The discovery of our lack of knowledge

Lecture in response to Van der Leeuw lecture Joshua Foer, 2011

Robbert Dijkgraaf

What is the state of our knowledge? Leonid Brezhnev was once asked to sum up the state of the Soviet Union in one word. His answer was: 'Good'. Then he was asked what his answer would be if he were allowed to use two words. This time he said: 'Not good'.

In the same spirit, I would say that at first sight the state of our knowledge is excellent. We are living in a golden age. Not only are increasing amounts of raw information becoming available, along with better technology to deal with it, but a growing proportion of the world population is making use of this information. The total amount of information on earth is now estimated at the truly unimaginable figure of a zettabyte, which is 1 followed by 21 zeroes – roughly speaking, the memory of a billion modern PCs. Google alone uses about one million servers. Right now over two billion world citizens have access through the internet to a large proportion of that immense amount of data, and more and more often they access it through a smartphone, wherever and whenever they like. How different from the times when only a handful of scholars had access to a library with rare manuscripts, tucked away in a monastery or palace.

However, at second sight it's not going well at all. As Joshua Foer rightly says, we are becoming increasingly careless with all this knowledge. We are forgetting more and more, or remembering useless, unimportant things. We are living like scroungers on the garbage heap of information, and every new day information is dumped on us through a wide variety of channels. We rummage around with a bucket, looking for something useful or stimulating.

One of the reasons for this 'epidemic of forgetting' is that there is so much more to forget. We are coming up against the limitations of our brains' capacity to remember, which is estimated to be 'only' a petabyte, that is, 1 followed by 15 zeroes. We use most of this memory capacity – fortunately – for our personal memories; the rest we fill up with random facts. Now, in 2011, our brain can theoretically store only one millionth of all the information in the world, and in practice that is much less. The fraction of factual knowledge we remember will only become smaller and smaller in relation to the growing amount of existing information. Any aspiration to have complete knowledge is doomed to failure. Ultimately we will know nothing about everything.

Knowledge is becoming fragmented

Information is not only increasing in volume, but also changing in character. The basis of knowledge is fragmenting and being scattered in steadily diminishing units throughout the world. Modern human beings construct their own worlds from these building blocks. The degree to which we share a solid foundation is constantly lessening.

There are several reasons for this fragmentation. In the first place, research, by definition, digs further and deeper. This results in huge growth in raw material, including a lot of chaff. In the eighteenth century Carl Linnaeus classified about twelve thousand plants and animals. In his time it was still possible to have a personal relationship with these species. From elephant to sparrow, from stinging nettle to pine tree, they all still had a certain degree of cuddliness. Recently the total number of organisms on earth was estimated at 8.7 million, the vast majority of which still have to be discovered and named. Who can still feel anything in relation to these millions of insects and microbes? Do they actually still all need to be named?

Secondly, modern technology enables this knowledge to be distributed more broadly and in small units, from the internet to text messaging and Twitter. In principle this is a positive development. In the thousands of years of the history of science, information technology has always played a stimulating role, from the invention of writing to the printed book to the internet. Making information available and sharing it keeps science and scientists on the ball. But because of the fragmentation we are losing sight of the main story lines that in the past have held the facts together. Everyone puts their own IKEA flat pack of separate knowledge components together.

Finally, we have globalization. Knowledge comes to us from further and further away. Cultures which not so long ago played only a limited role have now become prominent; take for instance how much more influential and widely known China has become. While multiculturalism may no longer be politically correct in many European countries, in the worldwide knowledge community it is thriving and growing.

These combined forces of research, technology and globalization will only grow stronger. It looks as though it is pointless to fight against fragmentation. Under the motto of 'more is less' – there is more and more, but we share less and less – the landscape of information will have only a few tall, very narrow peaks which everyone can see from a great distance. This is the – let's say – 1 per cent of knowledge that is shared by 99 per cent of people. Apart from these few peaks, there will be infinite numbers of slopes and foothills full of pebbles. This is the 'long tail' of the distribution, where diversity grows and thrives, but where it is difficult to rise above the level of individual interest. This is the 99 per cent of knowledge that is shared by only 1 per cent of people.

Because of the growing diversity of the supply, knowledge stays around for increasingly short periods. Whether or not something sticks depends on the strength of the glue. Teachers are very keen on putting Post-it notes on things; but as soon as someone opens a door the draught blows everything away. Why is it that after four years of classical Greek at school I scarcely remember anything apart from the alphabet (which as a physicist I still use every day)? The strength of knowledge also depends on the quality of the glue.

The Dutch canon game

If we want to put a stop to the Great Forgetting, we will have to have a keen debate about what exactly we want to remember and how. There is no doubt that without concrete knowledge of facts no true creativity or innovation can exist. What is less clear is *which* facts we should select from the overwhelming and growing supply of information. How can we find our way in this mountain of information? Where should we put signposts so that we don't get lost?

These questions are extremely pertinent when it comes to drawing up canons – something which has become a popular game in the Netherlands since the success of the Canon of Dutch History drawn up by a team led by Frits van Oostrom, former president of the Royal Netherlands Academy of Arts and Sciences, in response to diminishing historical awareness. In the Netherlands we now have canons of skating, of rock music, of the Green Heart and Christian education – among others. In fact, I have played the game myself. A few colleagues and I put together the 'science canon', a list of fifty subjects relating to science and technology that every Dutch person should be familiar with. We particularly enjoyed taking 'zero' as our first subject.

This wave of 'canonization', which is totally unknown in countries such as those to the south of us, has made a positive contribution to our thinking about knowledge. Of course any canon immediately leads to a fierce discussion about which topics should or should not be included in the list and about whether or not a list like this gives a distorted picture of what it is presenting. But to be honest, such discussions are mainly beside the point. All calls for more differentiation are appropriate and justified, but that differentiation can only take place if an initial broad overview has been made. The canon is not a caricature, but a sketch which captures the essence of a subject with a few strong lines which can later be differentiated in more detailed drawings.

Kissing numbers

A canon presents concepts which form the hub of the knowledge network – the centre where many spokes come together. These subjects, which can be viewed from many perspectives, have what mathematicians call high 'kissing numbers'. In crystallography, the kissing number is the number of neighbouring atoms touching a particular atom. The higher the kissing number is, the more 'beautiful' the crystal is.

When we were drawing up the science canon, a high kissing number was an important criterion. Take for instance the subject of 'evolution'. It has many facets. In the first place we can approach it as a general concept of biology – as the alteration of life on earth by means of natural selection. A second possibility is to examine a concrete example that captures the essence, such as the diversity of finches on the Galapagos Islands. A third possibility is the historical or personal approach – in this case Charles Darwin and his voyage on the Beagle, including all the practical and intellectual difficulties he was faced with. Then there is the angle of application (vaccinations against evolving influenza viruses) or of current affairs (such as the Dutch fundamentalist television presenter who lost his belief in creationism). All of these perspectives are useful and suitable for teaching purposes. Each sheds a different light on the theory of evolution and they can reinforce each other. Evolution needs to be part of our intellectual baggage.

This multiplicity is to be found even in the hardest of the hard sciences: mathematics. The American mathematician Bill Thurston, winner of the prestigious Fields medal in 1984 for his work on the classification of three-dimensional spaces, gave a wonderful illustration of this in an extremely interesting and against-the-grain analysis of the usefulness of proofs in mathematics. In this paper he states that in mathematics an exact definition is not always the only goal. To stimulate further progress it is also important for the context of the new ideas to be conveyed properly.

Thurston takes the concept of the 'derivative of a function' as an example. He starts with the formal definition taught in standard lectures. At first sight this definition is not very illuminating, since you need to have mastered many other elements first. Interestingly, in his article he also makes an error – a typo? – in the definition! But there are also other ways of thinking about derivatives. For instance, the second definition is about speed. We all have an intuitive understanding of speed and in this way we see that the derivative can also have a direction. Then there is the third definition: the slope of a line tangent to the graph of the function. This geometric representation makes generalization to several variables obvious. Then comes the fourth definition: infinitesimal change, a definition which comes close to the original thinking of Newton and Leibniz. But Thurston goes further. Definition 5 uses symbolic notation, definition 6 takes a linear approach, and definition 7 looks at the function under a microscope. The list goes on for quite some time. Each perspective provides a new insight. Ultimately we get to definition 37: 'the Lagrangian section of the cotangent bundle that gives the connection form for the unique flat connection on the trivial R-bundle for which the graph of f is parallel'*.* To most people – even to many mathematicians – this is mumbo jumbo, but at a certain point Thurston needed precisely this definition and no other.

So what is the *real* definition of a common concept like the mathematical derivative? It would actually be a shame to have to choose just one of these many definitions. It is better to use them all and to understand the connections between them. The great value of the concept of a 'derivative' is that there are so many aspects of it. This is why it also deserves a place in our baggage.

Globe or atlas?

There is a deeper reason why core subjects have many facets. To explain this I would like to compare two models of knowledge: the globe and the atlas. The globe depicts the earth from a planetary perspective as a 'pale blue dot' in the firmament of the cosmos. The globe represents a universal definition, a concept that is both comprehensive and unique. Think of the definition of a legal term; usually it will contain not one word too many or too few. Legal experts are not asked to give a different version or to elaborate on such a definition.

The other model, the atlas, presents us with a collection of maps. Each map covers a section of the subject. There is also a line indicating how one map connects to another. To understand the whole concept we need to imagine all the maps stuck together to form a globe. An example of an atlas model is the system of world languages. Each language is a page in the atlas and dictionaries help us to translate the concepts of one language to another. Translation is just an approximation. There is no one-to-one correspondence between – for instance – Dutch words and French words. Every language has untranslatable words, and often there are several alternatives for one word in the other language. You might say that humanity has developed this atlas of languages to capture the multiplicity of phenomena in the world. Esperanto does little justice to this cornucopia.

I suspect that few concepts can be captured as aptly in one definition as the globe. For most concepts we will have to resort to a knowledge atlas.

Trial and error learning

There is another good reason for concentrating on core concepts with many facets. Because our personal knowledge base is becoming more fragile and more random, the diversity of background knowledge is increasing. Education in the broad sense will have to take this fragmentation into account. You will all be familiar with those boards for toddlers that involve putting circular, square or triangular pegs into holes of the same shape. In the past everyone had the same square holes in their boards, but now there is a huge variety of shapes. As a result there is a greater risk that a schoolteacher will try to ram a square peg – no matter how carefully that peg has been chosen – into a round or triangular hole. This is why a teacher's toolbox needs to be well-stocked as possible. One student enjoys an abstract concept, whereas another is better off with an example from everyday life. If a handyman comes to fix something for you equipped with only a hammer, you will probably feel a bit worried. Or as the saying goes: if all you have is a hammer, everything looks like a nail.

The modern dream of medical science is *personal medicine*. No two diseases are exactly the same and ultimately every patient should have extremely individual treatment, with tailor-made medication, for every medical condition. The dream of teaching should be *personal education*. No two people are the same and ultimately every child should have an extremely individual learning pathway with tailor-made teaching methods. Unfortunately, for the time being both visions will remain dreams.

This phenomenon was highlighted in a report issued by the Royal Netherlands Academy of Arts and Sciences about a controversial subject: maths at primary school, and specifically long division. In maths teaching there are two schools of thought about this: the realistic and the formal. The first tries to teach sums by using practical examples, whereas the other is based on the clear, abstract structure of numbers. The Royal Academy's study showed that neither approach was superior, but that for some children one worked better than the other. The conclusion was that for the best results the teacher should be able to use both methods.

I have had the same experience myself in lectures for large audiences like this one. Particularly if the lecture is about an esoteric subject such as the expanding universe, the Big Bang, black holes or string theory, it is not easy to make contact with the audience. I try to carry the listener with me to my own level of understanding. That means that the destination of the intellectual journey is known, but the starting point certainly is not. Everyone in the audience will depart from a different place. My own approach often consists of the trial and error method. I try to use a wide range of metaphors and analogies in the hope that something will work, because apart from the big differences in prior knowledge, there are also big individual differences in styles of thinking.

For example, there is a striking difference between people who approach things – roughly speaking – from the left or right side of the brain. People in the first category think more in terms of language and have more affinity with a linear structure, a logical algorithm which works towards the goal step by step. People in the second group think more figuratively and have more affinity with a picture or diagram. For them a picture says more than a thousand words, whereas a picture drives the first group straight into logical quicksand. Research has shown that the two kinds of people are about equally common. This alone means that a speaker has to be extremely careful about illustrating an argument with an image. One man's meat is another man's poison – something that works very well for one person will be complete gobbledygook to another.

The broad spectrum approach occasionally goes even further: sometimes listeners think up an appropriate metaphor for themselves and then make the final leap all on their own. Once a man came up to me after a lecture and said he hadn't understood it at all until I had said ... But the funny thing was that I hadn't said ... ! I would really have loved to have said ... , because it was a fantastic metaphor. This person had thought of the right tool himself. I was now able to add this gift to my toolbox and use it to help other people in the future. Physicists call this phenomenon 'stimulated emission': literally this refers the process in which an atom emits a particle of light under external influence, but here I am referring to the creativity and imagination that led a person to take the decisive step.

My advice to stop the Great Forgetting and to restore our factual knowledge is as follows. First identify the universal core concepts: ideas, facts and events that are worth remembering. Then study them thoroughly and from many angles. Take your time and revisit these subjects regularly so that they eventually become familiar friends. Approaching a subject from different points of view increases the strength of the glue, because then at some point it will 'click' with everyone, regardless of their experience and background. This common canon would then be the basis of further knowledge acquisition.

We'll look it up

Well-chosen facts are just the first step in the edifice of our knowledge. No matter how carefully we select the relevant subjects and then commit them to memory, the next step is about the procedures we can use to find additional facts. Where can we look things up and to what degree can we trust the knowledge we find in this way?

These two kinds of knowledge, factual and procedural, play strikingly different roles in public debate. We like to use the first to judge other people, while we tend to apply the second to ourselves. Politicians who do not know when Charlemagne lived or tourists caught on camera who cannot point out their holiday destination on a map are very popular with columnists. The underlying reasoning is that MPs with inadequate knowledge of history are not fit to serve the interests of the nation and that those silly tourists do not deserve to be in such faraway places and should have stayed at home.

We prefer to use other standards for our own knowledge. It's enough to know where we can find the answer. Imagine, for instance, that you are asked a question about the operation of our parliamentary democracy such as: what majority is required in the

Lower House for an amendment of the Constitution? Even if you don't know the answer immediately (and it's two thirds), you will not feel less able or entitled to vote in the next election. If this fact is relevant, you'll just look it up.

Unlike other people's factual knowledge, we think of our own in the same way we think about common sense. Questionnaires show that nearly everyone thinks they have exactly the right amount of common sense. Less would be stupid, more quite unnecessary.

For example, some inquiries in a limited circle revealed that parents attach a lot of importance to knowledge of the capital cities of Europe. But this seems to apply mainly to the *old* Europe. When asked to name the capitals of countries in the *new* Europe such as Macedonia, Moldavia or Montenegro – to restrict myself to the letter M – these same parents have to think very hard. Although we believe our children should learn these capitals at school, the older generation evidently do not feel they need to do a refresher course in geography; these countries are just too new. That's the kind of knowledge we look up on Wikipedia. (In case you're wondering – the answers are Skopje, Chisinau and Podgorica.)

Knowledge from scratch

There is another way to acquire knowledge as well as remembering it or looking it up. You can work it out for yourself.

When our family lived in the United States, we made a big impression on our American neighbours with our ability to make pancakes from scratch, that is, with eggs, milk and flour, without using a packet of pancake mix. Scientists also attach importance to building up knowledge from scratch. The main ingredients are common sense and logical reasoning. For instance, in physics there is a pleasing category of problems known as Fermi problems, named after the legendary Italian American physicist Enrico Fermi. Fermi had a habit of asking his students questions that were hard to look up and certainly not worth remembering, but that you could easily work out for yourself. Famous examples of Fermi's questions are how many piano tuners are there in Chicago, how many molecules are there in the earth's atmosphere, and how long would it take to level Mont Blanc with excavators and lorries. After some practice, a few calculations on the back of an envelope produce a good rough estimate.

The reason why physicists think being able to answer questions like these is useful is not just because they are lazy and don't want to carry too much information around with them. It also develops a skill that in a way goes beyond factual knowledge. If you know how to make pancakes from scratch, you are no longer dependent on the availability of ready-made pancake mix.

Dark knowledge

The ability to generate facts instead of reproducing them is relevant, because the majority of facts are still unknown to us. One of the essential functions of knowledge is to be aware of what you don't know. For instance, since recently cosmologists have been convinced that only 4 per cent of the content of the universe is known. That is, 4 per cent of the cosmos is made up of stars, planets, gas clouds, particles and radiation which modern astronomy has identified and documented. The other 96 per cent consists of new forms of matter and energy, called dark matter and dark energy, whose nature and composition is still completely unknown. I like to call known unknowns like this 'dark knowledge'. We know that this knowledge could exist, but we don't yet know its content. Seen in this light, 96 per cent is a very reasonable percentage. In many other disciplines, the percentage is much closer to 100 per cent.

Then there are the 'unknown unknowns' – things we don't even know we don't know. Twenty years ago these same cosmologists were not aware of the existence of dark matter and energy. The discovery of this lack of knowledge was one of the great breakthroughs in science. If dark knowledge represents that part of the room that is not yet illuminated, then breakthroughs like this stand for the discovery of a door leading to a completely new room.

Joshua Foer's passionate plea to stop the great forgetting and to give facts a prominent role again is an important signal, especially because it is the voice of a new generation. We have treated our knowledge carelessly for too long. It's time for major repairs and a fresh look at what is worth remembering and understanding.

This is particularly important because we are on the eve of great breakthroughs. The true explosion of information is probably still to come. That is both a frightening and a stimulating thought. Precisely because there are still such vast areas of the mountain of knowledge to be explored, it is essential for us to have the right baggage with us – just those carefully selected facts which form the framework of our knowledge. It is only then that we can continue this journey, which may well be the most exciting in the history of mankind.

Robbert Dijkgraaf is president of the Royal Academy of Arts and Sciences and Professor of Mathematical Physics at the University of Amsterdam.