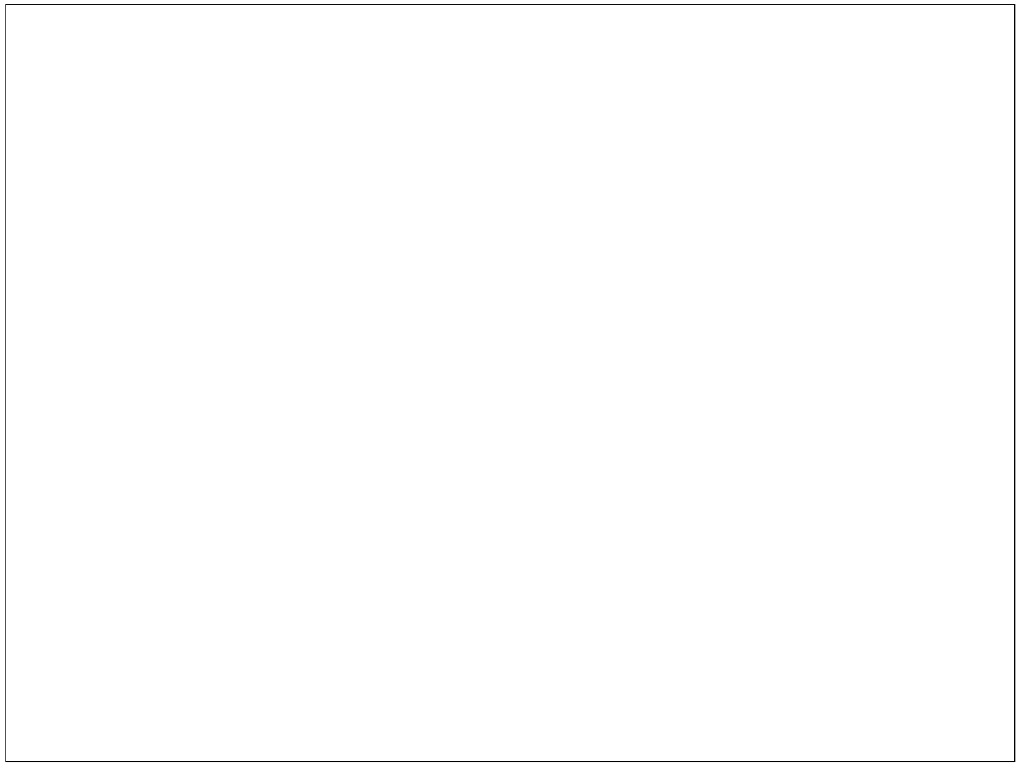

Introduction to Massive Open Online Courses (MOOCs)

Seb Schmoller
ICTP, Trieste, 30/9/2013

Contact details – seb@schmoller.net

0-0.25

Thanks for the opportunity to contribute to this workshop. It's great to back here at ICTP

Doing this talk is a bit of a leap in the dark.

My overall aim is to provide an……..

[NEXT]

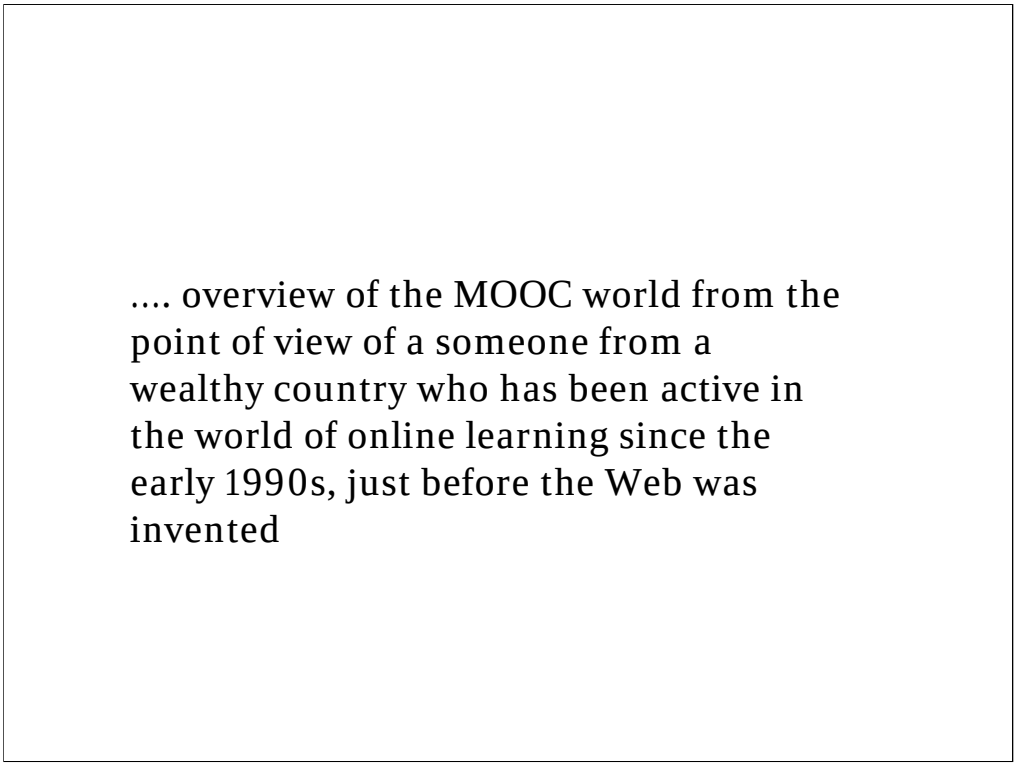

0.25/0.5

The talk is influenced by two factors. Firstly:

[NEXT]

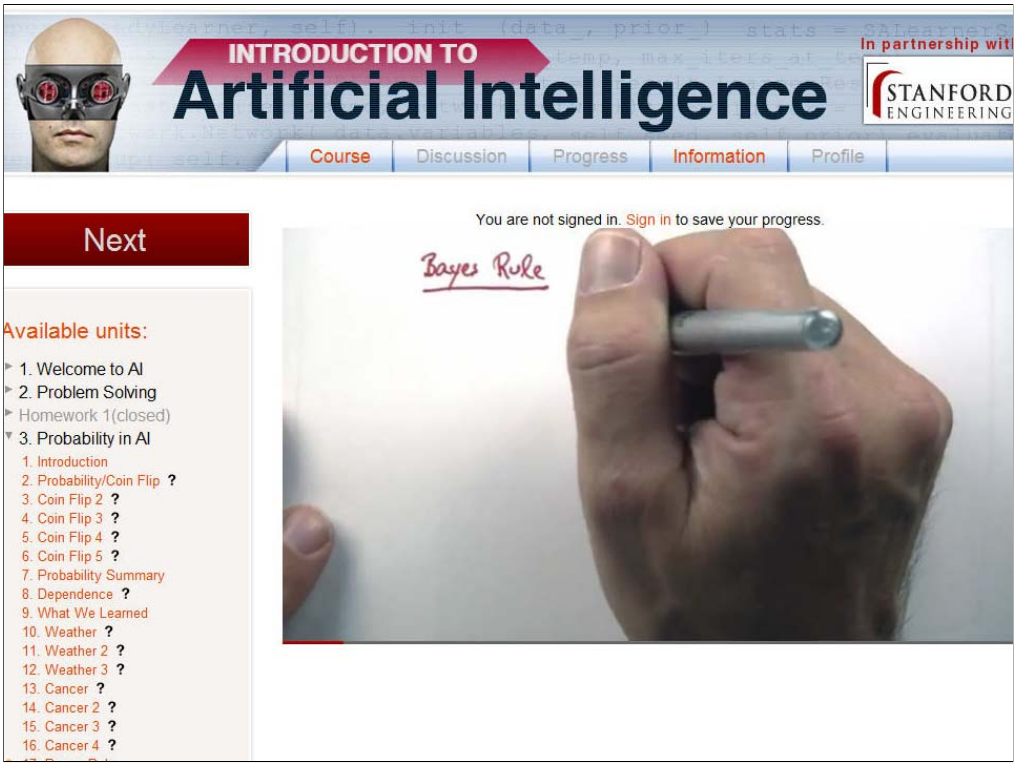

0.5/1.0

I had the luck to be a student and, as things turned out, participant observer in Peter Norvig's and Sebastian Thrun's 2011 "Introduction to Artificial Intelligence" online course

I was one of 23,000 from over 50,000 who actively started the course who completed the course.

Earlier this year I was also a student on Keith Devlin's much smaller

[NEXT]

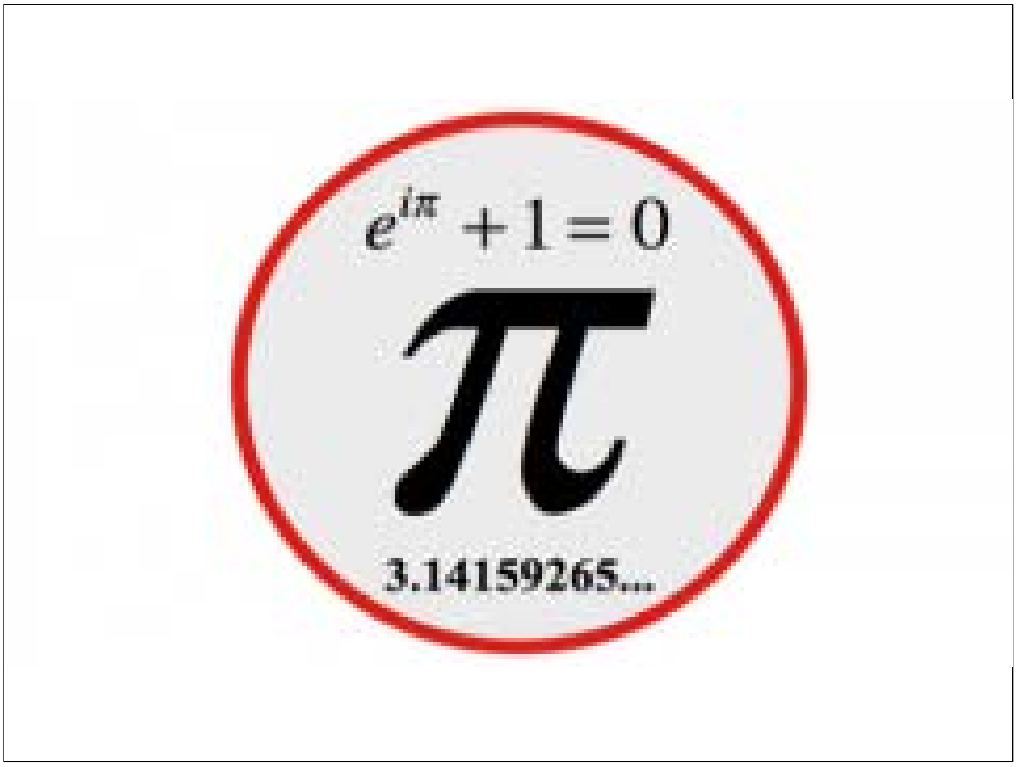

1/1.5

Introduction to Mathematical Thinking course, which was run under the Coursera banner.

The second factor that influences this talk is that with others I have got a concrete problem to try to solve with a MOOC, which is to

[NEXT]

.... create over the next two years, a free and open online maths MOOC suitable for the several million adults in England without intermediate level maths

http://www.ufi.co.uk/projects/cm

1.5-2.25

My perspective is as an enthusiastic, cautious, optimist.

Though I freely admit to having been blown away over the last two years by what I believe is a new phenomenon if only for reasons of the scale of what is going on, and whether or not there is also plenty of naïve hype about MOOCs. There most certainly is.

This is what I intend to cover in the talk:

[NEXT]

Structure

1. Backdrop to MOOCs
2. Some MOOC specifics
3. MOOC learners and their learning
4. Open questions
5. Conclusion
6. Resource links

2.25 - 3

I'll just give you a moment to read this list. My plan is for the talk to take a bit less than 30 minutes, which should give time for discussion.

I don't suggest that what I'm going to say is new.

It draws on the work of lots of other people.

It steers clear of things that I reckon will be covered in detail later in the workshop, for example high stakes testing, business models, and the broad issue of the meaning of the term "open".

So, to the backdrop.

[NEXT]

1. Backdrop to MOOCs

3.00 – 3.25

The next two charts show the relative frequency with which the terms “Online learning” and “MOOC” were used as Google search terms. They are made using “Google Trends”.

[NEXT]

Backdrop - Comparing some terms

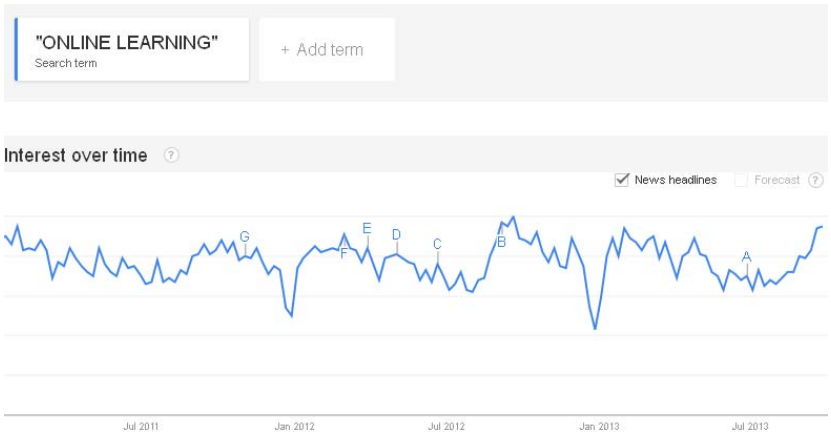

3.25-3.5

Here is “interest over time” in the term “Online learning”. Steady, with distinct lacks of interest between Christmas and the New Year.

[NEXT]

Backdrop - Comparing some terms

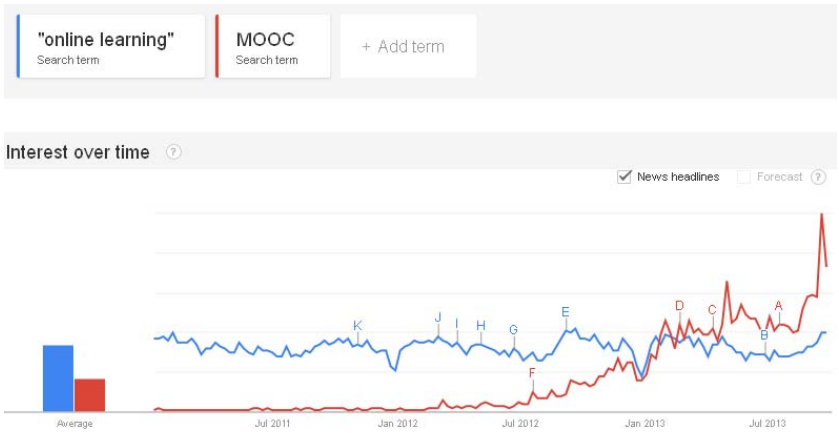

3.5-3.75

You can see here that use of the search term MOOC took off in January 2012, with its relative use exceeding use of Online Learning from early this year.

[NEXT]

Backdrop - Comparing some terms

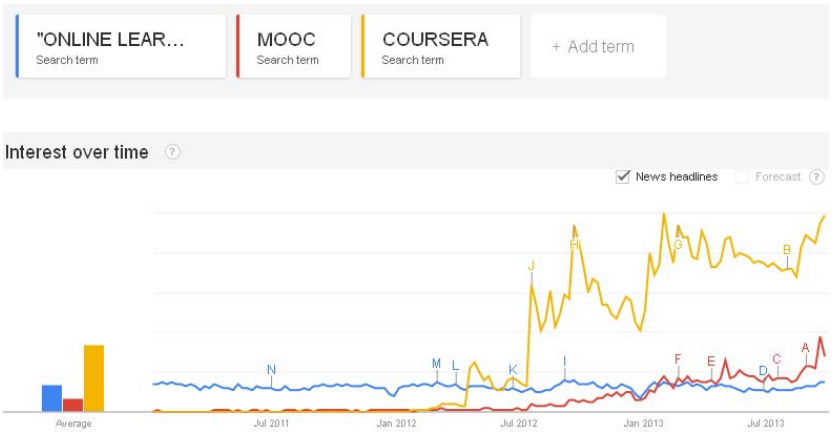

3.75-4

And the term Coursera (the Stanford University start-up company that is offering the largest number of MOOCs) has seen its use grow similarly.

In fact more searches are made with the term Coursera than with the term MOOC.

But……

[NEXT]

Backdrop - Comparing some terms

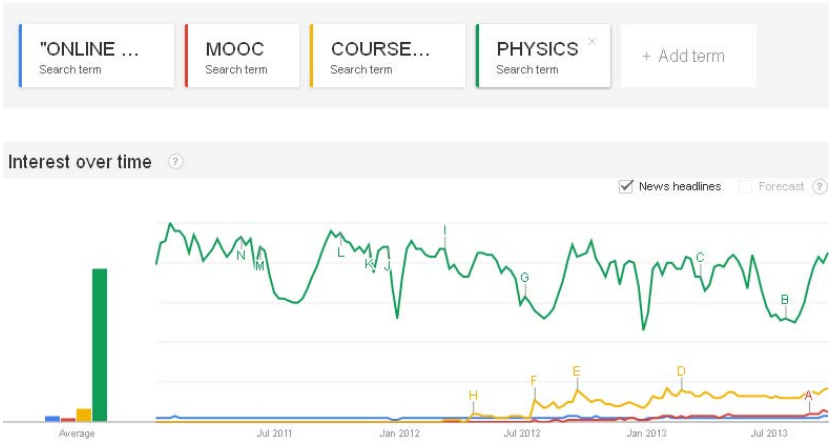

4-4.5

…. lets keep things in perspective…….

Nevertheless, something is going on.

What underpins the MOOC phenomenon is a big “mush” of other previous and parallel developments that relate to online learning.

Here are some examples:

[NEXT]

Backdrop – large scale online services that relate to learning

1. Google search
2. Wikipedia
3. Google translate
4. Duolingo – language learning system
5. Kahn Academy

4.5-5.5

Google search – our starting point when we want to find out about something – unless it is blocked

Wikipedia – the hugely popular and generally very reliable user-created encyclopaedia

Google Translate – an increasingly effective way to understand or at least get the gist of material in an unfamiliar languages

Duolingo – 10 million people are learning foreign languages with Duolingo, the brainchild of Guatemalan Luis von Ahn

Kahn Academy – a huge and extremely popular suit of self-study videos

What characterises these examples, which I've picked from the "mush"?

A key point is that a lot of them have been a long time in development. Here's the earliest version of Google Translate that I could find.

[NEXT]

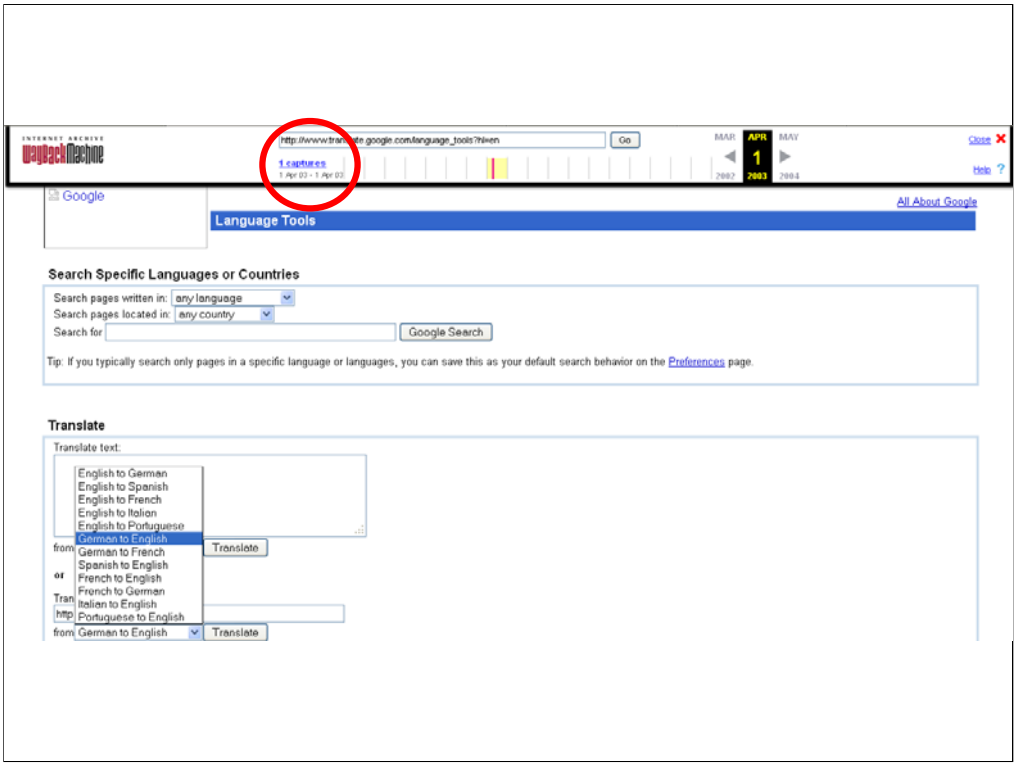

5.5 – 6.25

The number of language pairs it covers is now much wider, and the translation quality has improved a very great deal.

Meanwhile,

6.25 – 6.5

Duolingo has it origins

6.5 – 8.5

in von Ahn's reCaptcha system – and just as reCaptcha helps digitize hard to read bits of books, Duolingo is using language learners to translate Wikipedia.

I urge you to watch von Ahn's wonderfully funny talk on the subject.

But it's not just that these learning-related services that have been a long time in development. There are others that I've not mentioned:

- VLEs like Moodle or Blackboard have been in widespread use for over 10 years;

-online discussion groups pre-date the Web;

-the UK's Open University was running 12000 learner wholly online courses over 10 years ago.

And many of the **ideas** underpinning the design of online courses have been established for 10 years, with disciplines like psychology having many firmly established points (say about how memory works, or about the narrow bandwidth of consciousness) that are many decades old.

But the underpinning technologies have some other important characteristics:

Backdrop – characteristics of these large scale services

1. A long time in development
2. AI
3. Entrepreneurs from outside traditional education
4. Dependent on user-generated content or user-effort

8.5 – 9.5

I'll say a bit about #2 to #4.

Artificial Intelligence – a lot of the people behind the developments (Google's Peter Norvig, Duolingo's Luis von Ahn, say) are at the leading edge of the development of (artificially) intelligent systems – the same is true for leading figures in the US MOOC start-up companies.

Entrepreneurs from outside traditional education are to the fore. Take Salman Kahn for example; or Jimmy Wales, who co-founded Wikipedia (with a traditional academic scholar).

User-generated content is a very important facet of some of them (in the case of Wikipedia, reCaptcha, Duolingo): this thread runs through into MOOCs too.

To which we'll now turn.

[NEXT]

2. Some MOOC specifics

9.5 – 9.75

What makes a MOOC a MOOC?

MOOC specifics 1 – defining characteristic

.... a MOOC is an online course that has (or is designed to have) a large or very large number of learners, whose learning is **not supported** individually by paid teachers or other paid learning support workers

9.75 – 10.75

From the point of view of a learner it may **feel like** there is one-to-one support, and if the course is designed well, this ought to be the case; but the practical reality is that there are too many learners for any of them to actually receive individual attention by paid staff.

I did some **very** rough and ready calculations for the AI MOOC I was on. I reckon that on average, during the period the course was running, about 4 seconds of staff time was used by each hour of student learning time. That's perhaps 1/60th of the amount per learner hour in conventional university teaching, and 1/180th of the amount in conventional vocational education. And that was on the basis only of the 23000 people who completed the course, not the ~50,000 who were active on it.

But looking at other examples I think it is safe to say that the running and production costs per enrolled learner of a successful MOOC could be between 1/10 and 1/500 of those of a traditional course aiming to cover the same ground. Whether the learning would be as good as in a conventional course is obviously moot, but for me, this is the defining characteristic of MOOCs.

[NEXT]

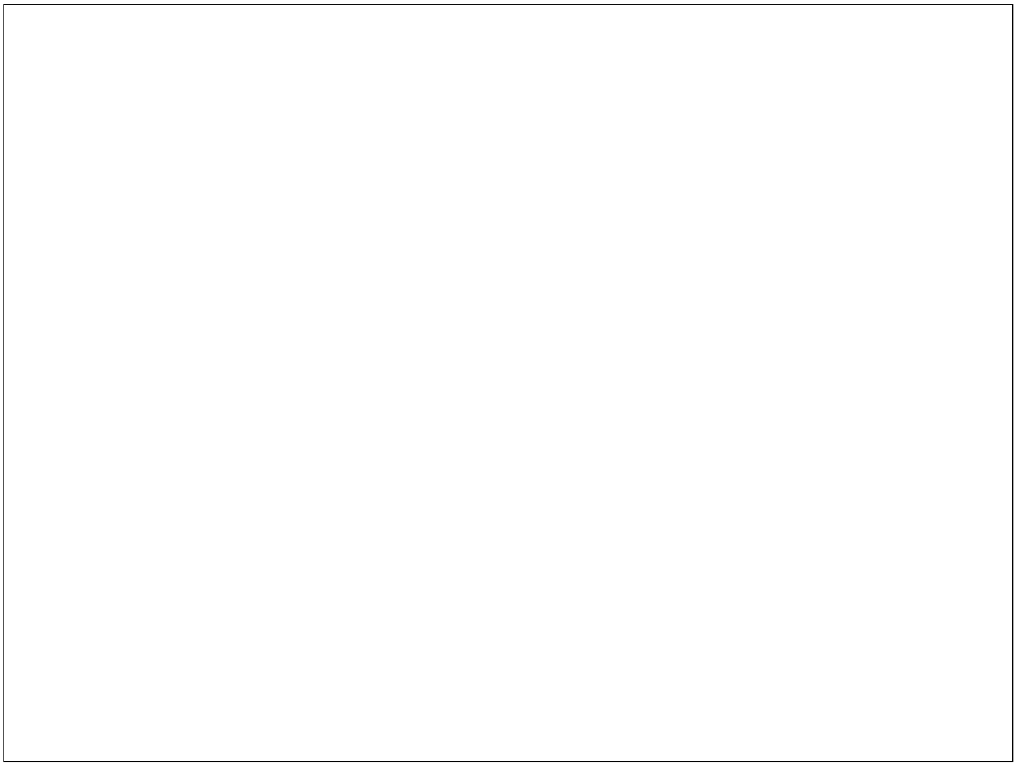

10.75 – 11.30

What can we say about types of MOOC?

Donald Clark's has done a taxonomy identifying eight types of MOOC, based on their learning functionality, not their origins. This is definitely worth googling for and I will include it in the resource links.

But at a simpler level it may be helpful to ask about the extent to which a MOOC has the following characteristics:

MOOC specifics 2 - characteristics

content is pre-built and interaction pre-designed or pre-programmed (possibly in a way that is responsive to individual learners' actions)

and/or

some or most content and interaction is created by participants "on the fly"

11.5 – 12.75

Examples of the first kind include most of the MOOCs offered by Coursera or by Udacity

Examples of the second kind include MOOCs run by communities of interest. For example, my old employer ALT organised a (small) MOOC called ocTEL with over 1000 learning technologists signed up to it.

But the key point here is that any MOOC can be designed to have a mix of features.

On the AI course, for example, there was a very active and supportive discussion board, run and moderated by students – so plenty of "on the fly" interaction.

On Coursera courses that use peer-based assessment, a lot of support is provided to learners by learners.

In the University of Alicante's UNIMOOC (on how to set up a new business), there is plenty of student/student interaction on projects.

[NEXT]

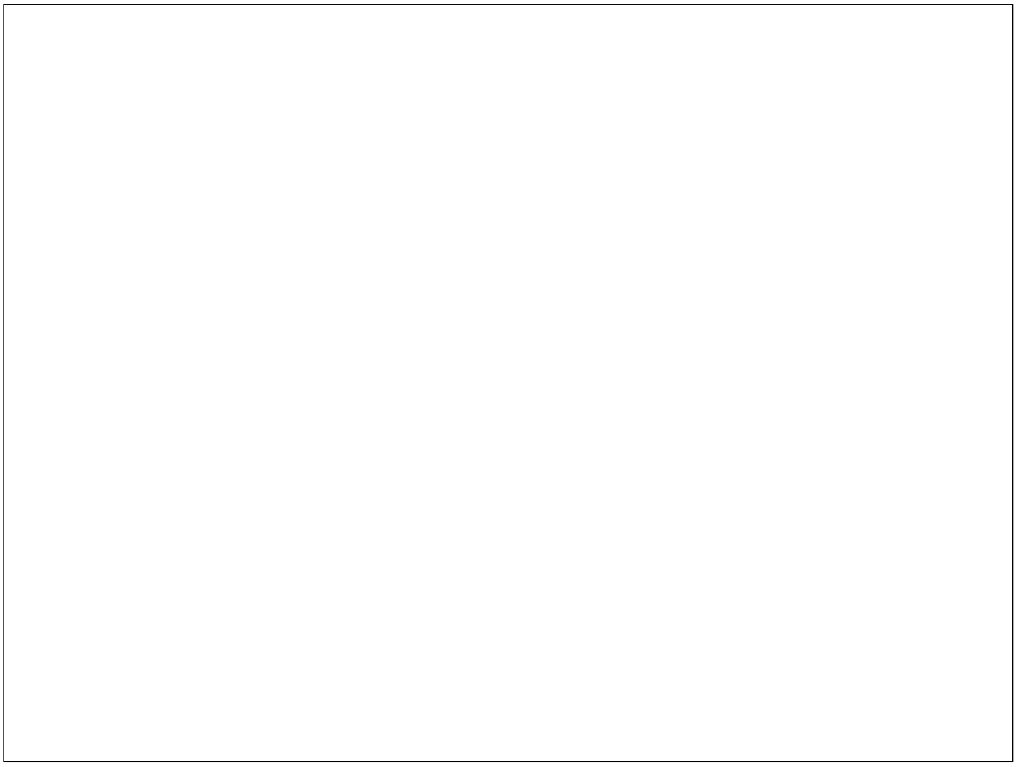

12.75-13

What about the throughput – that is, numbers of learners starting and finishing?

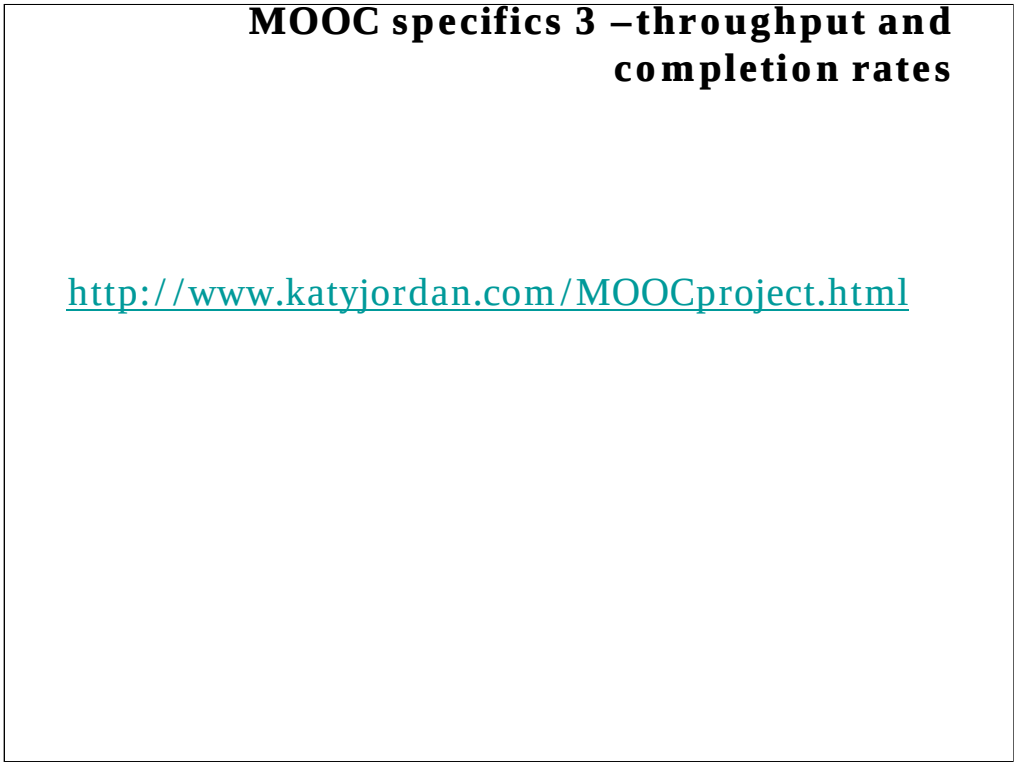

13 - 15.5

Some individual organisations have published data about their MOOCs, with the University of Edinburgh and Duke University good examples of this. In Duke's case they've provided lots of data about MOOC production costs as well.

But a lot of the broader data is being gathered informally. Katy Jordan, a researcher at the Open University has been doing just that. Her web site presents data about a range of courses, which she categorises according to the assessment methods used, the platform on which the course was run, the name of the organisation whose course it was, with data about throughput and completion.

Let's probe it a bit, depending on how time is going.

You can see the generally very low completion rates, and I am sure this issue will get discussed later on in the workshop. Personally I think it is almost irrelevant since the sign-up process for most MOOCs is so easy. What would be more useful to know is the proportion of those who get properly started on a course that go on to finish it. But even that ignores that when participation is voluntary, the meaning of completion rate is simply not the same as in a conventional course.

It would also be good to know who has learned how much, and at what cost!

Let's now take a look at ……

[NEXT]

3. MOOC learners and their learning

15.5 - 15.75

I want to start with an idea that I think is completely fundamental to thinking about MOOCs, their design, and how they help learners gain mastery. The idea is emphasised in two statements by different researchers, neither of whom are in the least bit “anti-teacher”.

15.75 - 16

Herb Simon a Nobel prize winning interdisciplinary US social scientist who was also one of the founders of Artificial Intelligence. He died in 2001 in his 80s.

16-16.25

Dylan Wiliam is a renowned UK expert on maths education and on formative assessment.

Learners and learning 1 - fundamentals

Herb Simon

"Learning results from what the student does and thinks and only from what the student does and thinks. The teacher can advance learning only by influencing what the student does to learn."

Dylan Wiliam (at the 2007 ALT conference)

"….. teachers do not create learning, and yet most teachers behave as if they do. Learners create learning. Teachers create the conditions under which learning can take place."

16.25 - 18

I'll read the quotes out.

These two statements get to the heart of a particular paradox which is that teachers tend to think it is ***the work that they do*** that causes learners to learn.

Whereas in fact it is **the work that the learners do** that results in the learners learning.

Only to the extent that a teacher induces a learner to make the right effort (that is, influences "what the student does to learn") does a teacher's work cause learning. Effective teachers know this, and base their practices on it.

Why does this matter for MOOCs?

The key questions, it seems to me are:

[NEXT]

Learners and learning 2 – the key questions

Can a really brilliantly designed and implemented MOOC:

• (Dylan Wiliam) “create the conditions under which learning can take place”?

• (Herb Simon) “advance learning .. by influencing what the student does to learn”?

And if yes, can it do it as or more effectively (and by implication more cheaply) than can conventional course?

18 - 19

Here’s the cautious view of Peter Norvig’s, the co-originator of the AI MOOC (albeit not in response to the exact questions I posed):

[NEXT]

Learners and learning 3

"My four main points are:

1. mastery learning with 1-1 tutoring is best;

2. there is a path to individualized online instruction that achieves that, but we aren't there (or close) yet;

3. motivation, not information is the key;

4. Simon's quote that learning is what students do, not teachers.

I would say that a MOOC is less effective at motivation than a good teacher, but a good MOOC is much better than a book, a bad teacher, or nothing."

(Source: comment 77 in http://tinyurl.com/b5jdbm2)

19-20.25

If you put me on the spot my answer to the questions would be:

[NEXT]

Learners and learning 3

1. Satisfaction.....

2. ..... is not evidence of learning

3. “Healthy worker” effect

4. Effort matters

5. There is a problem with generalisation

6. Protocol needed to consistently measure learning gain

20.25-23.25

1. The satisfaction rate evidence from those who complete MOOCs shows that a high proportion of MOOC learners think they’ve gained from their MOOC.
2. Being satisfied does not necessarily mean that you’ve learned much. (See work on the problem of super-teachers, whom everyone including the students think are doing a brilliant job: except not much learning is going on.)
3. Self-motivated people who’ve learned how to learn are not a promising population from which to draw generalisations about the effectiveness of MOOCs. (A sort of healthy worker effect)
4. There is evidence (from Udacity and San Jose State University) that the MOOC learners who make most progress are those who make most effort and persist most.
5. There is such a wide range of designs to MOOCs that generalising about MOOCs as an overall class of intervention is pretty pointless.
6. For MOOC providers it would be a good thing if widespread adoption could be got of a protocol which could be used to get normalised gain (or lack of it) judged. For every or many MOOCs.

[NEXT]

4. Open questions

23.25-23.5

Before finishing I'll say something about some of the open questions relating to MOOCs that bother me.

Open question 1

Will there tend inexorably to be only one or very few MOOC providers?

23.5 - 25

This does seem to be what happens with web-bases services – think of

- Wordpress and Blogger for blogs
- Google for search
- Facebook
- YouTube

So the answer is that there probably will end up being very few MOOC platform providers.

This will put a lot of power in a few organisations' hands and it is one reason why I welcome the announcement by Google and the US universities in the edX partnership that they will join forces to create, from 2014 an Open Source MOOC platform called Open edX, which, I believe will be for MOOCs what Wordpress has become for web-publishing.

[NEXT]

Open question 2

Will MOOCs remain a mainly "rich world" phenomenon?

25-25.5

Currently uptake is not restricted to (but is dominated by) the rich world; and rich world providers dominate the supply of courses.

And being a MOOC learner requires you to have connectivity and equipment that are not as ubiquitous as in the rich world.

But the gap is narrowing; and there is nothing to stop MOOCs being designed to work in "lower tech" ways, though working in an offline way is difficult if the MOOC needs to interact with "the platform" for it to work fully.

[NEXT]

Open question 3

Are MOOCs only good for a small (?) proportion who've learned how to learn?

25.5-26.5

So far this seems to be the case. But

i) How much does it matter? – even that small (?) proportion embraces a large number of people in absolute terms

ii) If it does matter, let's tackle it. That's where more sophisticated approaches like adaptive learning may have a part to play

iii) Is **learning how to learn** being muddled with **being motivated to learn**?

[NEXT]

Open question 4

Is what we are seeing mainly a hype-laden "bubble"?

26.5-29.5

Historically some people have always had a tendency to fall for "solutions" to educational and other challenges (educational TV, teaching machines, learning styles, one laptop per child).

There been a lot of boosting of the idea by

- People who've enjoyed being MOOC learners
- People who want to attack public education or make cuts to it
- News-savvy, gifted, and telegenic people like Sebastian Thrun, with large amounts of venture capital behind them, who catch the eye of politicians

I think there is enough that is solid and undeniable about MOOCs (not least the very large numbers of satisfied MOOC learners, and their overall scale) to be confident that this is no bubble. Which is not to say there's no hype to guard against.

If you are familiar with the "Gartner Hype Cycle" these diagrams may help.

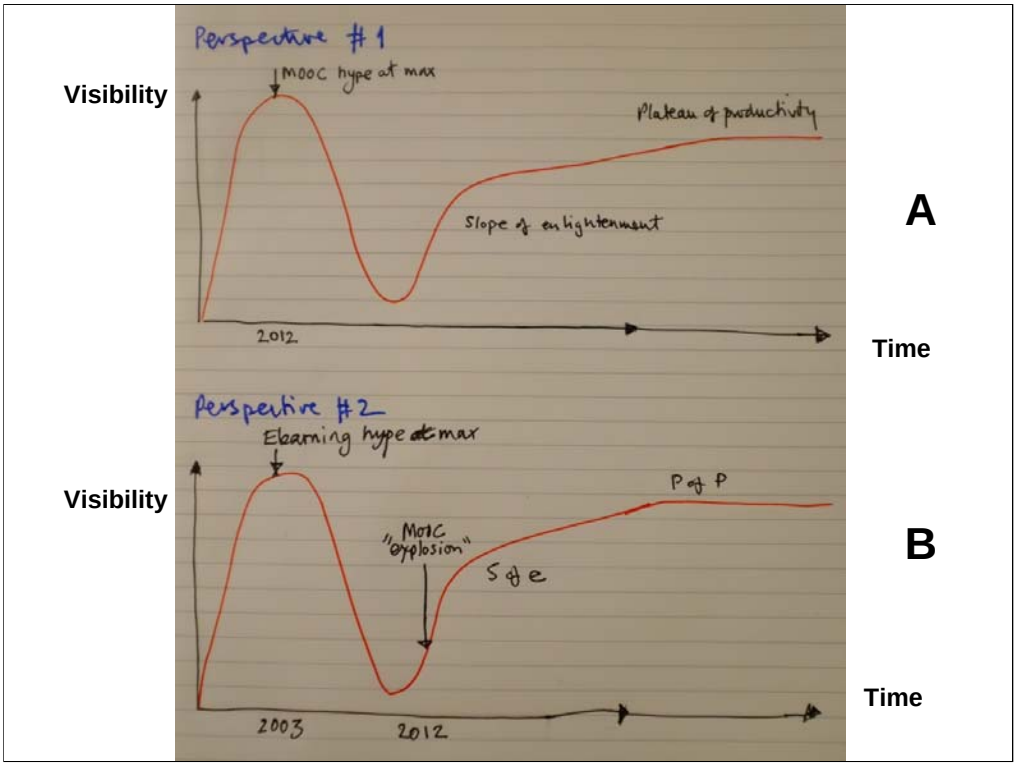

26.25 - 28

I think it is (just about!) arguable that MOOCs represent the coming together of a range of established but never hugely successful approaches that failed to gel from 2000 onwards, despite the predictions at that time, for example Eli Noam's 1995 "Electronics and the Dim Future of the University".

Under this view the MOOC "explosion" is happening on the so called "slope of enlightenment" rather than representing the "peak of inflated expectation", which was over a decade ago.

Thus B not A.

5. Conclusion

29.5-29.75

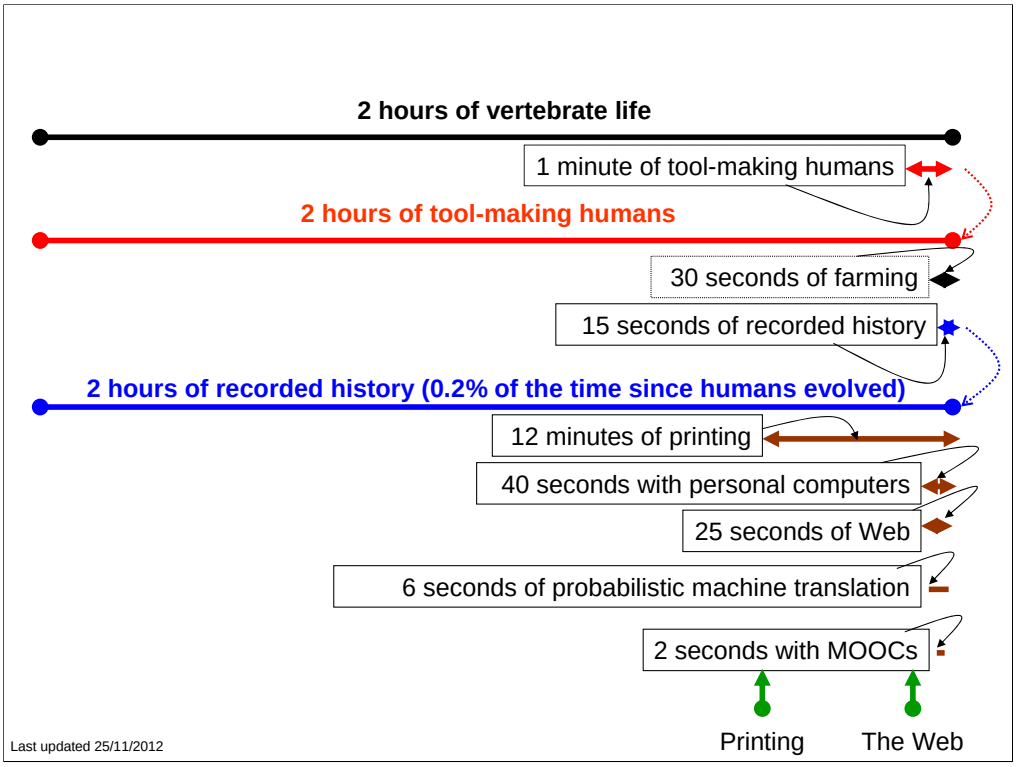

29.75 - 32

…….diagram, which is uses ideas in Richard Sennett's book "The Craftsman" based on those of the geneticist John Maynard Smith (who died in 1974).

1. "We" as a species (who evolved as hunter gatherers, not as book readers or Internet users) have not changed cognitively or biologically for hundreds of thousands of years.
2. We've only had recorded history for a small proportion of that time.
3. We've only had the tools and systems which concern us in this workshop for a tiny proportion of the time we've been recording history.

It took centuries for the print-based distribution and mediation of knowledge to become properly understood and for it to have its influence.

The thing is that we really are at the very beginning of a new phase in how humans – cognitively unchanged for the last few hundred thousands years, and unchanging for the next - create, distribute and mediate knowledge digitally, and learn.

MOOCs are a part – only a part – of that new phase.

For now, our job is to:

1 seeking to make MOOCs work;

2 understand the affordances of things like subject, level, MOOC design, learner-characteristics;

3 scientifically assess their impact.

6. Links – to follow….

Questions & discussion

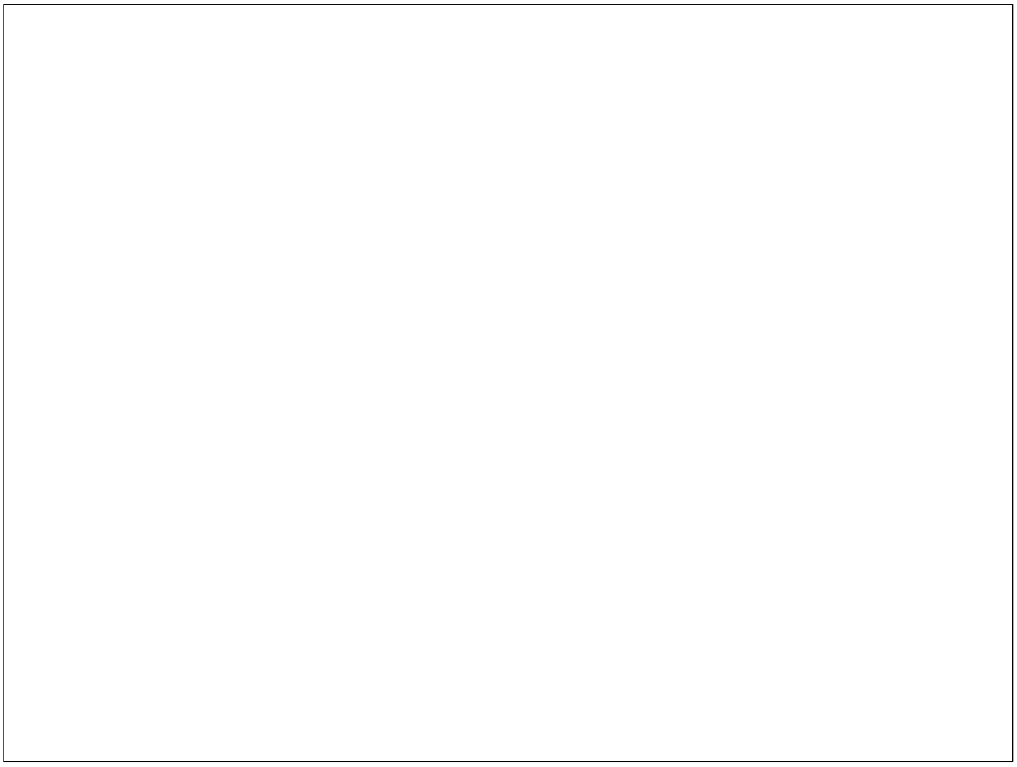

0.1 / 18.95

Breather before final slide

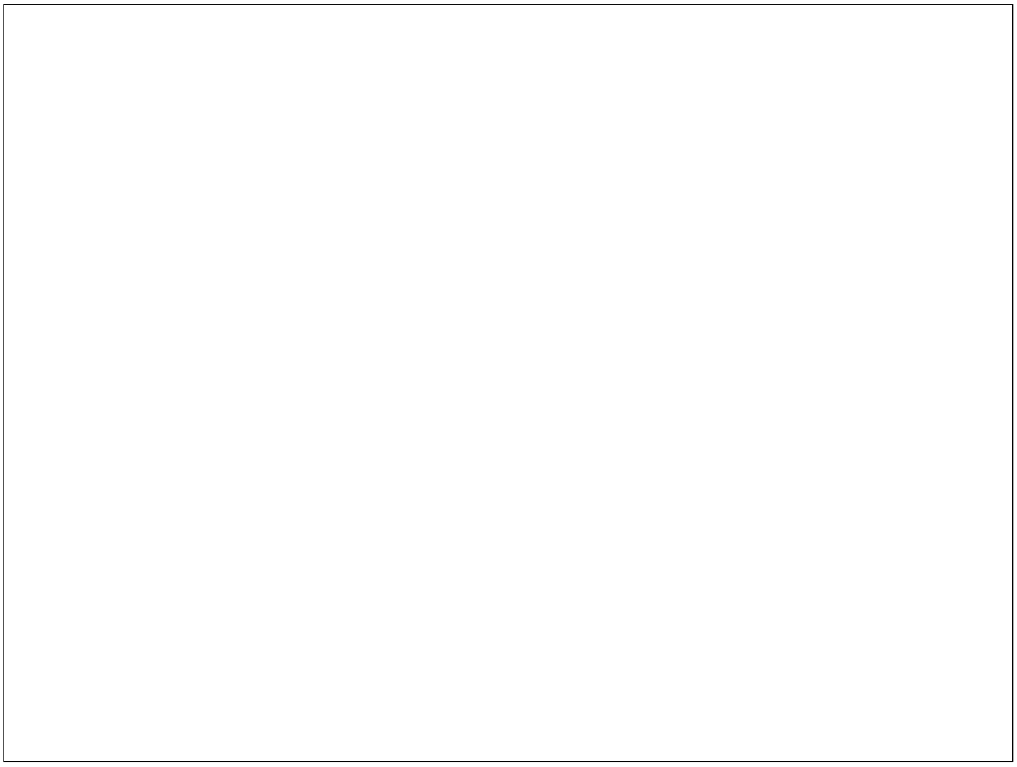

Final slide

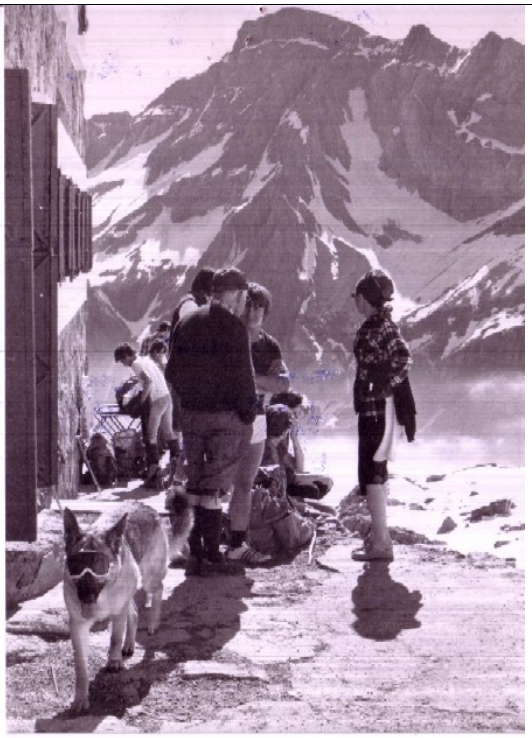

Refuge des Sarradets, Pyrenees, France, July 1979

In case a closing picture is needed