

Learning in an open world

6 September 2007

Transcript of a keynote speech by **Peter Norvig, Director of Research at Google**, at the 2007 Association for Learning Technology Conference in Nottingham, England. In the chair, Mike Sharples, Professor of Learning Science at Nottingham University. Slides for this talk are at http://www.alt.ac.uk/docs/altc2007_peter_norvig_keynote.pdf [13.5 MB PDF]. **Slide transitions are indicated in square brackets.**

This text transcript is at http://www.alt.ac.uk/docs/altc2007_peter_norvig_keynote_transcript.pdf [75 kB PDF]. An MP3 recording of the talk is at http://www.alt.ac.uk/docs/altc2007_peter_norvig_keynote_audio.mp3 [23 MB MP3]. Slides and video of the talk, captured as an Elluminate Live! Session, is on the ALT-C 2007 web site at <http://www.alt.ac.uk/altc2007>.

[1] Thank you all for coming here. Now I'm a technologist, as Mike pointed out I've been a teacher and I've been a life long learner. But before this meeting I had no real contact with the official *Learning Technology* discipline. And so I was rather apprehensive when I got the invitation from Seb to come here and talk. I said 'gee, before I could do that I'd probably have to read ten or twenty years worth of research papers just to catch up,' to be able to talk to you guys, and that seemed kind of intimidating. But then I was very encouraged when I talked to a friend of mine, Hal Abelson, professor at MIT and educator, author of probably the world's best book on computer programming. [2] And he told me you only need to read one paper. So that was a big relief. I figured that was a reading assignment that I could handle, and so I was able to accept Seb's invitation.

[3] This is the one paper that Hal told me to read. It's a paper by Benjamin Bloom from 1984 on what he calls the *two sigma problem*. And you probably can't read the screen very well but what he points out is a comparison that he and his students had done of three different learning methods.

The first is what he calls *conventional teaching* [4] in which students learn the subject matter in a class with about thirty students per teacher and the teacher lectures and gives tests in the way that we're used to.

Then the paper compared that with a model of Bloom's called *mastery learning* [5], where the idea is you do the same kind of lecturing but instead of having tests to separate out the good students from the bad students you keep on testing until the students actually get things right. And so you test to identify what their errors are, then you go back re-teach them on that and have them keep going until they learn all the material. This second approach is more teacher intensive and more personalised, and therefore more expensive. And as Dylan Williams was pointing out in the previous keynote this is something that perhaps we don't have the resources to do, maybe this is too expensive for us.

The third approach is the *one-on-one tutoring* approach [6], which we know is too

expensive. The results that Bloom came up with in this study was that there were significant effects to both these augmentations, and Bloom says you get about one standard deviation of improvement from using the mastery learning and two full standard deviations from using the one-on-one tutoring [7]. So according to Bloom we have the answer to how to make our students succeed; it's just that we can't afford it. And now the challenge is: is there something else that we can do that's going to be as effective as one-on-one tutoring that uses some other kind of technology. In 1984 we didn't have all the technology that was available to us today. Is there something else we can do now that can answer Bloom's challenge from so long ago?

First I want to talk a little bit about what exactly conventional learning is like and then contrast that with some of the other possibilities. So I learned a little from another educator, Frank Rhodes, ex-president of Cornell who said "*in the basic business of teaching resident students universities have not diverged much from the method of Socrates -- except that most faculty members have now moved inside.*" [8]. So what do we know about Socrates? Well, maybe he had reasons for thinking that being inside could be a corrupting influence. So he did indeed like to teach outside, and we know that he's famous for his questioning approach, that he used questions in many ways, one was to keep the inquiry focused on what he thought was the lesson to be learned, another was to keep it appropriate and personalised to a learner. Another: he used it in sort of a bullying way to show the learner that he really knew what was going on but wasn't able to tear it out of themselves [9]. So maybe some marks off for harassment there, but it was a different time. [10]

But note that in this one-on-one questioning interaction Socrates is really much more like a tutor and not like a conventional teacher. So Rhodes was really wrong, it's not that we haven't moved forward at all from Socrates over two thousand years, it's that we've moved backwards from his tutoring approach to a non-tutoring approach that's not as effective. Now let's fast forward a millennium and a half. [11] Here we are, this is conventional learning in the year 1350. The lecture was a fantastic innovation in 1350 because books were very expensive things that had to be made by hand by monks so you couldn't really afford to give books to all the students. But in a lecture one person could read this material and get it into the minds of all the students at once, so the lecture was the broadcasting of the time. Now you would think that although it was a great technology for 1350, by a hundred years later the lecture probably should have died out, because with the new learning technology that Gutenberg came up with, it became cheap to have books, so why have this one guy in the front of the class reading from the book when all the students could be reading from it simultaneously.

Now, I'll leave it to the historians to sort out why that didn't happen, why the lecture didn't die out in 1480 or so, and I'll concentrate just on the problems with lectures. And note that we don't need any of today's technologies and studies, although there have been plenty of well documented studies. We can just go back to 1350 and they've identified the problem right here, it's this guy: the one who's sleeping in class. [12] So, even in 1350 we knew that attention wanders when you lecture, people fall asleep and you can look around at the other people in the room and only about a third of them are really paying

attention. The others, if they're not sleeping, they're starting to nod off or they're talking to their neighbours. That's what happens.

In places where educational results really matter regulations are actually put in place to limit lectures. For example in the United States, the Federal Aviation Administration is responsible for the teaching of pilots, and they put in regulations that say no lecture can be longer than twenty minutes maximum. Why do they do that? Because they know if they don't that's what happens: planes crash. [14]

They found out that there were bad results and they put regulations in place to, to take care of the problem. Unfortunately in other industries where maybe a lapse in attention is not as dramatic the lecture survives unregulated. Now I notice that this particular enlightened conference for most of the sessions does respect the twenty minute limit, unfortunately there is one exception to that: the keynote lectures. So that bodes ill. We'll just have to hope for the best. We're only about eight minutes into the talk so maybe catch about five minutes more then wander out, talk to your friends and come back at the end. I'll try to stop early and leave plenty of time for a more effective method of interactive question-and-answer.

Okay, so one of the big problems with lectures is that you just can't do it too long or it starts to get stale. There's another problem: the lecture is suited only to the auditory learning style. We know that there are multiple methods of learning, and as an aside I was glad when I went to my daughter's third grade parent introduction night, where the teacher explained what was going to go on in the classroom. She started to talk about different kinds of learning styles. The teacher said that we recognise some of the kids are visual, they respond to written words and diagrams, some are auditory, they respond to spoken words and music, some are tactile, they respond to body movements and touch. And we try to individualise the instruction to be appropriate to each child and the modality that they prefer. I thought: this is great, here's a teacher that's making use of the latest research, my child's going to get a good education. And then she went on and said 'and in fact you can diagnose which type of learner each student is just by listening to how they talk -- if you're a visual learner you'll say things like "I see what you're saying." If you're auditory you'll say things like "I hear you." And if you're tactile you'll say things like "I feel that you're right."' And I was thinking 'Nah, that can't be right. It can't be that easy.' And then I was heartened when Google hired T.V. Ramen [15] who's a very well noted and respected engineer, happens to be blind, and within my first five minutes of talking to him he said 'I see what you're saying.' So, so I think my intuition was right on that one and it still is difficult to diagnose exactly what learning style a student is using, and they may prefer different ones at different times.

But it's right that you don't want to rely on a modality of instruction that just goes after one of those learning styles. So that's another problem with the lecture. Now one more thing about the lecture or the conventional approach is that it lacks portability. Let's compare it to another form of delivering information, a musical performance. [16] Now back in 1350 if you wanted to hear a lecture, well you were lucky if there was one place to go to hear a lecture within a days hike. And that's what you were stuck with. If you

wanted to hear music you probably had a few more choices but basically it was whatever troubadours happened to be in town, you could go and listen to them perform music or you could perform it yourself with your friends and family. In the modern world of course, here in the twenty first century, if you want music you're no longer stuck with your local troubadour, you have thousands of choices and you can get the very best performances by the very best musicians in any genre you want, replayed at any time. [17]

But contrast to what's happened with the lecture. If you want to hear a lecture, up until the very final years of the twentieth century you only had the choice of the live performance, you had to hike down to your local performance place which was a university and hear the lectures there. Now in the last few years the lecture has become a little bit more portable and you can shift in time the lectures from your local university, you can even get some lectures from universities round the world. But we've not moved to this model where there are star performers. You know, we could have imagined a model where if you want to learn about linguistics you said 'who's the best linguist? Oh it's this guy Pinker, [18] I'm going to go get him, put him on my Ipod, I'm not going to be stuck with whoever happens to be at the local university.' But for some reason we didn't go down this route and, and I can't really explain why.

Let's look now at which learning technologies might help solve the two sigma problem. [19] We saw the problem with the conventional approach, with the lecture based approach. How can we get away from that and towards the tutoring-based model without all the expense of hiring a separate tutor for every student. I'm going to talk about some of the things that Google is doing in this area as representative of the kind of technologies, of course this conference has covered many others. I'm going to start off with the more formal approaches and then move to the informal and think about why they may actually be a better approach.

First, Google is offering Google applications, education edition. [20] That's free email, chat, calendars, web pages, and documents, and branded with your university or educational institute, and help in putting it all together. Up until two days ago I thought that this was just a collection of online applications but now I know that if you integrate it all properly it's actually called a *virtual learning environment*.

Here's a little sketch [21] on the back of a napkin that was drawn by Adrian Sanure who's in the IT office at Arizona State University, and he called this like technology for an advanced alien culture. What did he mean? Well, he noticed as you see on the left here that there was an increasing gap. As information officer at a University he said "we've got limited budgets, we can only spend so much, and it seems like every year the best of industry, the highest technology gets farther and farther away from us, and we can't spend enough to keep up." But he said "if we move to this model on the right where instead of us trying to buy our own technology every year we ally with a partner, now we're suddenly on our partner's growth curve rather than our own growth curve and we can go as far as they go." And that's the approach he's taken, you can see there's lots of things that comes with this or you can get help with. There are tools at the school level like

setting up a website. [22] And there are tools at the classroom level like planning a lesson using Google Earth maps. This is in a partnership with Amnesty International, one of the lessons plans we provide, and there are many other types of lesson available. [23]

Here's another partnership and this is fairly new, just launched last week. [24] We've teamed up with Creative Commons and with the Packard Foundation from Hewlett Packard to create an index of all open courseware. So any university or other institution that publishes their courses on the web can find their way into this index, and we'll search it and give you back the results. So here you see a search for solar energy, you get the MIT course first, that's appropriate because they're a good engineering school and they're at the forefront of this open courseware movement so maybe they deserve to be first. Then there are other results from Worldwatch, from other schools and so on, a combination of different results that you can browse through. This is useful for the teacher trying to put together a new course of their own, and for the students trying to supplement the course they're taking, and for the home learner who wants to put together their own course.

Many of you are probably familiar with *Google Scholar*, which gives you an opportunity to search over academic journals. So if I search for *ivory billed woodpecker*, [25] I can see the recent journal article from *Science* that caused a lot of excitement, I can see an older book that talks about some of the background behind it. And so on. But some, some of our users are not familiar with Scholar, so this is not a destination that they would go to, it's hidden to them. One of the things we've been doing recently is trying to bring all our results to all our users without them having to know where to go to. So if you know to go to Scholar, great, that's what you want, you're in good shape. But if you don't know, we still want to get you the right mix of scholarly and non-scholarly material. Here's an example of a search using the normal Google search for *bigtable* [26] which is an implementation that some of our engineers have come up with for a very large type of database table used over parallel machines. And there's a mixture of results here: there's papers that would have been in Scholar, these academic papers, there's a Wikipedia entry, and the third entry there is a video, and then some blogs, and then some more academic papers and so on. The point is that if you thought about this on your own and tried to narrow where you were going to search, you probably would have said 'well this is an academic thing, I'll go to Scholar.' But then you would have missed the video and some of these other entries. You probably wouldn't have thought first that video is the kind of result that's associated with this kind of query, but in fact we have the engineer who designed this system giving a technical talk and that's probably exactly the kind of result you want. So in some cases you can be outsmarting yourself by narrowing yourself down to a specific search.

Okay, and of course, regular Google search works pretty well for finding subject matter for courses as well. All those things are available for the teacher, for the formal student, and for the informal student. [27]

Now let's get back again to Bloom's challenge and remember what he said -- the average student under tutoring was two standard deviations above the control -- and try to

understand exactly what that means, because it really is quite extraordinary. [28] A standard bell shaped curve, the control, by definition, 50% are above the mean, if you define that being beyond two standard deviations means doing really well then two percent are doing really well. When you move to mastery learning things are shifted approximately one standard deviation to the right, you can see the outlined bump there getting a little bit of separation from the control but maybe not that much. But still quite a change in that now 84% are above the control mean and 16% are doing really well. So you multiply by eight the number of people that are doing really well. [29]

Once you go to tutoring, two standard deviations, now you've really got a big advantage, you can see the two humps are widely separated, 98% of the students are above average and 50% are doing well. That's an extraordinary accomplishment if it really holds up. Some more numbers from the paper [30], in the control, the conventional teaching, students were on task 65% of the time, so essentially one third of the time their attention span was wandering. If you believe the research represented by that 1350 painting it was more like two thirds of the time attention was wandering. So, not sure which one to believe, but we'll go with Bloom for the moment. But you switch to mastery learning, you pick up a little bit of that time back, and are now at 70% time on task. Tutoring, you've got 90%, so you've really got their attention and maybe that's where most of the gains come from.

Bloom also reported correlation numbers, so in conventional teaching there's a .60 correlation between the aptitude (that's before you give the course you test them for background knowledge and the achievement) and the results of taking the course. That's saying 60% of the result is determined before you even get started. You're just allowing the top students to stay at the top as much as pushing them on. With mastery learning the correlation goes down to .35, and with tutoring down to .25. So in part that's saying everybody has a fairer chance to advance. In part it may also just be a statistical artifact because if everybody is squished off to the right then they all have about the same score, so now the correlations disappear because everybody's doing excellent.

In the paper Bloom had four main suggestions on what you could do to achieve these two sigma results without tutoring. [31] One, he said careful review of previous material, for example if you're teaching an algebra two class make sure you go back over the algebra one material. And do this with what he called enhanced cues, participation, and reinforcement, which is just a sneaky way of getting the tutoring back in: it's going over the materials, seeing who's getting it and not getting it, and if they're not getting it giving them special attention. Bloom also said a student support system with groups of two or three students studying together works well, of course this was back in the days when everything was done offline, maybe more effective now done online. Special programmes for reading and study skills, that's a remedial programme, and he said that computer learning can work but only for very motivated students. And maybe that has changed now as well.

Let's look a little bit at what kids today are doing really well, are out there on the right past the two sigma mark. [32] There's a number of ways you can achieve this kind of

level of high achievement. Some kids do really well on computer based training materials. And this is a slide from a site called *Hey Math*, which provides computer based training for algebra. [33] Dylan William talked about a similar programme at CMU; he said the results are very effective but it took twenty years to develop. There are a few subjects like algebra where it seems like it makes sense to put in that twenty years effort or however much effort it takes. It's complex material but it is very self-contained and you can put together the materials once and for all and really nail it, and then you can teach it to a lot of students. And so I think there are these few areas where that makes sense, but for most areas it probably doesn't make sense.

And even for things like algebra you have to think about how well these materials you put together are going to port to different kinds of students. So for example this lesson seems like it'll work well for somebody who's a football fan [34], but what about this kid [35]? Say he prefers baseball, or this kid who likes cricket better, or this kid [36] who doesn't actually know anything about organised sports but he's been herding and counting his family sheep since the time he was four year old. These are all very different types of learners, maybe trying to fit them all into the same funnel, computer based training with a specific theme is not going to work that well.

What other kids are doing well? Here's a kid who had a good summer [37], I don't know how much you heard about him over here but in the States he was hot. This is a seventeen year old high school kid who in his spare time over the summer unlocked the iPhone, traded the result for a car and then packed off to college. He said he spent a few hundreds of hours and he had the help of a few online friends. This is just the exact opposite of computer based algebra training. There's no lesson plan, in fact there couldn't be a lesson plan because before this guy did it nobody knew how to accomplish what he tried to do. Maybe the Apple engineers did but they weren't telling. There's no real theory behind it, it's all practise, or I guess actually there is some theory but he learned it as he went along. When he started out he was pretty handy with a soldering iron, that gave him the confidence to get started. He got into it, he soon realised 'gee, there's a bunch of this software stuff I'm going to have to know, I don't know much about that. But I got some friends online, maybe they can help me, maybe they can lead me in the right direction.' And together they started to work on it, he learned what he needed to know. It was collaborative, it was just-in-time learning and again he was being tutored whether he knew it or not by his friends: they taught him about software, he taught them about soldering, and between the group of them they were able to come up with the answer.

So that's a successful result, but is it duplicable? Well I don't know. Is he an exceptional kid or was he just lucky? Or was it just the determination to get started on a project that was important rather than any attributes that he had to have? I don't know.

Here's another example, this is Olin College [38], a college that's dedicated to every kid having this kind of iPhone-like experience. It's a new college, they just graduated their first graduating class last year, it's an engineering college located in Boston. The idea is that instead of theory-heavy lectures in segregated disciplines and individual efforts, they go after design exercises and interdisciplinary studies and team work.

The curriculum integrates these disciplines with practical projects. It's a do and learn approach. Olin's president Miller said "Students start out with an audacious project, which would in many institutions be heretical, except we do that deliberately. Because, after all, when you get hired in a corporation, that's the first thing that happens to you: they give you a challenge for which you've not had the prerequisites. It's all about learning how to learn. So we do that here from day one."

This was the first assignment for the freshmen class coming in: they divided the students into small groups and they assigned them the task of (in eight weeks) designing, building and demonstrating a pulse oximeter. That's this little device that you clip onto your finger and it tells you your pulse and it tells you your blood oxygen levels without pricking your skin [39]. The professors showed the groups the commercial unit and then they referred them to some patents and some technical documentation. And then the faculty's plan was to watch the students and see where they failed and step in and guide them towards a solution. And the problem was they never failed, they just got started, they said 'we understand what this problem is, we understand what we don't know.' And they went off and they tried to learn it. In eight weeks all the teams got the project done. And one of the side effects was that at the end of eight weeks they were now receptive for the theory that they didn't have. In other classes you would have spent two years learning about the theory of transistors and so on before they let you do a project. Here they skipped all that, threw them in at the deep end of the pool, said 'go ahead and do the project.' The students learned just enough, but at the end of that eight weeks they said 'you know I'm really interested in transistors now, I want to know some of that theory.' And then they started to teach the theory behind it.

One more example [40]. This is probably cheating because it is from science fiction, not from reality, but this world that Vernor Vinge imagines takes Olin college to an extreme. This is a world in which fictional high school students are doing the equivalent of a years worth of graduate level research in a few weeks time. And they do this because in this is novel of the near future, the students are permanently wired in, they've got wearable computers, and are on the grid all the time, and communicating with their fellow students nearby and across the world. And they're doing Internet searches by blinking their eyebrows without having to type on anything and seeing the results projected.

Many of the adults who are coming from the present day into this world of maybe twenty years in the future are lagging behind, and so the adults are sent to high school for remedial education because they can't keep up with the high school kids. This is a possible future, who knows if it's going to work out that way.

Let's go back to Frank Rhodes [41], who has some answers to the question of how important really is the social networking aspect? Is it just if you put the kids in connection with each other and with some potential teachers, is that all they need or is formal learning a vital part of the equation? Rhodes said that without community knowledge becomes idiosyncratic, the lone learner studying in isolation is vulnerable to narrowness, dogmatism, and untested assumption. And learning misses out on being

expansive and informed, contested by opposing interpretations and refined by alternative viewpoints. Without community personal discovery is limited. Okay, so I think it's clear that this is important but the question is can it be the only thing, or can it be the main thing? I think now we're ready to answer the question, the challenge from Hal [42] and from Benjamin Bloom [43] of what education should be like, how we can achieve this two sigma effect.

Let's try to synthesize these examples and here's the best I've come up with. I think the best way to get to this ninety percent time on task individualised learning is to concentrate on these points [44]: one, centre the education on engaging real world projects. The kids' feelings about the project are probably more important than the actual words that are coming out of the teacher. And if they're excited about it they're going to get better results, and it seems like a project based approach is the best way to get them excited.

Exploring in teams seems to help. We can't afford individual tutors but we can afford to put people together into teams, you can get that for free and you get much of these effects, of the two sigma effect of the tutors just by having people work together and help each other. Now the teachers are serving as facilitators and they can point to the theoretical knowledge when it's needed. And part of the point here is that it's needed much less than you think. That means that the teacher now is free to become a tutor because the kids are doing most of the work, the teacher is acting as a facilitator and one teacher can be spread out over thirty students and still have enough time to cover them all. And finally different students learn differently but let them figure it out from the world full of information, don't try to create all the materials ahead of time and make one set of materials for everybody.

We've gone through these transformations in the way information has been used: we started off with the lecture when there was one copy of the book in the world and the lecturer was reading it, the students were learning it. We got the revolution from Gutenberg that now students all could read the book. And more recently we've had the revolution of the Internet where the students now can do two things: one, they can access all this information from around the world, all the books and all this other kind of information. But more importantly they can access each other, they can access fellow students and other people throughout the world. When we put that together, that's a revolution that I think comes the closest to being a scalable solution to Bloom's challenge. So thank you and I'll stop lecturing and make the rest of this interactive.

(Applause)

Mike Sharples: Thank you very much indeed Peter. We do have time for questions, and I hope that you're willing to take questions now. We've got a roving mic' here and while you're thinking of what questions to ask what we'll do is start with one of the questions that have come in from Illuminate, and it's from Panos in Edinburgh: 'Web 2.0 technologies seem to be very busy social spaces, and not necessarily educational and learning places. It appears that it makes students to be busy with things not related to

learning but only to information sharing. Is this education? Any comments?'

A: So from my point of view, the point of these technologies is to connect the people, not to be the only place where they can act. So once you found the right person that you want to work with, you can do that in person if you happen to live nearby, you can talk on the phone, you can use web 1.0 technologies of email or whatever you want, and I don't think you need to stay within the bounds of the social networking site or wherever you started. I think it's making the connection that's important and what you do with it then -- there's many options.

MS: Other questions, yes, there's one there in the centre at the top.

Q2: This is sort of coming back to the Bloom paper which is sort of one of my little hobby horses. The, when, at the end of the Bloom paper he suddenly starts talking about 'well what about higher mental process think?' And you sort of start to realize that all of the studies in the classroom that he was doing was on students learning how to do things that they're taught the exact process step by step of every single thing they have to do along the way. And you also discover quite how bad his control group was because he found that setting homework makes a half a sigma improvement. Marking the homework makes another improvement. Actually teaching them stuff based on the, the, the stuff that you found out makes another improvement. Lots of educational technology, things like the cognitive tutors that Dylan mentioned yesterday, focus quite heavily on shown learning games by doing the Bloom like stuff where you train people in a set process. Something that very few systems seem to do is let students try out their own ideas and kind of react to that and sort of show them the consequences of their ideas. And it seems to me that's the only way that you can teach someone in situations where you don't already know the process. And for most interesting questions we'll come across in life we don't already know the exact process we'll need to go through. So anyway it was sort of just something that struck me, and I wanted to make that point and see whether you had any, any thoughts on it.

A: Yeah, I certainly agree with you and I went through the same process. I read the paper and at the end I said 'you know, this seems like it really applies to rote learning but does it apply to more creative learning?' And I said 'oh rats, Hal Abelson lied to me, I probably have to read more than one paper.' So I was concerned about that but I think you're right that the important parts of learning are not in rote learning. I think Bloom did emphasize that too much and I think we want to move away from that, and I think that as you point out and the theme of this conference points out, you're losing control. You're no longer saying 'well there's this narrow path that you can follow and we know exactly what the right and wrong answers are. And if you turn right or left that's okay because other people have gone down that path before and we know how to follow you up.' You start to get into areas where there are no black and white answers and where the student may be exploring something where we haven't seen the answer before. And I guess part of my message is that you should embrace that rather than fear it, and be willing to go there and trust that the students are going to learn something and maybe something new, and maybe something you as a teacher don't know about or have to figure out as you go

along, but go ahead and let them take that path.

MS: Okay, another question. We've got one down here, while the microphone's coming down we've got another question via Illuminate from Adrian in Manchester, referring to your Creative Commons search engine he says 'I'm wondering what criteria were used for ranking the Creative Commons open courseware materials, the one where MIT came out the top.'

A: Yeah, so I can't tell you the details, I can tell you it was built on Google custom search, which is a facility that we offer to anybody, anybody can create a custom search, this one was done by this partnership with Creative Commons. There's two main things you can do in a custom search: one is you can rule sites in or out, you can say 'I only want to take results from the following set of sites that meet these criteria,' or 'I don't want to take results from these other ones.' And the other is you can influence the core Google rankings up or down by saying 'the following terms are important to me as I rate these results.' Those are the tools that we provide. I have no idea what this particular custom search has done in terms of how it's used those tools to provide the ranking.

Q: And are the criteria generally set by Google or in consultation? How does that work?

A: The criteria are set by whoever publishes the custom search. We give you the tools in which you can state the policy that you want to implement, but it's up to you to state that policy.

Q: Okay.

Q3: Okay, I'm interested in the star performer idea and why that hasn't really taken off, and what might help it to, or is it perhaps an outdated notion?

A: I'm not sure why that is. It just seemed odd to me that in music we've gone so far in that direction. Sure, you can still go to your local live performance but most people most of the time they've got their head phones in and they're all listening to the same top forty groups. Or maybe they're not; some people have very varied choices but they're listening to recorded music rather than seeking out the local ones. I don't know, maybe the universities had a lot of power and they kept things that way. Until very recently I think it was very clunky to try to view video educational materials online. It's still better being there in real life, you get the chance to ask questions and so on. But I've been in lecture halls that were twice the size of this one and there's not a lot of interpersonal reaction going on there between the audience and the professor, so I'm not sure that you give up that much by putting the materials online. And of course we could go so much more beyond that. If you are going to go to this star model you could have much higher production values and put in animations and simulations and so on mixed in, it's expensive to do that, it seems like it's probably worth while doing but it hasn't happened and I don't know why.

MS: One over there, and one there.

Q4: Thank you, the start of your talk caused me some amusement because you used two exactly of the same images that I used in my own professorial inaugural lecture some years ago, which was the image of the people asleep in the lecture theatre and the Socrates wandering around outdoors. In my lecture I used a third image which was a group of Neanderthals who were driving mammoths off a cliff. And actually I think it was quite pertinent because what was going on was there's a whole bunch of young Neanderthals learning how to drive these mammoths off a cliff with a couple of old guys stood towards the back observing what they were doing. And it struck me that that is actually how an awful lot of real learning works, you've got people getting stuck into real world problems, learning how with each other with a bit of expertise on the side they can draw on. And I think the model that you put up towards the end is far closer to the Neanderthal than it is possibly to the Socratic. And indeed, you know, the tools we have now enable us, may enable us to adopt a old and effective model of learning, I wonder what you think of that.

A: That's interesting, I hadn't thought of going back to the Neanderthals. But I did think of going back to maybe the late nineteenth century or early twentieth century where we had the apprentice model being more common. I look at Olin college and see that as being something like the apprentice model. They're still in a college, they're not owned by a corporation but they're on task doing and learning, and it seems like that's a good model. And maybe we were trying to protect our young people too much by keeping them in school, keeping them segregated from the real world too long. I think especially with all the influences they have now they're ready to grow up earlier and we should have them out there if not in a formal apprentice programme at least something that's approaching it where they're interacting with the real world more.

MS: We've got another question there and then one at the back and then that'll have to be it.

Q5: Thanks for the refreshing lecture Peter, I'm Kanishk Bailey from universitas twenty one global, and I have noticed that there's a trend in which people have started differentiating between a search engine and a knowledge depository like Wikipedia. And Wikipedia has become such a huge volume of knowledge and there's a trend in which people would like to go to Wikipedia to search out the knowledge they want rather than going to Google for searching information of any sort. So do you think that Google is facing competition in future or about to face competition in the future from Wikipedia? And what are the strategies which Google would like to have in order to maybe venture into the same kind of a knowledge creation format rather than knowledge searching format of business? Thank you.

A: We love Wikipedia, we're big fans of theirs and I forget the exact statistics but I think we give them something like a third of their traffic. So people are going to them directly but a substantial portion of people who get to them are coming to them through us. And I don't really see that as a threat, people going away, if they know they want to go to

Wikipedia I'm happy for them to go directly there but most people I think will continue to go through more general sources because they know they don't want all their information from one place, they want to have a choice. And so I think people are going to continue to do that.

MS: Okay, we've got one last question at the back there.

Q6: My question is a follow up to that but you showed us the Google VLE at the start, made a comment about that, and one of the things I feel is that the Google techniques and algorithms provide us with too much information. I'm wondering if you see users and user generated content as a way of providing more refined and sophisticated search techniques.

A: Okay, so the question is: is there too much information and can users help refine that? I think the problem with too much information is that there really is that much information, it's not just a problem of overloading people with the presentation. So we'll show you ten results and we'll say 'this is one to ten out of a million.' Nobody really cares that there's a million results and they're not going to go to the end and look at that all, so the results are up there and really the question is how much time do they have to investigate this area? How much time are they willing to put in? And can we find the good stuff for them in that allocated amount of time? And given that, I think that there is room for user generated commentary to help that process. To an extent Google has always been driven by user generated content: there are users who happen to be webmasters who publish material. There are other users who are webmasters who link between the material and that provides the sort of votes that we go on to judge appropriateness. We also go by looking at our users to judge what's important and so on. It's done by algorithms but all these algorithms have as their inputs actions that ultimately come from the user. Now I think we can use, in addition to those sources that we're already using, we can use some more explicit ones, people voting 'yes, this is a good site for this topic or for this key word or for this area.' And getting more people than just the people who traditionally have had access as webmasters, trying to open up that more, democratise it more so that other people can put in their voices as well and have them counted. That's an area that we're certainly looking at.

Mike Sharples: That's all, unfortunately, that we've got time for. I'd like to thank Peter Norvig very much for an entirely appropriate and stimulating and thought provoking talk, and entirely free of PowerPoint bullet points as well. So thank you very much indeed for rounding off the conference in such a stimulating way.

(Applause)