too much about whether your answer is the same as everybody else's, or whether your answer is "correct" or not.

The game is a card game. You are shown four cards. Each card has a letter on one side, and a number on the other side. You see only one side of each card; you see "A B 2 1" on the visible faces of the four cards.

Your task is to check to see if the following rule is obeyed by the cards.

The rule is simply this: if a card has a vowel on one side, then it has an even number on the other side.

Now, let me ask you, which of these four cards would you *need* to turn over to check that the rule is being obeyed? You may choose more than one card if you wish.

Now let me describe a second game.

Imagine you're working as a bouncer in a bar. There's recently been a crackdown on underage drinking.

As the bouncer you have to enforce the following rule: for a person to be drinking legally, they have to be 18 or older. Your boss has warned you to be on the lookout for people violating this rule.

Well, you've just started work for the night, and you see four people at the bar.

The first person is drinking a beer.

The second person is drinking a diet coke.

The third person is quite elderly, definitely over 18. You can't see whether they're drinking alcohol or not.

The fourth person is just a kid, definitely under 18. You can't see whether they're drinking alcohol or not.

My question to you is, as the bouncer, which of these people do you *need* to check out, to see whether the rule is being obeyed correctly? You may choose more than one.

There are two interesting things here. First, when given this test the overwhelming majority of people choose the wrong cards to check in the card game, and the right people to check in the bar game.

Second, the two games are actually the same games, but in disguise.

Let's go back and look at the card game, and see what the common mistake is. Overwhelmingly, people think you have to turn the first and the third cards over.

Now, it's correct that you have to turn the first card over. You have to check to see whether or not there is, in fact an even number on the other side.

But it's simply not true that you have to turn the third card over. It doesn't matter whether there is a consonant or a vowel on the other side. Either possibility is consistent with the stated rule.

But you do have to turn over the fourth card. If there's a consonant on the other side, then the rule is okay. But if there's a vowel on the other side, then the rule is broken.

What's especially interesting is that the four-card game and the bar-room game are actually exactly the same problem. If we identify a person's age – over or under 18 – with being a vowel or a consonant, and their drink – alcoholic or non-alcoholic – with being even or odd, then you can quickly convince yourself that the two situations are exactly the same.

This experiment has been performed numerous times around the world, usually with essentially the same outcome: people get the bar game right, and the card game wrong. One common conclusion is that while people are extremely good at reasoning about social situations involving other people, as the concepts involved become more abstract they find it much more difficult to reason. Their intuition begins to go haywire even in relatively simple situations like the card game. In slightly more complex situations it stops functioning altogether.

I'd like to relate this psychological experiment to the process of scientific research. When I say "doing science" or "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. I don't mean trying to learn something out of a textbook, like you do in school.

Imagine you're an early scientist. How could you convince yourself that all matter is made up of atoms? Just how would you go about doing that? Indeed, what would inspire you even to come up with the hypothesis that matter is made out of atoms to begin with?

In the modern day, scientists 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 computers can do?"

What are the key learning difficulties involved in this process?

To explain this, let me talk a bit about how our ideas about how the Universe functions have changed over the years. Thousands of years ago, our ancestors explained everything in very anthropocentric terms, using social metaphors in which there is a spirit or God for every rock, tree and animal.

A few hundred years ago, Newton, Descartes and other early scientists moved away from social metaphors to more abstract concepts, like gravitational force and inertia.

As time has gone by since then, our understanding of the Universe has improved still further, but we've had to introduce abstractions that are more and more foreign to our human intuition, and further away from the social ideas we understand so effortlessly.

A priori, there is no reason to suppose the laws of Nature are particularly accessible to the human mind. The mind developed as a way of furthering competition with our fellow living beings, and only incidentally to understand the laws of the Universe. It's good at recognizing faces, making tools, and understanding social situations. It's not good at solving mathematical equations or other abstract reasoning. In doing science we are borrowing equipment developed for another purpose, and we should not be surprised if sometimes the equipment is not cut out for the task.

Even to understand a relatively simple piece of modern technology, the laser, we need to pile up abstraction on top of abstraction, starting with pretty concrete things like atoms and electrons, before moving on to more ephemeral concepts like photons, coherence, stimulated emission, and population inversion.

The interesting thing is, and this is the point I want to stress, most scientists aren't much better than average at when it comes to intuitively working with such abstractions. We're stuck with the same equipment as everybody else. That, you see, is the special challenge that faces the scientist: using equipment

not cut out for understanding abstract concepts to understand a Universe whose fundamental laws are abstract concepts.

What scientists do master is a large variety of skills that let them, slowly and methodically, work with such thickets of abstraction. To read a single page of scientific work may typically take me an hour or two. The reason why is that I can't just effortlessly follow long lines of logic involving long chains of abstraction. Instead, I have various systematic ways of moving through it, and rather slowly and meticulously understanding the ideas.

I hope that description conveys to you some sense of what makes doing scientific research a tricky business. Scientists are constantly asking questions about how the Universe works. Even phrasing the questions, let alone answering them, involves tricky combinations of abstractions.

So my talk will be from the point of view of a practitioner who is 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 I still find science plenty tough enough for me.

In preparing this seminar, I tried to imagine that I was writing a letter to myself 10 years ago, when I was just beginning my career as a scientist. 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 as a scientist.

Interestingly, the principles I identify don't involve the abstract side of science at all. They are general principles about how to learn effectively. It's as though all the abstraction in science is merely an obstacle that makes doing it a little more difficult, but which one can cope with by learning appropriate skills. Once those skills are learned, the core problems of learning are much the same as in other areas of life, and it is those core problems that I'm focusing on for the remainder of the talk.

Everything I will say is common sense, all the fundamental ideas are ideas you will have heard before. The principles and ideas I describe certainly aren't new. But that doesn't mean these ideas are common practice. Hopefully this seminar will illuminate both the importance of the principles, and some of the difficulties involved in following them.

Obviously, in limited time there is much that must be omitted on a subject as large as tough learning. There is one important omission I'd like to at least mention. The three principles I describe are focused largely on individual actions, and ignore the larger social context – the norms and institutions of the

societies in which we live, and how those affect the applicability of the principles. This social context obviously has an enormous impact on tough learning, and deserves separate treatment.

Let me conclude this introduction with a note describing my hope for this talk. Alan Kay, arguably the father of modern personal computing, is fond of saying that "A change of perception is worth 80 IQ points." I am hopeful that my perspective as a scientist will be a little bit different from the usual perspective applied to the problems of tough learning. This is, of course, dangerous – perhaps some of the ideas here are naïve, or ignore important difficulties. But I hope that the unusual perspective will prove illuminating to at least a few of you.

First principle

Probably the single most intimidating and paralyzing thing for scientists beginning their research careers is a sense that they are not able to make a significant contribution.

Beginning research for the first time is incredibly intimidating. You are expected to say something new and original about how the world functions. How on Earth can you do that, when you're surrounded by people who know everything you know, and more?

Many people respond to this realization with a sort of frenetic paralysis. Often they move frantically around from topic to topic, trying to find something to which they can make a contribution, but always coming back to the fundamental problem, namely, that they feel inadequate to the task.

In my opinion, the single most important principle of effective learning is that it requires a strong sense of purpose and meaning.

When one has developed such a sense of purpose and meaning, it motivates you in all you do. Absent such a sense, it is extremely difficult to make a contribution.

Where does such a sense of purpose and meaning come from? For many people, it comes from being able to make *unique and significant contributions to something larger than themselves*. Whether it is a nurse helping elderly people in their last days, a writer completing a significant book, a parent raising a child, or a scientist solving a research problem, in each case the person is making a unique and important contribution to a larger community and cause than themselves.

In scientific research, and in many other activities, the way to achieve this is by developing *unique abilities* that enable one to contribute in *unique ways*. Research is, by definition, the process of discovering things no one else has discovered. By developing unique skills and talents, a scientist is able to make truly unique contributions to the research community.

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 which fits them into a community, such as the community of medical doctors, of mechanical engineers, of primary school teachers, of scientists, 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 automata, 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 creative contribution to your community, to something larger than yourself.

To develop effectively, we need to balance all three of these activities. Many people – I have certainly been guilty of this - concentrate on the first and the third of these activities. They find themselves frustrated in trying to achieve the third, simply because they haven't developed abilities that enable them to make a unique contribution.

The second activity is difficult for two reasons.

First, it requires one to act on one's own, to say "I choose to develop this ability, even though it makes me different from others in my community", and to follow through on that choice.

Second, even if one has the will to develop in this way, it can require a raft of self-development skills. These skills are simple enough in themselves, but many of us lack them. If we take the trouble to master these skills, then we will be able to develop abilities that give us the power to make a unique contribution to the larger community.

Second Principle

Let me move onto the second principle.

Have you ever seen people playing soccer?

If you watch adults play, they'll be nicely spread out across the field.

By contrast, if you watch a bunch of eight year olds playing, they run around in a tightly bunched 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.

The example of kids playing soccer is all very well. Let me give you one closer to adult experience.

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 heard – informally - 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, the same informal source told me that close to forty percent of the applications in the next year were in those areas.

That's a four-fold increase in the number of applications!

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 more likely, at least to me, is that many of those people saw their own short-term interests best being served by changing their so-called research interests in this way.

If we *lack* a clear and complete long-term vision of what we'd like to achieve, and how we'll achieve it, 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 also 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":

> *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 keeping this independence of solitude 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.

The difficulty in doing this is that the only reward for doing it is, in fact, long-term. In most organizations, nobody will come and pat you on the back for developing and then constantly revising a long-term vision, that is, a description of what your goals are, and how you are going to achieve them. That is something that you need to take responsibility for. It means developing the skills necessary to develop such a vision, and committing regular time to it.

You may ask what the difference between the first and second principles is. The first concerns purpose and meaning, while the second relates to having a long-term vision. In fact, it is possible to follow either principle without obeying the other.

The difference is that purpose and meaning is about *why* you do things, while long-term vision is about *what* you do, and *how* you do it. You can develop a strong sense of why you're going to do something – you might develop a compelling desire to win a gold medal in the 100 meters – but without knowing what you'll need to do, or how to do it, it won't happen. Similarly, you can decide what you'll need to do, and how, but unless that deep inner desire, purpose, meaning – the why – is present, it won't happen.

Incidentally, this distinction between the first and second principle is ignored in many time-management programs. Sure, it's possible to manage one's time in more effective ways. But unless you have a deep sense of purpose about what you are doing, you may simply be doing unimportant things, faster.

Third Principle

I'd like next to describe a striking moment in a seminar given by the self-improvement 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 assiduously taking notes, paying close attention, and so on, paying much closer attention on average than those not 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, changed the way some of the people in that room looked at themselves. Those people then learned at a rate perhaps two or three times faster than the people around them. He achieved this simply by changing people's social role so they see themselves as teachers.

This example is an expression of the third principle, namely, that the most effective way of changing your own behaviour is to change your social role, if necessary, by *creating social 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 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 the Junto members were asked questions like "Have you lately heard of any citizen's thriving well, and by what means?" or "Have you lately observed any encroachments on the just liberties of the people?" 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!

Suppose you want to get fit. What's the best way to do it? Many people join a gym. Did you know that gyms make more money in the first two months of the year than through the other ten combined? But most of those people never show up beyond February.

By contrast, imagine you join a running club, and concentrate in the beginning not just on running, but equally on making new friends, and becoming part of the social life of the club. Do you think it would be easier to get fit this way?

This principle is obvious, but it's surprising how few people take full advantage of it. People will complain loudly about the difficulty they have in changing some behaviour, or learning some desired new behaviour. Why then don't they create a social environment that will support and reinforce the change they desire?

I think the answer to this question is that such creation requires a whole raft of skills.

As a scientist, I've frequently wanted to learn some new sub-field of science. Doing this on my own is, I find, extraordinarily demanding. So, beginning a few years ago, I began inviting interested people to participate in informal courses that I'd teach about those subjects I wanted to learn. The theory was that this would force me to learn those subjects, and it'd be useful to other people as well.

The first few times I did this it was pretty dreadful. Not because I was especially inequipped to teach, or to learn. The difficulty was that I lacked the skills necessary to create the social environment. Remember, this was an informal group, not a University class or some other more formal group. We had issues with commitment, accountability, and social dynamics, issues that I was unprepared to deal with. It was several years before I learned the skill set necessary to effectively create and maintain such a social environment.

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 one of the most effective ways of changing deeply ingrained habits and establishing new behaviour is to either create a social environment that will promote the behaviour we're trying to learn.

These three principles are obviously not all there is to tough learning. There are many other important principles and skills. But, if I could write that letter to myself of ten years ago, 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/blog/

Thankyou all for your attention.

Acknowledgements

Thanks to Michael Bremner, Allan Quigley, Andrew White, and, especially, Jennifer Dodd, for many constructive comments that have greatly improved this essay.