The Royal Academy of Engineering

Philosophy of Engineering

Volume 2 of the proceedings of a series of seminars held at The Royal Academy of Engineering

The Royal Academy of Engineering

Philosophy of Engineering

Volume 2 of the proceedings of a series of seminars held at The Royal Academy of Engineering

ISBN 1-903496-78-0

October 2011

Published by
The Royal Academy of Engineering
3 Carlton House Terrace
London
SW1Y 5DG
Tel 020 7766 0600
Fax 020 7930 1549

Copies of this report are available online at:
www.raeng.org.uk/philosophyofengineering

Registered Charity Number: 293074

Table of contents

Foreword

The Royal Academy of Engineering wants to move engineering to the centre of society by highlighting the crucial role that engineering plays in shaping our lifestyle and culture. The contribution that engineering has made to intellectual history is central to this role. Engineering has had an enormous impact in developing tangible benefits from the complex body of knowledge that humanity has developed. Relativity theory and Darwinian natural selection might be cited as pinnacles in the ever-progressing ascent of human knowledge, but we should add the development of the World Wide Web and space exploration as examples of the awe-inspiring level of understanding that has been reached. Engineering has made an overwhelming contribution to our understanding of the way the world works and how to make the world work for us.

One of the main aims of this seminar series was to develop an appreciation of the nature and role of engineering knowledge. There is great value in developing a better understanding the nature of engineering knowledge and of engineering itself. This allows us to raise the profile of engineering by demonstrating its role in developing sophisticated knowledge. This can bring a sharper understanding of engineering method, which can be of great value to engineering education. It can also enable a clear formulation of what engineering *is*, in order better to convey its value through public engagement.

As well as a better appreciation of the nature of engineering, the series aims to show that there is much fuel for philosophers if they look to engineering for examples. Philosophers may find in engineering enlightenment on the kinds of questions that they have struggled with for centuries and no doubt philosophers will also find new issues to engage them. Engineering work on artificial intelligence and information technology can, for example, enlighten the philosopher's questions about the nature of thought, consciousness and language. The engineering process of synthesis and construction can inform metaphysical questions about what the world is made of, how it can be broken down and what its fundamental elements are.

There is, however, no point in engaging in a philosophy of engineering unless it has a use - no engineer embarks on a project unless there is an end purpose for what they are working on. The objective of this series is to demonstrate the complexity and richness of engineering and the extent of its influence on human progress. It can be used to send out the message to society that engineering is an important, rewarding and worthwhile profession. In addition, the skills of philosophers in constructing and delivering clear arguments could be of great use to engineers. If philosophy of engineering can help to cultivate such skills in engineers, then engineers will have a stronger voice with which to convey that message.

Engineering is a broad, interdisciplinary field and has links with the social sciences and humanities as well as the natural sciences. The basic aim of the seminars was simply to get engineers and philosophers together to share ideas and to identify research areas of common interest. The Royal Academy of Engineering hopes that this will be the beginning of a fruitful collaboration and that this, and the first volume on philosophy of engineering, will provide food for thought for philosophers and engineers alike.

Dr Keith Guy FREng
Chair of the philosophy of engineering steering group

Introduction

This is the second volume of a series of papers presented at The Royal Academy of Engineering, for its series on philosophy and engineering. The first part of this volume tackles some deep philosophical issues on which engineering has a bearing. These include the nature of mind and language; conceptual distinctions in science, engineering and common sense; engineering's influence on the fundamental categories in nature, and the nature of knowledge.

In 'Philosophical engineering', Nigel Shadbolt argues that our modern world is distinguished by a coming together of science and engineering. Questions that we once attempted to understand using pure reason and philosophy are now addressed by empirical, scientific methods. The science of thermodynamics allowed a radical breakthrough in technologies from the internal combustion engine to the refrigerator. The science of aerodynamics enabled enormous advances in the design of aircraft. Sometimes the engineering preceded the science and on other occasions the science was needed before the engineering could commence. However, in the modern era there has always been an intimate connection between the two. His paper examines how this connection can be seen dramatically in the digital realm. In particular, the World Wide Web has produced an information fabric that requires analysis to understand its properties and engineering to ensure its effectiveness. The Web also embodies real human knowledge in a way that suggests we are entering a new era – an era where questions that were essentially philosophical in nature are appearing as engineering challenges. How can meaning be characterised and shared between computers on the Web? Can intelligence emerge as a consequence of large number of interactions and transactions on the Web? These and other questions are reviewed together with the likely future development of a new discipline.

In 'When water does not boil at the boiling point', Hasok Chang reveals the surprising finding that the boiling temperature of pure water under fixed pressure is not constant, but is dependent on the nature of the vessel in which the boiling takes place, and even more on the amount of water. These variations were well known to physicists and chemists in the late 18th century and throughout the 19th century, but seem to have been largely forgotten in modern science. Chang shows that in contrast, mechanical and chemical engineers have preserved and developed this knowledge, at least concerning the effects of surface quality. In standard thermal physics the presumed ontology of boiling is encoded in the sharp liquid-gas boundary in the phase diagram. For engineers working on heat transfer, the "boiling curve" has "surface superheat" as its main independent variable, which at the outset abandons the assumption of a sharp temperature at which the phase transition occurs. The two frameworks of analysis are incommensurable with each other, which raises interesting questions about the relation between science and engineering, and between engineering and our everyday experience of the world.

In 'Ontology in engineering', Peter Simons describes how engineering and philosophy share breadth, abstractness and a reflexive attitude. He explains that while the uses of philosophy for engineering are most obvious in ethical issues, ontological analysis is a potential source of mutual help and insight for philosophy and engineering. Ontology is the maximally abstract, go-anywhere theory of everything. It aims at a categorial framework for any subject matter. Though philosophers trade in its internal disputes, Simon argues that, away from the specialist journals, ontology can provide relatively neutral analyses of concepts and objects fundamental to engineering: part/whole, structure, function, life-cycle, emergence, process and product, needs and requirements, success and failure, design and planning. Negatively, this can prevent conceptual foul-ups, whether informal or as enshrined in IT models. Positively, it can enhance conceptual transparency and inform tools for managing complexity, and can thereby be a valuable tool for engineers.

In 'Constructionism, the maker's knowledge tradition and knowledge engineering', Luciano Floridi argues in favour of a constructionist approach to knowledge. Put simply, this is the view that an epistemic agent knows something only if that agent is able to build that something. Alternatively, an agent qualifies as an epistemic agent when she is not just a mere user but a producer of knowledge. The maker's knowledge is knowledge of the ontology of the artefact. Floridi gives a broad account of what is meant by constructionism, with some references to the philosophical tradition that has inspired it, the so-called "maker's knowledge" tradition. He then shows how the tradition applies to computer science, arguing that constructionism is a form of knowledge engineering, showing finally how the Turing Test can be correctly read as an application of the constructionist method.

The second part of this volume highlights the critical importance of the responsibility of engineers to consider the effect of their work on the environment. Climate change has meant that the wider impact of technology on the environment cannot be ignored. However, protecting the environment is a complex matter involving balancing different values and priorities.

In 'Engineering sustainability: Synergies between science, precaution and participation', Andy Stirling explains how risk management in engineering is conventionally based on quantitative expert techniques, often referred to as 'science-based' risk assessment. These methods are typically presented in stark contrast to what are often held to be relatively ambiguous, unscientific (and even 'politically correct') 'precautionary' or 'participatory' approaches. He argues that, likewise, current high-level debates in areas like energy policy and nano-engineering tend to view innovation as a matter of linear technological progress. Engineers have in the past characterized innovation as a 'race' along a path that is essentially predetermined by science. In this view, precaution and participation can again appear as distractions or obstacles. Drawing on recent interdisciplinary understandings of risk and uncertainty and the nature of the innovation process, Stirling's paper critically examines these well-established assumptions and explores some practical policy implications. Far from being in tension with scientific rigour or technological excellence, precaution or participation are argued to provide necessary elements of a more robust approach to the governance of technology.

John O'Neill, in 'Engineering and environmental values', notes that the use of the concept of engineering in the biological sphere often elicits strong negative ethical responses. Practices such as 'bioengineering' and 'genetic engineering' are criticised on the grounds that they are 'unnatural' or involve humans 'playing god'. O'Neill's paper examines whether there is a reasonable basis for such responses. He considers some parallels that can be found between arguments against 'bioengineering' and arguments that have been articulated in political theory against the concept of social engineering. He also raises questions about the assumptions about the nature of engineering that both sets of argument involve.

In 'Children of Martha: on being an engineer in the environment', Roland Clift argues that engineers have a different relationship to environmental issues than scientists: engineering implies a commitment to action rather than analysis. His paper focuses on the responsibility of the individual engineer rather than the profession as a whole. The resulting concept is that the engineer should be a technical expert who acts as a social agent, not merely a technician.

Part 1: Engineering, Metaphysics and Mind

1.1 Philosophical Engineering

Professor Nigel Shadbolt

Professor Nigel Shadbolt FREng is Professor of Artificial Intelligence in Electronics and Computer Science at the University of Southampton, where he is Head of the Web and Internet Science Group. He is currently researching the next generation of World Wide Web methods, technologies and standards. He was one of the originators of the interdisciplinary field of web science. Between 2000-7, he was the Director of the £7.5m EPSRC Interdisciplinary Research Collaboration in Advanced Knowledge Technologies (AKT). AKT was particularly influential in establishing the viability and value of web-based semantic technologies. He has also been heavily involved with the commercial exploitation of his research. In 2006 he was one of three founding Directors and Chief Technology Officer of Garlik Ltd, a company specialising in consumer products and services to put people and their families in control of their own digital information. In 2009 the Prime Minister appointed him and Sir Tim Berners-Lee as Information Advisors to the UK Government to transform access to Public Sector Information. This work led to the highly acclaimed data.gov.uk site that now provides a portal to over 7500 datasets. In 2010 the coalition government appointed him to the UK Public Sector Transparency Board that oversees the continuing release of government data. He also advises government in a number of other roles. In its 50th Anniversary year 2006-2007, Nigel was President of the British Computer Society. He is a Fellow of the Royal Academy of Engineering, and the British Computer Society.

Professor Shadbolt's background originally was in philosophy and psychology, and it was the questions that philosophy posed that really drew him ultimately into trying to think about modelling intelligence and the nature of intelligence. This led him into artificial intelligence and a career now involved with a variety of forms of engineering.

The rise and rise of science

We should remember that not many centuries ago much of human knowledge was gathered by methods of reflection and introspection. The organised characterisation of bodies of knowledge, as exemplified by the scientific revolution, started to get under way when people began to ask questions that could admit to empirical falsification. Hypotheses were formed following the collection of data; and empirical methods were employed, backed by theory. Copernicus, for example, was recognised for his attempt to use detailed recordings and from these to try and understand how they might change our understanding of, in his case, the movement of planets. In fact, his empirical methods were somewhat imprecise, and although he did a great deal of outstanding observational astronomy, he also based some of his findings on classical results that contained errors (Moesgaard, 1973).

One of the most striking examples of the move towards trying to explore and understand the world in terms of systematic investigation was Vesalius, the Flemish physician and anatomist who produced a set of recordings and a quite wonderful description of the structure and function of human anatomy (Vesalius, 1568). His work really did reformulate our understanding of a whole branch of the natural world.

We are all aware of Galileo's work, particularly in mechanics. His work on ballistic trajectories was a piece of applied science, but he was also capable of deriving all sorts of interesting proofs. His method of proof often involved geometric and diagrammatic methods (Drake, 2003), so his ways of coming to understand the world relied on representations and languages quite different to those we are now used to. Isaac Newton, of course, was a wonderful expositor of both theory and its embedding in practical empirical observation.

This is how science came to grapple with the questions that had previously existed only in the philosophical realm and gave these disciplines a firmer footing, where questions could be addressed using new methods and techniques. It wasn't obvious that systematic empirical analysis and experimental refutation of hypothesis could be applied to such questions; it took these huge figures to establish the importance of the scientific method.

But as knowledge became systematised it is interesting to see how the iconography of science has developed – those scientists we can actually recognise from their images. We would unfailingly recognise an image of Albert Einstein. Alan Turing the great mathematician and father of AI would be recognised by many. And so too the Nobel laureate Richard Feynman, in a beautiful example of the scientist applying his mind to solving an engineering riddle. This was the Challenger disaster; where he demonstrated, to the astonishment of the audience, that O-rings became rather brittle when subjected to low temperatures by dropping a piece of O-ring into his glass of iced water, demonstrating beautifully the power of simple provable demonstration.

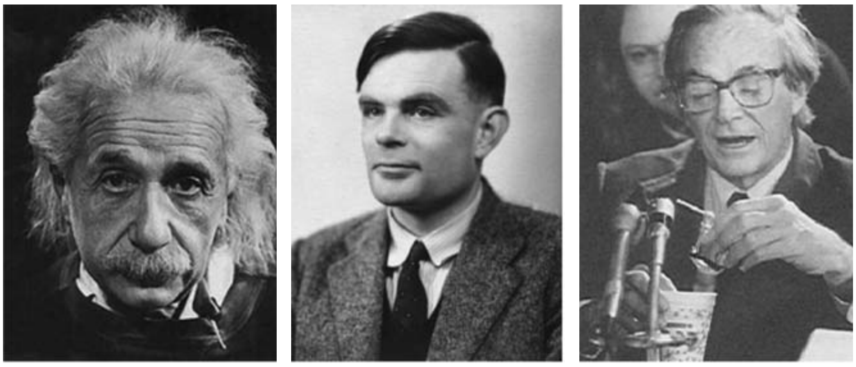

Figure 1. Photographs of Albert Einstein, Alan Turing and Richard Feynman

Female scientists may not be as well recognised but they have played a huge part in the rise of science – the three below all having won the Nobel prize. Madame Curie is recognisable, as the iconic scientist and winner of two Nobel prizes. But I fear not so readily recognisable is Dorothy Hodgkin who used X-ray techniques to determine the structure of man biochemical structures including penecillin, or Maria Goeppert-Mayer who first suggest the nuclear shell model of the atomic nucleus.

Figure 2. Photographs of Madame Curie, Dorothy Hodgkin and Maria Goeppert-Mayer

The ascent of engineering

In the history of engineering, it is not so much the individuals who are recognised but the devices they created. Engineering emerges rather through the tradition of artisan craft, in the ability to construct useful and exquisitely engineered devices. The reason I make this point is that I suspect one of the reasons that philosophy has found it hard to find a place in engineering is that it is rather unlike science, which became such an iconic activity that it deserved a piece of philosophy associated with it. Engineering was about things and devices and philosophers may have wondered whether they were an appropriate subject for philosophy.

…into celebrity?

However, in the Victorian era, engineering stood in some danger of becoming iconic. The Industrial Revolution made celebrities of engineers like Brunel with his bridges and his steamships. But usually it is the device that dominates. The jet engine, computer and mobile phone have changed our world out of recognition. I am not so sure we would be able to recognise their orignators - figures such as Frank Whittle, Presper Eckert, or Martin Cooper. The jet engine, the computer and the mobile phone have transformed of the world we live in. Yet, despite this, a serious belief that these devices and the processes associated with them deserve philosophical investigation has, I would say, been rather slow to arise.

Scientist or Engineer?

The exception is AI, which is interesting. This may be because it gets to the heart of some challenging questions around metaphysics. More recently people have begun to wonder if the web can be called "intelligent" in any meaningful sense. Tim Berners-Lee is constantly asked, and asks himself, whether the work he does is science or engineering. I submit that when we look at the modern information fabrics we have built in computer science, we are talking about philosophical questions, because perhaps rather more directly than the jet engine, or indeed the mobile phone, our information processing systems do something very interesting with part of the core content of philosophy.

It also needs to be said that science and engineering are two faces of the same ingenious creative capability in humankind. Joseph Black's theoretical work on latent heat clearly laid a foundation for refrigeration engineering. In a very similar way, Page and Brin's Google page-rank algorithm is based on some rather smart mathematical analysis. It is based on the understanding that the Web of links can be characterised as an eigenvector that can be understood as a relevance vector that shows how the most important pages on the web could be calculated from their dependence on other important pages. To get a convergent algorithm where that happened required some insightful analysis. It also required some innovative engineering to build the network of servers that could hold the indexes and distribute them. Here as in so many cases one sees an interplay between these two faces of science and engineering.

Science and engineering: a web science perspective

One of the things we had been thinking about in the Southampton/MIT Web Science Research Initiative is that considering web science as either engineering or science is too restrictive. In the information space that is the web, we typically find the cycle of an idea that is born, a technology to support it, and a social dimension that makes it take off. In fact, typically you have to start with an issue, which in the case of Google was how to find relevant pages in billions of pages on the web. The idea was that you could formulate the importance of pages and recursively build an algorithm to rate importance across the web. The technology, the design, was a very smart insight as well. The social element that ultimately allowed the phenomena to get a grip was to add link incentives, ultimately using advertising. The early social phases of Google were simply academic search requirements at Stanford when the early Google systems were being revealed. From this grew the idea to index the web, to build support vector machinery, to use an eigenvector algorithm to understand the link structure, then implementing and incentivising this system by auctioning search terms used in Google in a market; from this you get the Google phenomena.

On the web, we see properties that we don't entirely understand. In particular, where micro principles give rise to emergent behaviour, or macro phenomena. The idea that the macro phenomena arise from micro behaviour is a beloved idea in science. One of the things observed repeatedly in the web is, if an application or idea is really successful you know because people start to spoof it or spam it, trying to take advantage of its massive outreach across the planet. For example, large amounts of the web are automatically constructed to try and convince the search engines and the software robots out there that these are legitimate pages, where in fact they are plants – automatically built - to try and boost the particular rankings of people's interests, whether it's pornography, services or products.

This happened when email use took off; it has happened in blogs also. There are around 120,000 blogs created a day, and about 7000 of them are so-called 'splogs' or spoof blogs. Fake blogs are created because, increasingly, bloggers are also doing advertising-based sponsoring of content on their blog, so there is an advertising model behind that too. It fits a very similar model to our Google analysis. To try and understand why blogs took off just about six years ago you have to look back to some very simple technological innovations, such as the track-back. The track-back was a piece of internet protocol that allowed anybody who used somebody else's blog reference to ping back to the originating blog to say 'I'm referencing you, you can make a link to me'. The track-back allowed the construction of a reference network very quickly. This social citation process was at the heart of the blog's success as a piece of technology.

Understanding the behaviours in this space is part of the challenge of web science, an activity that has gone on for some time, though not explicitly under that name. Part of our initiative is a call to action to get people to understand that, in taking the web as an object of study, it has engineering, scientific and social aspects that need to be understood. So why call it web science? In truth we could not quite bring ourselves to call it philosophical engineering or webology. But what I have wanted to emphasis is that the distinction between science and engineering is not a dichotomy but a complementarity – and in the case of the web and web science we absolutely need both perspectives. And so far this has been essential. Take for example the shape and structure of the web - the topology of the web.

Early results were genuinely surprising findings, for example: the scale-free nature of the web; that it's not normally distributed; the link structure; and that some nodes are massively over-represented in terms of their link structure than others, which means the web doesn't look like a random network. It means, for example, that it's susceptible to attack – if you take out the hubs you can fragment it. In contrast, a random network is pretty tolerant to about 85 percent removal of the nodes; you will still often find a path in a web of a certain size and complexity.

This analysis and understanding of topology has been an area where the mathematicians and physicists and the web engineers have been busy trying to understand its properties. We believe there is a whole slew of analogous areas where we need to understand both how to engineer for scaleable systems and to understand the scientific properties thereof. I would call this the continuation of science and engineering by other means. Where is the philosophy in this?

The web: socio technical and emergent

I now turn to the particular research area that I've been working on for the last few years: the semantic web. Remember whenever you see the word website, it has a deeply socio-technical and emergent characterisation. If you look at all the web 2.0 developments in the blogosphere, such as Wikipedia, Second Life, Facebook, and so on, you are seeing new social practices which take advantage of very simple engineered protocols, which when they have scaled to millions of users generating millions of pages of content then give rise to interesting analytic questions: for example, can you track the blogosphere? A whole area of technology, Web analytics, looks to see whether it can analyse the tens of millions of blogs and detect new ideas being promoted: for example, how far the iPhone is being talked about in this massive online distributed set of individual commentaries. Can you detect the memes or trends or conversations that are emerging here? There is interesting work on just these sorts of questions.

Making the web semantic

It is worth stating how simple an idea the semantic web is. It's not about smart agents taking over the web and doing your bidding unaided. It's about taking metadata standards and making them a little richer. So when we look at a web page, for example a conference web page (Figure 3), four billion years of evolution in our eyeballs can make interesting inferences about it. We can identify the fact it relates to a conference, and conferences come along with all sorts of standard expectations: they have a location, a start date, an end date; they'll have people register for them; they might have pictures of keynote speakers.

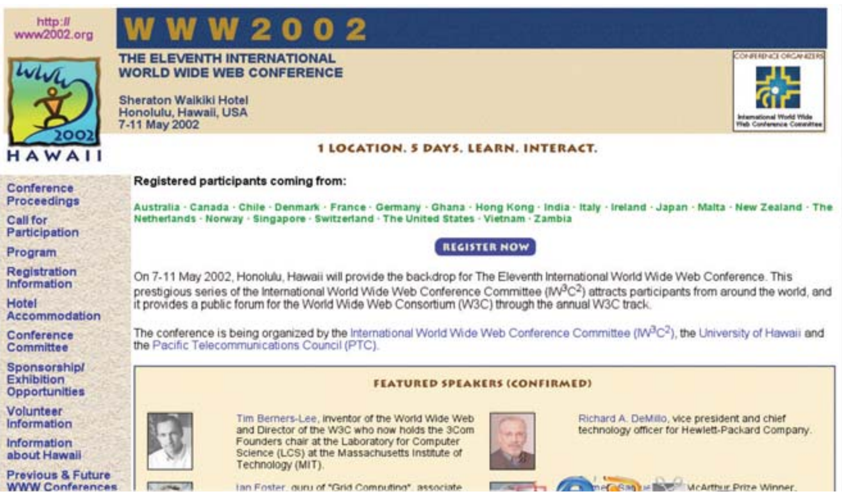

Figure 3. A web page advertising the eleventh international World Wide Web conference

This simple description of objects and their interrelations could be discussed in any kind of AI class on knowledge representation. It would suggest that a simple object representation would be the thing to do here. You could capture some of the knowledge that is in the page, and in a way the early semantic web is all about trying to produce very simple internet protocols that capture in machine-readable ways those simple object classification schemes and simple relationships and expected properties. That's all it is. It is about using the simplest AI engineering representation language, a triple-based language which is essentially a subject-verb-object kind of structure, but each of those elements in those subject-verb-object sentences points to a resource on the web. It is a pointer, a URL, if you de-reference it you will get to the originating object, the piece of data. The semantic web is about turning a web of documents into a web of data. A web page with this sort of annotation can be searched and queried by smart programs in a different way. And if all the pages of conferences on the web used this simple standard we would have new emergent possibilities. I could search the web of all the conferences that occur between one date or another, or which feature a particular keynote speaker, or are about a particular topic in a specific region etc.

Figure 4. The eleventh international World Wide Web conference web page with the object classification schemes highlighted

In other words, it's about developing standards, languages that enrich. That page without the semantic metadata would simply tell you where to split the paragraph, what font size to use, what colour to use, where to insert the photograph. The semantic additional layer is telling you, go to this page and look up the conference ontology – ontology is just an organised classification of terms and relationships – and if you want to look up the little mini-ontology about academic conferences it will tell you what objects to look out for, what properties those objects have. You can then look in that web page and understand what RDFS subclass publication means. A publication is a kind of thing that will be associated with a speaker at a conference, for example.

The whole point about the semantic web is you can annotate anything. It's not just web pages that will have all this web standard metadata language, and there are now web standard adopted metadata languages to describe the semantic content. You can do it with publications, databases, scientific structures and so on. Even people, patients, as we will see in an example I will describe later.

Sharing meaning on the web: ontologies

In the engineering community 'ontology' refers to an agreed terminology, an agreed conceptualisation of objects and relationships that have understood properties. There are communities of practice who need to share information at a richer level than just describing the superficial properties of a domain, or who need to avoid years of work to construct an information exchange data model between one relational database and another. The ontology expressed as a web standard is a lightweight way of establishing a data exchange format. So in the life sciences, in proteomics, genomics, where companies would like to share some of their data across the web, the concept of an ontology is starting to feature strongly. We are seeing it in e-health, we are seeing it in manufacturing and e-defence.

Realist stance

We are not completely naïve about our sense of meaning here. The final phase of the talk is to try and explain why I think there is something very interesting happening here that has quite deep philosophical roots. When you look at the kind of semantic languages that are at work on the web or the semantic web, they embody what people would in philosophical terms refer to as a 'realist' or 'positivist' tradition. It is the idea that our language engages directly with reality, that there is a simple one-to-one mapping between objects of the language and referents (objects denoted) out there in the world, whether these are missiles or aircraft, genomic or chemical structures. That tradition has a noble tradition in philosophy back through Aristotle, Leibniz, and the early Wittgenstein. It is the idea that you can build a linguistic description of the world that stands in a one-to-one correspondence with the physical world.

This is a powerful idea and has been a core idea in philosophy. In the late 19th century, people like Frege, Wittgenstein and Russell were heroically trying to capture significant amounts of the meaning of language in terms of formal logical calculi. Could you characterise the essential meaning of terms using a logic-based approach, because the language was then going to stand in a very direct relationship between terms and objects in the world? I remember slaving over Russell and Whitehead's Principia to appreciate the power of that approach, and theirs was an heroic endeavour.

The interesting thing about that tradition is that when computers emerged and could carry programmable language instructions, one of the places people looked for an interpretation or a semantics of what was intended in the programme symbols, they used the work of Alfred Tarski. Tarski semantics, for many classes of theoretical computer science and programming semantics, is seen as a way of rooting what you intend by your symbols to some representation of the world, or the world you are trying to model.

That is at the heart of one strong tradition in AI. Nils Nilsson is one of the most famous exponents in AI of what is called the logicist tradition. A model theoretic approach says, here is a logic, this is what you're going to model in the world. It's an abstraction, it's an approximation but you're clear about what it means.

The semantic web has elements of this approach in it. In fact, Ian Horrocks and others who build description logic inference engines for the semantic web intend to inject a huge amount of denotational meaning into new kinds of logics, which are going to provide the semantics for these web semantic languages. But when you actually build a real semantic web application, it isn't just like this.

Constructivist stance

There is another stance, called constructivism: it alleges that the positivist, realist stance is a misconstrual. Constructivism argues that language is about a complex set of norms and social conventions. This tradition includes Husserl and Heidegger, and the later Wittgenstein, who came to reject all of his earlier views about how you could use logic as an appropriate characterisation of meaning and language. This constructivist view, that understanding meaning is about understanding rules of use, the context of use, also has its adherents in AI.

In the early days of AI, people like Terry Winograd were building programmes where the only way you could understand the meaning of the programme was to understand the procedures it executed. In fact, Winograd went on to become a supporter and admirer of this hermeneutic tradition, the tradition that language is a practice.

In the semantic web there are emergent signs of this. If you look at the modern penchant for tagging on the web, there are very large numbers of people associating terms with objects. There is no precise, agreed semantics, rather the overall statistical distribution of people's choices give you things called tag clouds, which are representative of people's associations between words and objects. They don't claim to capture everything, however.

Autobiographical thread – recurring theme realism and constructivism

My own journey through all this has always equivocated between a desire to be very neat and clean and logical, through to the recognition that the real world is always a more complex and more constructed space. In my PhD work (Shadbolt, 1984), on the one hand I was modelling formal systems using logics like Montague Semantics, and on the other hand looking at real dialogue and realising these two just didn't line up. In expert systems you could pretend that you could characterise rule-based reasoning, but when you looked at the richness of expert practice you realised how exquisitely tuned and modifiable real knowledge was. In the semantic web the fundamental use of machine-readable metadata looks to be in the logical positivism, tradition and some people say you can't possibly work on that basis because you can't have one language for all the concepts you need to embody on the web. In fact, very few people claim that there is going to be one monolithic ontology that is going to describe everything of interest on the semantic web; it requires small communities of interest to build local agreements. So that the semantic web is comprised of many ontologies that represent different and distinct vocabularies to describe and conceptualise the world. Moreover, what is happening more recently is that the whole activity is becoming, in some sense, more statistical.

MIAKT – realism meets constructivism

To give you one concrete example before I end, we have been doing some work with Mike Brady's group at Oxford on medical images and signals, in the area of breast mammography. We were looking to build an information management system using semantic web approaches (Dupplaw et al; 2009). In an area like this you have multiple experts, with different conceptualisations of the world. An expert in X-ray imaging and an expert in ultrasound and an expert in MRI will use different vocabularies, they see the world differently, they don't always have the ability to translate between these vocabularies. There are real challenges, and one ends up building multiple ontologies in a domain like this.

Our approach was to build a set of web services, some of which embodied intelligent services like image registration, segmentation, finding areas of possible concern, creating a simple patient record management system and so on. You could do this by building an elegant hierarchy of objects, the kind of structure built in the era of classic expert systems: disease type, subtype and so on, and properties associated with the diseases, and it would look clearly Tarskian. You will have divided the world up into a set of subcategories that maps into some set of external phenomena you observe. But when you really look into what happens in a medical context, those classes are actually constructed using all sorts of social institutional norms. Some health services classify their diseases according to their fundamental economic model of how they operate. For example, in the US, whether insurance companies will pay claims against illnesses is one of the determinants as to what objects are discerned, what divisions are made in the disease hierarchy. In the protocols that medical professionals use to identify objects of interest, they have to follow quite complex institutional norms. So you can look at a hierarchy of disease and say it looks like there is this object and this sub-object, but when you really look they have this more constructed character (Hu et al., 2007).

There is a particular object that is looked for when trying to diagnose malignancy in breast disease and the NHS has guidelines about what to look for to identify a particular kind of cell. There is a particular problem with the type of cell that is often a false positive; it is easy to identify it as being malignant when it is benign. The process for establishing whether or not you are likely to be seeing a malignant condition involves a detailed set of staining procedures and reproducibility criteria, going back and looking at the sample again.

If you start to fiddle around with those procedures and protocols, you get no stable intuition about whether what you're seeing is of category x or y. A slight departure from the protocol in use at a given hospital would reduce both between-expert agreement, and even an expert's own ability to tell whether he is looking at one class or another. So the reality is that, unsurprisingly, our way of understanding the world is massively conditioned by the rules of use and the procedures we adopt.

Norms, indeterminacy and vagueness

That insight provides a quite deep potential insight for philosophy. One of the reasons that the philosophers left behind the logical positive view of meaning was that humans were clearly more complex than it suggested, and meaning instantiated in language seemed much more complex than a simple logic-based denotation account. In practice, one of the things that has been remarkable about the engineering of computers and information systems is that we have found ourselves in a situation where we actually built machines which were premised on a logicist view of how symbols could engage reality, so we have built devices essential to our wellbeing on which that is the epistemological basis for agreement. We then come along as the users of these systems with a much richer conceptualisation and are able, in some sense, to add a richer layer of understanding and appreciation, discounting in some cases, sometimes being critical of, or questioning of, the allocations made.

When you try and put contentful meaning-based systems to work you see that disciplines, whether in e-science or e-health or e-commerce, do not have clean boundaries. It is hard to decide just where your domain of discourse ends and somebody elses takes over. This was true before the web, but it becomes very apparent when you see it happening at web scale. These ideas have been around and interested philosophers for a long time; the idea that you can't write down a definitive set of procedures for all time because new understanding, new insights, may change your fundamental notion of what objects to pay attention to, what objects matter. Modelling and abstraction are rampant through these application areas, and when people get into these systems they reinterpret their practice and their conceptualisations through time. Their view of how they understand this material changes over time too.

Summary

In summary, the large metaphysical questions remain, questions such as what it is to be in the world and of the world. However, I contend that our science and engineering are taking some questions which were clearly originally largely philosophical in character into a practical context. Engineering is having an impact on how we understand what it is for objects to contain meaning and from devices that now support us in so much of our problem-solving, our machines, our information fabric, questions arise about what notion of meaning they embody. As our science and engineering evolves, new philosophical positions and possibilities emerge and a new kind of scientific enquiry can be born too. This is particularly the case with web-based knowledge and semantic-based processing.

Christopher Longuet-Higgins once said that AI is the pursuit of philosophy by other means, and that is quite a nice idea to carry away. I think that science and engineering of the web, or AI or information, are new arenas for the pursuit of serious philosophy by other means.

References

Drake, S. (2003) Galileo at Work: His Scientific Biography. Dover.

Dupplaw, D., Dasmahapatra, S., Hu, B., Lewis, P. and Shadbolt, N. (2009) A Distributed, Service-Based Framework for Knowledge Applications With Multimedia. *ACM Transactions On Information Systems*, 27, 4.

Hu, B., Dasmahapatra, S., Dupplaw, D., Lewis, P. and Shadbolt, N. (2007) Reflections on a Medical Ontology. International Journal of Human-Computer Studies Special Issue on Ontologies.

Moesgaard, R. S. (1973) Copernican Influence on Tycho Brahe. In Jerzy Dobrzycki The Reception of Copernicus' Heliocentric Theory. Springer.

Shadbolt, N. (1984) Constituting Reference in Natural Language: The Problem of Referential Opacity. PhD thesis, University of Edinburgh.

Vesalius, A. (1568) De Humani Corporis Fabrica. Franciscum and Criegher.

1.2 When water does not boil at the boiling point

Professor Hasok Chang

Hasok Chang is Hans Rausing Professor of History and Philosophy of Science at the University of Cambridge. At the time of the seminar he was Reader in Philosophy of Science at University College London, where he taught for 15 years until August 2010. He received his PhD from Stanford University in 1993 with a dissertation on the philosophy of quantum mechanics. His current research focuses on philosophical analyses of historical developments in physics and chemistry in the 18th and the 19th centuries. He is the author of Inventing Temperature: Measurement and Scientific Progress (Oxford University Press, 2004) and Is Water H2O? Evidence, Realism and Pluralism (Springer, forthcoming), and co-editor (with Catherine Jackson) of An Element of Controversy: The Life of Chlorine in Science, Medicine, Technology and War (British Society for the History of Science, 2007).

The metaphysics of boiling?

This paper is about ontology, dealing with a specific issue in the realm of material ontology. Specifically, it is about boiling. In the course of research I was pursuing in the history of thermometry, I discovered that the boiling point, contrary to common belief, is not really constant, even under fixed pressure. I learned this from historical sources and then I went into the lab to test out whether what these 200-year old sources said was true and, for the most part, found that it was. Then, trying to understand how this happens, I was guided to the engineering literature rather than the physics literature, with which I was quite familiar. I discovered that there was a very great difference between the ways that physicists and engineers looked at the phenomenon of boiling.

The physicist's phase diagram and the engineer's boiling curve

I will begin by giving a very graphic illustration of that difference and then I will go back to tell the story from the start. Here is the physicist's typical phase diagram, which you will all have seen, separating the solid, liquid and gas phases (Figure 5). There is a sharp line separating the liquid and gas phases, which defines under fixed pressure the temperature at which boiling takes place.

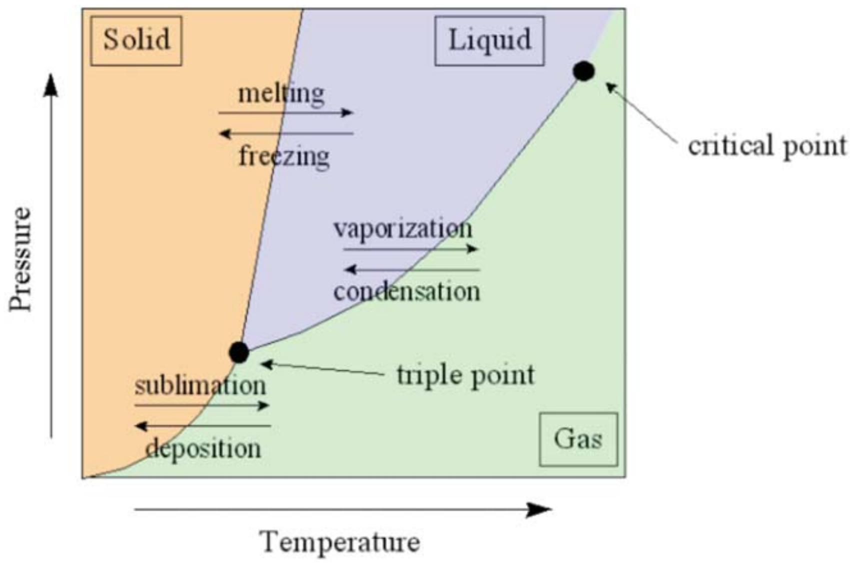

Figure 5. The Phase Diagram

Then, in the engineering textbooks on boiling, I learned about what they call the boiling curve. (Figure 6)
This generally appears in the context of chemical or mechanical engineering, on the topic of heat transfer rather than phase transition. The engineers plot the rate of heat transfer against what they call the surface superheat. Surface superheat is the excess of the temperature of the heating surface over the usual boiling point. In the physicist's phase diagram there is no equivalent thing.

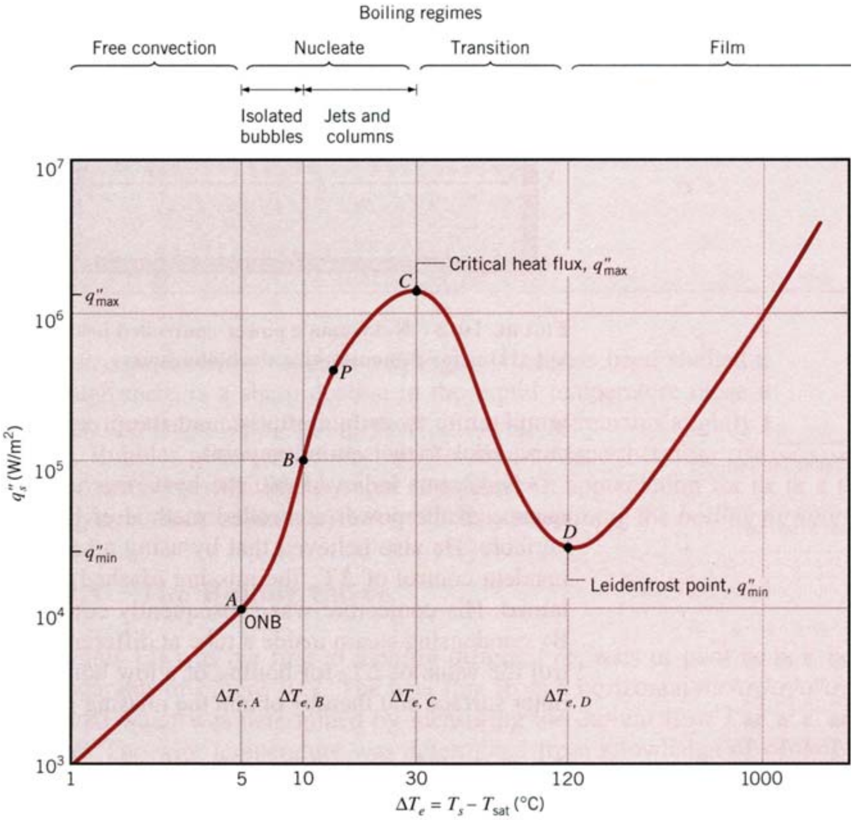

Figure 6. The Boiling Curve

I learned about this in the course of researching my book, *Inventing Temperature*. It was about 10 years ago that I first stepped into the area of temperature when, having done my PhD on the philosophy of quantum mechanics, I got fed up with talking about things which people said they did not understand at all. I therefore decided to go into something really easy—everyone knows what temperature is in some sense, and they at least know how to operate a basic thermometer to measure the temperature of something. I decided to talk about something like this and make it easy on people.

The origins of thermometry

Thermometry began with the finding of so-called fixed points – reaching the boiling point is the chief one, and everybody knows that boiling happens at a fixed temperature. Here is a passage from a curious little book, *The Complete Idiot's Guide to Economics* by Tom Gorman, where the author is trying to draw a contrast between social sciences, where things are uncertain, and physics. He says,

> "The findings and knowledge produced by a social science generally cannot be as exact or predictable as those of physics or chemistry."

As an example, he pulls this up:

> "For instance, if you put water in a saucepan on a stove, you know with certainty that it will boil when it reaches 212 degrees Fahrenheit."

However, this is not quite what my historical work was revealing. Here is a table listing all sorts of different phenomena which the early pioneers of thermometry had used as fixed points (Table 1).

Person	Year	Fixed Points ("and" indicates a two-point system)	Source of information
Sanctorius	c. 1600	**candle flame and snow**	Bolton 1900, 22
Accademia del Cimento	c. 1640?	**most severe winter cold and greatest summer heat**	Boyer 1942, 176
Otto Von Guericke	c. 1660?	**first night frost**	Barnett 1956, 294
Robert Hooke	1663	**freezing distilled water**	Bolton 1900, 44-45; Birch [1756] 1968, 1:364-365
Robert Boyle	1665?	**congealing oil of aniseed, or freezing distilled water**	Bolton 1900, 43
Christiaan Huygens	1665	**boiling water, or freezing water**	Bolton 1900, 46; Barnett 1956, 293
Honoré Fabri	1669	**snow and highest summer heat**	Barnett 1956, 295
Francesco Eschinardi	1680	**melting ice and boiling water**	Middleton 1966, 55
Joachim Dalencé	1688	**freezing water and melting butter, or ice and deep cellars**	Bolton 1900, 51
Edmond Halley	1693	**deep caves and boiling spirit**	Halley 1693, 655-656
Carlo Renaldini	1694	**melting ice and boiling water**	Middleton 1966, 55
Isaac Newton	1701	**melting snow and blood heat**	Newton [1701] 1935, 125, 127
Guillaume Amontons	1702	**boiling water**	Bolton 1900, 61
Ole Rømer	1702	**ice/salt mixture and boiling water**	Boyer 1942, 176
Philippe de la Hire	1708	**freezing water and Paris Observatory cellars**	Middleton 1966, 56
Daniel Gabriel Fahrenheit	c. 1720	**ice/water/salt mixture, and ice/water mixture, and healthy body temperature**	Bolton 1900, 70
John Fowler	c. 1727	**freezing water and water hottest to be endured by a hand held still**	Bolton 1900, 79-80
R. A. F. de Réaumur	c. 1730	**freezing water**	Bolton 1900, 82
Joseph-Nicolas De l'Isle	1733	**boiling water**	Middleton 1966, 87-89
Anders Celsius	by 1741	**melting ice and boiling water**	Beckman 1998
J. B. Micheli du Crest	1741	**Paris Observatory cellars and boiling water**	Du Crest 1741, 8
Encyclopaedia Britannica	1771	**freezing water and congealing wax**	Encyclopaedia Britannica, 1st ed., 3:487

Table 1. Summary of fixed points used by various scientists.

You can imagine the philosophical problem here, where they were trying to judge which natural phenomena always occur at the same temperature, before they had any thermometers. How would they determine that? All kinds of phenomena were imagined to be constant in their temperatures – things like 'blood heat', as Isaac Newton called it, which was the body temperature of a healthy human being. We might chuckle at that and say that he did not know any better, but even that was a great improvement over things like so-called 'most severe winter cold' and 'greatest summer heat', with no indication of where and in which year; the melting point of butter; the temperature of the wine cellars in the Paris Observatory, and so on. This was quite a problem.

By the late 18th century, there was a fairly well-agreed custom of using the boiling and freezing temperatures of water, but people were also worried about these points too – so much so that, in 1776, the Royal Society in London appointed a special committee to look into this matter of fixed points of thermometers. This committee was chaired by the great Henry Cavendish.

One can see why they were worried, looking at a thermometer made by a man called George Adams, the official instrument maker to King George III (Figure 7).

Figure 7. George Adams Thermometer

The wonderful thing about this thermometer – apart from having four scales, one of which is upside down starting with zero at boiling – is that it has two boiling points. One says, 'water boils ['boyles'] vehemently', and the other says, 'begins to boil'. 'Vehemently' is at 212 degrees Fahrenheit, while 'begins to boil' is at about 204 degrees Fahrenheit. That is a huge range. I decided there had to be a simple experiment to check this out.

Is the boiling point sharp?

I did a simple experiment to investigate why this range was given for the boiling point of water. I took an ordinary beaker filled with water and heated it with an ordinary Bunsen burner, monitored by three different thermometers. One was a platinum thermometer, working on the principle of electric resistance; the other two were mercury thermometers (one an ordinary one found in teaching laboratories, the other a Beckmannn thermometer, graduated to 1/100 of a degree centigrade). When the thermometers read about 96 degrees, there was already a good deal of activity in the beaker with vapour bubbles coming up but collapsing in the body of the water because the water was not hot enough.

At about 98 degrees, I observed the surface beginning to break. At that point the water had certainly begun to boil, and it was boiling quite well, at well below 100 degrees. The digital thermometer was quite jumpy during this process, but its readings were broadly consistent with those given by the two mercury thermometers.

Another observation I made was that the temperature began to creep up as I kept boiling the water. It would rise to close to 101 degrees in an ordinary glass beaker. My conclusion was that the boiling point actually is not sharp. This is very easy to observe and you have all seen it when you boil your kettle of water at home.

Joseph-Louis Gay-Lussac (1778 – 1850): boiling in different containers

This was only the beginning of the trouble. The next phase came with Gay-Lussac, who published a very curious paper about boiling in 1812, in which he said that pure, distilled water, under normal pressure, boils at exactly 100 degrees in a metallic vessel, while in glass it boils at 101.232 degrees centigrade. Again, I decided to investigate whether that actually happens. By now I already knew that the temperature of boiling water in a glass beaker creeps up quite a bit above 100 degrees, but how would it work in different vessels?

Here, once again I found that water was boiling quite well in a normal glass beaker, with a temperature of just over 100 degrees. However, in an aluminium vessel I found that the water was boiling at a temperature of only 99.3 degrees.

I also used a saucepan covered with Teflon, and it turned out that Teflon is extremely good at exciting bubbles. Bubbles formed very eagerly at quite a low temperature, even though the water was nowhere near hot enough for the bubbles to rise. However, the water was boiling well at only 98.7 degrees.

Using a ceramic mug I observed quite a different kind of behaviour with bubbles initially coming up from only one spot. The surface had trouble producing these vapour bubbles, even with the temperature at 101 degrees, and it continued to rise to 102 degrees.

Gay-Lussac was therefore right about the difference between metal and glass, although he did not know how right he was about the effects of various surfaces because he did not have things like Teflon. By this point, therefore, I was quite intrigued by all of this. You might have thought it was easy, but you need to establish what counts as boiling – which turns out to be one of the most complicated phenomena we have in basic physics and physical chemistry.

Jean-Andre de Luc (1727 – 1817): heating water slowly

There are two further kinds of experiments which I shall discuss briefly, both based on work by Jean-Andre de Luc. He was a very interesting character: a Swiss businessman from Geneva, but also a physicist, geologist and theologian who eventually settled in England as a tutor to Queen Charlotte. He lived in Windsor, gave the queen a lesson once a day and then, for the rest of the day, he was able to do what he wanted. Among other things, he was a member of the Royal Society committee on thermometry.

De Luc said to himself that, when chemists normally boiled water, they were not really getting at the pure phenomenon, in two senses. The first was that when they used an open flame, it was enormously hot and did not heat the water evenly. What you should do to see what water really does near the boiling point is to bring the whole body of water to the same temperature uniformly. You should not have a situation in which the water would be very, very hot right next to the surface being heated by an open flame, while the rest of it would be much cooler. He said, therefore, that you must heat it slowly.

When you heat water slowly, it will not boil because it loses too much heat from the surface and so, in order to enable it to boil in that situation you must use a thin-necked flask so that the surface area of the water is very small. In the chemistry lab I was working in, I found lots of chemists' volumetric flasks which have exactly the right shape. I used a hot plate which, although very hot, was nothing like an open flame. In this situation the water began boiling normally, but continued with bubbles only coming from a couple of places while the temperature crept up. It rose to nearly 104 degrees, with just very large occasional bubbles, throwing the water out of the flask. I observed the same pattern heating water in a bath of graphite, which was a way of heating water even more gently.

One final experiment that I will discuss is based on subsequent work by De Luc. In my experiments, I observed how the boiling behaviour progressively changed. Initially, boiling started with lots of little bubbles, as we normally observe it. Then gradually, there were fewer, larger bubbles as the temperature rose higher. What was happening there, at least according to the theory of the time, was that the act of boiling took more and more of the dissolved air out of the water, which made the water 'stickier', as it were. The dissolved air somehow facilitates the formation of the vapour bubbles and, when there is less and less air, the water has a harder time forming bubbles, with the result that the temperature increases – because the formation of vapour bubbles is how the water loses heat. That is why you have higher temperatures and fewer bubbles and so on.

De Luc then considered what would happen if he took the maximum amount of air out of the water. He tried all kinds of different techniques to achieve this and, in the end, he used a kinetic method. You all know what happens when you accidentally drop a can of coke or something: mechanical agitation will tend to release the bubbles, and so shaking is what he did.

This is De Luc's description of what he did. He said:

> "This operation lasted four weeks, during which I hardly ever put down my flask, except to sleep, to do business in town, and to do things that required both hands."

(He did not say what those things were.)

> "I ate, I read, I wrote, I saw my friends, I took my walks, all the while shaking my water....." (Jean André De Luc, *Recherches sur les modifications de l'atmosphère* (Geneva, 1772), vol. 2, p. 387.)

When he heated this carefully de-gassed water in a gentle bath of hot oil, he found that it went up to 112 degrees, without any boiling – but then it exploded.

I tried to reproduce this phenomenon, although I used a different method of de-gassing. First there was nothing going on and then, at the point where the water reached 107.5 degrees it exploded.

Engineering and the variable boiling point

How do we make sense of this and is there a good theory to explain all of these effects? This is where I came into engineering. As I mentioned, I was doing these experiments in the chemistry labs at UCL. The chemists were very kind, and they were talking to me about the theory because they were quite intrigued. There were two physical chemists, Steve Bramwell and Mike Ewing, who were specialists on phase transitions and such things and they suggested that I should go to the engineers to discuss these effects. So I then discovered an entire floor of the library at UCL which I had not seen before, the engineering floor, which I had never had any reason to visit. Sure enough, I discovered there a whole section of textbooks about boiling, which in itself was very enlightening (for example, Frank P. Incropera and David P. DeWitt, *Fundamentals of Heat and Mass Transfer*, 4th ed. (New York : John Wiley and Sons, 1996).)

I will give you an initial impression about the difference between science and engineering – and then I will correct myself about that shortly. The initial impression is that scientists just idealise things but the engineers really have to know what happens. If you are trying to cool down a nuclear reactor or any such thing at the correct rate, then you will not worry about a nice phase diagram, but you will worry about the actual rate of heat transfer and how that is affected by the precise conditions under which you make the reactions. That is what I initially thought and there is a grain of truth in it, but I will say more about it later.

As it turned out, the engineers have a theory, particularly to explain the effects of the surface quality of vessels on boiling. Why is the bubbling different in all these different containers? The answer seems to be that you need micro-pores on the surface in order to facilitate bubbling and this has to do with the basic physics of bubble growth. The key issue here is the force of surface tension, which tends to close up bubbles. Today's physicists are only thinking about balancing the vapour pressure with the external atmospheric pressure pressing down on the water, and not really thinking about the fact that a bubble with the two forces balanced out exactly will actually close up because of surface tension.

The surface tension, JJ Thompson and others worked out, was basically inversely proportional to the radius of the bubble. In the limit, if you are trying to grow a vapour bubble in liquid water, starting from nothing, you cannot do it, because the surface tension will be infinite. It will never go to zero but, if there is nothing to grow it from, it will be very close to that.

In order to instigate the boiling process, you need a solid interface with what are called nucleation sites. Micro-pores on the solid surface will be filled with air, vapour, or sometimes even vacuum if the surface is hydrophobic enough. Those little pre-existing bubbles of gas or vacuum will serve as convenient nucleation sites – which is why, when you have a glass of beer, the streams of bubbles come only from particular spots and not from all of the surface of your glass.

That was a great education for me, having been educated in physics and not engineering, to know that physicists had basically nothing to say about this but the engineers really had a theory. However, even engineers do not seem to have a good theory about the effect of the dissolved air in the water. There are some computational physicists trying to model that sort of thing but, to my knowledge, they have not got very far.

Why didn't we all know about the variations?

This led me to two further questions. Most of us boil water on a daily basis and so, if there are such great variations in the boiling temperature and behaviour as I have just shown you, then why did we not all know that before? The answer is that some of you knew this very well, but others of you probably did not. This is partly because we are lazy. Ever since I did those experiments, I can no longer look at a body of boiling water without fascination! Every little saucepan in the kitchen is different and, once you start looking, you will see these differences.

There is a bigger factor in that the everyday kind of boiling is really something very particular. We normally boil water in great open-faced containers, which implies that you also need an intense source of heat and so you tend to use open flames, and that is what produces the so-called 'normal' boiling behaviour that we all recognise. However, with slight changes in the boiling setup it is very easy to produce what is called super-heating. There have been real-life situations in which super-heated boiling was the norm rather than the exception. In the 19th century, there were great problems with boiler explosions because, when you recycle the same body of water, over and over again, making steam and getting all the air out, the water will have a hard time boiling, to the extent that you have mighty explosions in the boilers, killing people working around steam engines. The 19th century engineers came up with the solution of putting in what they called 'boiling chips' to make the boiling smooth, and those are still used in chemistry labs today.

Even more prosaic and modern is the use of microwave ovens. If you type in 'super-heated water' in Google, you get lots of sites talking about this, including a number of American fire brigade websites. These warn people not to boil their water for instant coffee in the microwave because there you are heating the water directly, and not through a very hot surface. The water is then very prone to super-heating and once you put in a spoonful of instant coffee it 'explodes'. Just slight variations in what we do, therefore, will make these great differences.

This leads to a question of practical purpose. Why is it that we boil water, and what are we interested in when we study boiling? There are very different purposes. The typical engineer's interest is the problem of heat transfer; the typical physicist is concerned with the abstractions of phase transitions. You and I, in everyday life, worry about getting things to boil in order to cook things. Every different situation will produce different phenomena.

Ontological issues

To wrap up, let me return briefly to ontology. Having seen all of these things, I have found myself asking the question, what is boiling? In fact, I am not the only one. Apparently de Luc was so captivated by this phenomenon of boiling and the great variations in it that, when his book – which was mostly about meteorology – was just about to go to press and he started doing these investigations into boiling, he ended up adding an 18-chapter supplement about this boiling phenomenon.

He concluded that there were six different kinds of phenomena that people called 'boiling', which were really very different from each other. These included the ordinary, normal boiling which we recognise; and then there is the very erratic, explosive event that I showed you with the volumetric flask; then going down to just very fast evaporation from the surface when the water is super-heated – but then, in between the big bubbles, is it boiling or not? Do we know what boiling is? De Luc concluded that he did not – although he did not really push that point in the Royal Society committee on thermometry. However, had he done so, it would have caused some mayhem.

But if we do not know what boiling is, of course we do not know what the boiling point is either. Which brings me to another point: if the boiling point of water is so indefinite, how was it that thermometry could be stabilised by using the boiling point as one of the fixed points? What do we mean now by 100 degrees centigrade? There is a long answer to that which I will not give you, except to say that physicists realised that the boiling point of liquid water was not such a stable point, and they then moved on to using the temperature of steam and on to other, more sophisticated things.

The other main point to make about this comes from my chemistry colleague, Mike Ewing. When I showed him all these variations in the lab, he let out a big sigh and said that these were precisely the kind of things that he tried to avoid – and Ewing and his colleagues are very good at avoiding them. This means that they know precisely what situations they need to use in order to produce fixed, reproducible boiling behaviour with a fixed temperature.

I would like to close with a brief comment on how the purposes of engineering and scientific practices shape the ontology of even everyday phenomena like boiling. Engineers and physicists – and different types of physicists and different types of engineers – are looking at very different aspects of what happens when you get water to a high temperature. There are different conceptual frameworks that they apply, so that the engineer's boiling curve cannot be expressed on the same basis as the physicist's phase diagram. They are dealing with different problems—for instance, the boiling curve is not talking at all about the temperature of the water. It would be very interesting to apply the same kind of analysis to more complicated phenomena, but a great lesson that we can take is that even in the most mundane, ordinary phenomena, there are great problems of practical or applied ontology.

For further discussion and video footage of experiments, see Hasok Chang's online paper, "*The Myth of the Boiling Point*": **http://www.ucl.ac.uk/sts/chang/boiling/index.htm**

1.3 Ontology in engineering

Professor Peter Simons

Peter Simons studied mathematics and philosophy at Manchester and then taught at the Bolton Institute of Technology and the University of Salzburg in Austria. His specialisms are metaphysics and ontology and their applications, the history of logic and the history of philosophy in Central Europe. He is the author of Parts and Philosophy and Logic in Central Europe from Bolzano to Tarski and has published over 200 papers. From 1989 to 2001 he was ontology consultant at Ontek Corporation in California and at Leeds he collaborated with the Department of Mechanical Engineering and the Keyworth Institute. Since 2009 he has held the Chair of Moral Philosophy (1837) at Trinity College Dublin, where he is also Head of Department.

Introduction

There are certain kinds of philosophy that you can see straight away would be relevant to engineering. One of these would be ethics, of course, because engineers face actual ethical problems from time to time and sometimes they might call on the advice of a so-called professional. Another way in which philosophers can contribute indirectly, perhaps, to engineering, is in their input into the design of information management systems, IT systems. However, I shall not talk primarily about these, but more about more direct ways of interaction between philosophers and engineers.

I shall begin this paper by saying briefly what ontology is. Secondly, I shall ask, what do ontologists investigate? The next question is, what use is ontology? I will give some examples from my own experience of how ontology might be useful and suggest general ways in which ontologists could help engineers and how they even might have done so in the past. Finally, I will suggest how engineering can help ontology: in other words, what do I, or people like me, get out of a collaboration with engineers?

Ontology: what?

There is a book with the glorious title of *Philosophia prima, sive ontologia, methodo scientifica pertractata, qua omnis cogitationis humanæ principia continentur*, meaning 'First philosophy or ontology, treated by the scientific method, in which are contained all principles of human knowledge', by the German rationalist philosopher Christian Wolff. Published in 1729, this was the first widespread appearance of the word 'ontology' – in this case, the Latin version, *ontologia* – in history. The word was invented in the late 17th century in Germany but it was Wolff who made it popular. Wolff was the most influential philosopher in Germany up until Kant in the late 18th century. He wrote two philosophical systems – the first, in German, in which he single-handedly invented most of the philosophical vocabulary that Germans use today, and the second in Latin. His German system and Latin system each fill very large bookshelves.

Here is my translation of Wolff's definition of ontology: 'Ontology, or first philosophy, is the science of entities in general or insofar as they are entities.' What we might take from that is that ontology is a very general, broad, wide-ranging subject; indeed, there is none broader. This idea of a maximally general discipline or science is by no means new and, in fact, you find it given with almost the same words in the fourth book of Aristotle's *Metaphysics*, where he talks about a 'science of being as being'. He describes such a science as dealing with the first things and the first principles. Incidentally, the word 'metaphysics', which encompasses ontology, is not a word that Aristotle used. He called it 'first philosophy'. The word 'metaphysics' was actually invented by Aristotle's editors about four centuries later when they collected together a rag-bag set of books to do with very general things and put them after the books on nature. The books on nature were called *Ta physica*, which comes from *physis*, meaning 'nature', from which we have the word 'physics'. The books after the books on nature were called *Ta meta ta physica* – the books after the books on nature. So the name 'metaphysics' is merely a librarian's invention, whereas 'ontology', the word which fits the notion perfectly, is actually a very late coinage.

In Wolff's way of dealing with things, and I think this has not been bettered in overall strategy, metaphysics has two sides. On one side, metaphsyics is the theory of being as being, or ontology, which equals *metaphysica generalis*, and that contains principally the theory of categories and principles which are applicable to all things. In the early 20th century. This was called 'formal ontology' by the German philosopher, Edmund Husserl.

On the other side, Wolff distinguished a number of slightly less general disciplines dealing with the mind, the material universe, and God, and he called these psychology, cosmology and theology. Rather than delve into the interstices of 18th century rationalist metaphysics, I prefer simply to say that, in addition to ontology, there is what one might call *systematics*, which is the study of the diversity of the domains of being and how they are interconnected. That was called by Husserl material or regional ontology. If you are interested in the most general concepts that are specifically applicable to engineering, then you would be doing the material or regional ontology of engineering.

Categories

I will now turn to categories because dealing with these is what gives ontologists any purchase that they have in being able to help engineers. What are categories? They are general concepts, dividing and classifying things at the most fundamental level and they are therefore extremely abstract and widely applicable. If you get two metaphysicians – especially category theorists – in the room, then you have a dispute. In fact, between any two metaphysicians, there will be at least three opinions and probably more! So as to how many categories there are and which ones they are, no two philosophers will agree.

Secondly, there is no consensus about how you arrive at the categories. We are not quite sure how Aristotle did it, because we only have the result from Aristotle and not the method, but it was plausibly reconstructed in the 19th century by French and German historians that he got the idea for the categories by reflecting on the grammar of simple Greek sentences. The categories of the 18th century, from Kant, were obtained by reflecting on the forms of logical judgment. So both Aristotle and Kant used what I called a logico-linguistic method. That is one of various ways in which you might come up with the categories. I will illustrate the options by mentioning some choices. *A priori* or *a posteriori*: you might think, for instance, that you could derive the categories simply by reflection, *a priori,* without regard to experience, or you may think that you need experience to help you decide which categories are applicable. Is the investigation of categories immune from or open to scientific advance? My answer is that it is open to scientific advance, though Kant thought it was not. Logico-linguistic or hypothetico-deductive: do you arrive at the categories by a method similar to that of science, or do you arrive at it by reflecting on language and logic? The predominant tradition in Western philosophy, even today, is the logico-linguistic method, but I do not think it is right. Fallible or not: are the categories something that you get right once and for all, or do you have to crunch, work, grind, sweat and groan at them, as we find out more about the world? The answer is the second, unfortunately. Are the categories something that stay the same over time and have no history, as Peter Strawson, the recently deceased Oxford philosopher, said of descriptive metaphysics, or do they promise to revise our ways of conceptualising the world? The answer, in my view, is the second. Finally, do we investigate the categories one at a time, in bits and pieces, as analytic philosophy does in short articles in the journal *Analysis*? Or do we try to string them all together and produce a connected system? Answer: the second is the only possible way you can do it and everything else is just mucking around in a sandbox.

I have offered these oppositions as mutually exlusive, but in practice they work more like polar opposites, and you can get cases which mix both kinds of approach, the *a priori* armchair approach and the *a posteriori* empirical approach. When I first started working at Ontek, we worked like logicians, but the chief breakthrough in learning to do things differently came in 1991 when I read the great biologist Ernst Mayr's book on the history of biology. I said, "Look guys – what we are doing wrong is that, instead of following the paradigm of logic and mathematics, which is what philosophers do, we should be following the paradigm of biology, which goes out and investigates raw similarities, differences and events in nature, and then makes sense of them, using an extremely complicated classification system." We studied, analysed and abstracted the principles of the biological classification system and that analysis was used to revise the ontological structure, so that it was *a posteriori*, but it was done the hard way – we had found out that it did not work the other way, or that it was too inefficient. However, what we were of course aiming at was an ideal language for representing the facts as we categorised them, and that is why it has to be fallible, because you have to zig-zag back and forth between the two approaches.

So now here are some examples of categories from the formal and from the material side:
Formal: identity, existence, part/whole, dependence, material, property, feature, boundary, number, quantity, relation, order, situation, structure, function, process, event, change, similarity, persistence, location.
Material: time, space, cause, emergence, mind, agency, artefact, intention, choice, design, manufacture, need, requirement, cost, effectiveness, success, strategy, institution, species, biological character, engineering feature, geophysical feature.

As you get down into the material categories, or the concepts applicable in specific domains (I have tried to do these in the direction of more general to less general) you will notice that categories are beginning to creep in which might be of interest to engineers, such as agency, artefact, intention, design, manufacture, need, requirement, cost-effectiveness, success, strategy and so forth. These are all highly general concepts which can benefit from the kind of expertise on thinking at stratospheric levels of abstractness that philosophers bring to their job when they are doing it properly, but which are still relevant to the concerns of engineers. I actually think that engineers are in some ways closer to philosophers than either are to scientists – they have a rather similar take or attitude to the world.

Ontologists are philosophers, so they argue with each other and in particular, they argue about categories. They argue about which ones they are, which ones you need and which ones you can do without. Can you do with fewer, or do you need more? Should you go for a more parsimonious explanation, or is the parsimonious explanation explanatorily inadequate? If it is inadequate, we need to add some more categories. If you can do without it, then all well and good; get rid of it and do what William of Ockham said: cut it away with your ontological razor.

So ontologists then fight about which categories you can get rid of or eliminate in relation to which other ones. Then finally, when each individual has decided, at least on a Thursday afternoon, which ones he or she likes, they will then spend the rest of the weekend wondering what the principles are that govern them, before completely throwing it all out of the window on Monday morning and starting again. So that is a rough sketch of what ontologists are about. Huge amounts of heat are generated, with slightly less light in the journals, and it is not just serious but also great fun.

What use is ontology?

What use is ontology, inside and outside philosophy? Inside philosophy, it is useful for generating controversy, keeping the subject bubbling and ontologists in a job. I am being flippant, of course: it is a serious subject. What use could it be outside? Let us put the arguments and disagreements aside and suppose that, somehow, ontologists are tamed and put to use. How could they be helpful to engineers, among others? They are used to producing conceptually tidy, consistent, complete and well-formulated conceptual systems, so they can bring those skills to bear. Tidiness and hygiene consisting, for instance, in not having different words for the same thing and not having a word which is ambiguous and so on. If all goes well, they can introduce additional analytical depth and perspicuity to the treatment of concepts which other people take for granted. They can help to eliminate redundancies and repetitions or overlaps, which can be tidied up. All of this is another form of conceptual hygiene, you might say. They can simplify things down – not so that they become so simple that they are brain-dead, but so that there are no unnecessary repetitions and redundancies.

Philosophers are trained to be articulate. They can and should bring this articulacy, this ability to formulate their thoughts as clearly as possible in as straightforward a way as possible, to the study of whatever subject they are working on, whether that be ontology or engineering or applied ethics. You can hope to rely on the articulacy of ontologists, as you can hope to rely on the articulacy of other philosophers.

Because philosophers deal with high generality, they are sometimes capable of seeing connections where other people, who are more specialised, do not. This is one reason why the government, in its infinite wisdom, has occasionally hired philosophers to chair Royal Commissions – people like Sir Bernard Williams or Dame Mary Warnock. They were specifically chosen because, as philosophers, they were used to dealing with these inter-disciplinary matters in such difficult subjects as surrogate motherhood or pornography. Finally – and this is another aspect of the articulacy – they are, or should be, able to express themselves clearly, succinctly and, we hope, entertainingly.

Those are some very general virtues that ontologists are supposed to be good at. That might just qualify them merely to be quiz show hosts on television, so let us see whether they can do more than that. Let me take three examples where, *qua* ontologist, I have personally had input into collaborations with engineering entities, engineering collectives or firms and so on, where the ontological input has had some effect. These are three very sketchy examples but they are real.

Example 1: multiple bills of materials

The first example, and certainly the most important for many purposes, was a study of the so-called multiple bill of materials problem in application to aerospace engineering. Bills of materials are engineering breakdowns of the structure of complex artefacts and it turns out – and this is something that surprised me when I encountered it for the first time – that one bill of materials will not do for engineering purposes: rather you usually need three or sometimes more for any particular artefact. You need one for the engineering design, one for the manufacturing and one for the maintenance, and sometimes even one for the retirement. The problem is that they look at different aspects of the breakdown of the artefacts and the way in which they are put together, taken apart and messed around with, and they give you incompatible accounts of what these artefacts are made of. Actually reconciling these differences turns out to be a severe financial and conceptual problem for large enterprises. We came across an actual case where different bills of materials for a well-known aircraft type differed by 40%, and there were no links between them, so any engineering changes had to be propagated by hand.

The United States Air Force Manufacturing and Technology Division, in its wisdom, hired Ontek Corporation, for which I was working at the time, to look into how to rectify or at least partly alleviate the multiple bill of materials problem. They thought, quite flatteringly, that we would be able to produce an ontological account which would give us additional bite over the standard solutions that were at that time suggested.

The analysis of the problem took place in the early 1990s and appeared in an article called "Aspects of the Mereology of Artifacts", published in a book I co-edited with an Italian colleague in 1996. 'Mereology' means the part/whole breakdown, or theory of part and whole, here as applied to artefacts. In this article, the chairman of the company, Charles Dement, and I, analysed where the major discrepancies lay between different bills of materials. Our analysis was then applied to the multiple bill of materials problem and written up in a technical report that was submitted to the United States Air Force five years later, under the term *MEREOS*, which was the name of the USAF programme.

Example 2: databases

The second example is a slightly less grandiose one but it was in fact a part of the requirements for getting *Mereos* off the ground. There was the apparently trivial-seeming problem of reconciling the mutual incompatibility of databases. Databases have grown up in different decades with different structures and they do not talk to one another well. It is rather like you trying to talk to someone whose only language is ancient Hittite, for instance – it doesn't work too well. If you are running, say, an Access database on a desktop computer, or an Excel spreadsheet, or a FileMaker database, you will not necessarily be able to do queries on an Oracle, or a DB2 or IMS database, because the languages and the whole architecture are incompatible.

This would be fine if enterprises only had one database apiece but, unfortunately, they often have generations of them stacked up in their management. The example that we investigated was Alcoa, which had something like 14,000 different databases going back decades and very few of them "talked" with one another. This meant that data had to be gathered from the different databases by hand, thereby nullifying half the importance of databases – which is that you do not have to do everything by hand because the machine is supposed to do it for you.

The way in which we went about tackling this, which was principally conceptual, was to analyse the different architectures and structures of data, using an account of what we called facts (not to be confused with what philosophers call facts). This was the main thrust of the work that Ontek did in the 1990s. One of the things we produced was a cut-down version of our major database management system we called Information Access Vehicle (IAV). This was piloted in a demonstration through Northrop and Westinghouse in 1993 but they decided, I think wrongly, that it was too complicated to work at production scale. It was interesting, however, that we could do online queries and update an IMS database, which it would normally take an ancient IBM (and an ancient IBM programmer) to do, from a Macintosh desktop, because it went through a kind of translation module, which was based on our formal ontology.

Example 3: Lockheed Martin

The third example is more recent. There is a very long and complicated story about how we got into this, so I will not go into how we got from software design to enterprise design, but we did. During the late 1990s, Ontek was tasked with helping with the re-organisation of a large aerospace concern in Georgia. The question they faced was how to turn the enterprise, with no Cold War to finance inefficiencies, into a modern, commercially viable manufacturer of military aeroplanes.

Our advice was that they needed to exploit a formal analogy – in our view, in fact, an identity at the most formal level – between systems engineering and enterprise engineering. We decided that the best we could do would be to treat the enterprise itself as if it were a product or a manufactured item, and use the same principles that system engineers apply to systems to the enterprise itself, rather than leave the redesign in the hands of business people. We applied that to the strategy problem at Lockheed Martin Marietta in 1999 but, unfortunately, while we were in the process of doing that, Lockheed Martin in Georgia was taken over by Lockheed Martin in Texas and the whole thing was shelved because the Texas management did not take over our work.

How ontologists can help engineers

In what ways can ontologists help engineers, if all goes well? This is based on the premise that people are even interested in starting with letting ontologists through the front door – which is not always a good idea.

First of all, it is nice to be asked to help, especially with something other than making tea. Hard problems are what ontologists ought to be set at. Secondly, they should be trusted, at least in part, and given information embedded in the enterprise, as far as required, and made part of the team. It is important that they are made part of a team, otherwise they will go into a corner and write papers for the *Journal of Philosophy*, which is what they do in their spare time anyway. They should actually be made to cooperate and talk to other people, and these should be the right people – the people who know what they are talking about and who are sympathetic. So those are the passive things that can be done to help them. They should be invited in, given an office and made part of the team, and they should speak to the right people.

What do they do to earn their microdollars? They should bring to bear their special confusion-clearing glasses. They should bring a new perspective to old problems – and I will give you some reasons as to why their perspective is likely to be new at the end of this paper. They should be open-minded and they should listen and learn, rather than just supposing that they can go in and preach and get everyone doing their philosophy with machines. They should be prepared to put aside their petty philosophical squabbles. They should not expect to do everything by Thursday, or be expected to. However, the main thing they can be is not a lawyer, accountant or manager – that is how they can bring a different perspective. That is why I think we should be friends, because I am not an accountant, a lawyer or a manager, and neither are you, so we can help each other. However, that's philosophers in general. What about ontologists?

The specific virtue of a category-based ontology is that it acts as a kind of universal checklist. If you have a concept, process, object, relationship or whatever that you only dimly understand, such as engineering change, or function, or efficacy, an ontologist can not only conceptually analyse what the rough-and-ready word currently means (any good analytic philosopher should be able to do that), but by breaking down the objects investigated into their ontological basic factors or components, because ontology in Wolff's conception applies to anything and everything, it holds out the promise that nothing important has been left out. It's the systematicity, the universality, that counts. Otherwise you can never be sure something important isn't missing from the analysis. Of course because the ontology is fallible, the result is not 100% guaranteed, but it is likely to be better than an *ad hoc* approach.

This is how engineers can get something out of philosophers, especially ontologists. I must say that I am almost certainly likely to receive a better reception from you than among my own colleagues in philosophy, who tend to be sniffy, dismissive and supercilious, or even downright angry, at the thought that a philosopher should even think of getting his hands dirty by messing with people who do real stuff. Fortunately, that attitude is changing somewhat.

How engineering can help ontology

Here is why it would be good for philosophers to get out more – and I do not just mean to the theatre or the pub. Engineers can give philosophers hard and real problems to solve. They can give them areas where they can see whether their philosophy can make a real difference to the formulation of theories and conceptual structures.

Philosophers notoriously have no one who can tell them they are wrong, except for other philosophers – and they do not count, because they are in the same business and are competitors. However, if you try to build something, even if it is only a mock-up, using a philosophical theory, and you find it does not work, then that is a good indication that the philosophy needs amending somewhere.

It can broaden the minds of philosophers, some of whose minds need to be broadened. It can take them out of the ivory tower, and it can give them something practical and useful to do. I regard all of this as for my benefit and many of my colleagues would say, 'No, thank you.' I do not say that we should put them in the countryside for two decades, as they did in China during the Cultural Revolution, but it would not necessarily be a bad idea to give them something like national service at Rolls-Royce, for example! It could help to bridge the two cultures, which do not need bridging so much in other countries.

In the United States, Germany, Italy and France in particular, there is not as massive a divide between the arts and the sciences, or between the arts and the practical disciplines, as there is in this country. This is something that is very strong in Britain, and particularly so in England, and it is a legacy of the aristocratic attitude to the scientific revolution. This means that there is still a kind of class divide between people who do things and people who think about them. There are of course precedents for that in the ancient Greek philosophers as well, none of whom liked to get their own hands dirty – that is what they had slaves for.

1.4 Constructionism, the maker's knowledge tradition and knowledge engineering

Professor Luciano Floridi

Luciano Floridi is Professor of Philosophy at the University of Hertfordshire (Research Chair in Philosophy of Information and UNESCO Chair in Information and Computer Ethics) and Fellow of St Cross College, University of Oxford. In 2009, he was elected Gauss Professor by the Academy of Sciences in Göttingen, awarded the Barwise Prize by the APA, and elected Fellow of the Society for the Study of Artificial Intelligence and the Simulation of Behaviour.

Introduction

This paper will start at the beginning with Plato, and I will tell you a little about the maker's knowledge tradition and how people, especially philosophers, are talking about who has better knowledge of the artefact in front of you. I will provide a little background on the tradition itself, including Hobbes's characterisation of constructionism.

There are some consequences which I will then illustrate quickly, before moving to a more general understanding of the idea behind the picture of knowledge as something that we at least partially build. Knowledge engineering and computer science will be my next step, which I will try to summarise in a number of principles. The names are just labels and it is the concepts behind them which matter. They are: construction and control, constructability, confirmation, non-descriptivism, economy.

Towards the end, I will move on to discuss a particular application of this idea of knowledge as something that we build. We will visit the famous Turing Test before drawing some general conclusions.

The beginning – Plato

In the beginning: Plato. Imagine that you have an artefact and you want to know who knows the artefact better. We will follow Plato's example and consider the lyre, but it could equally be a computer or other artifact. Who really knows it? Is it the imitator, the painter who paints the lyre? Or is it the maker, the artisan who made it? Or is it the user who plays it?

From Plato onwards, one of the standard answers has been slightly counter-intuitive. If I were to consider this question as an undergraduate, say, then the natural answer is that it is the person who built it. He knows it most. However, according to Plato, it would be the user who possesses better knowledge than the maker.

This is not your typical philosopher trying to be controversial. He had good reason for his answer. He thought that the user was the one who could go back to the maker, the engineer, and say, 'Look, it does not work in the way I would like it to work, and you will have to do something else.' So, for quite a while – several centuries in fact, and even today – some argue, especially within philosophy departments, that if you know how to use the artefact, then you really know it. You come first, and the maker or engineer who actually built it comes second in terms of knowledge of the object.

All this was actually provided as a general argument, and I will spare you the details, against the idea that the imitator knows the object best. Plato goes through this argument with different objects – but they are always artefacts, or things that one has built.

What is rather puzzling about that is picked up quite nicely in this quotation. This is in the dialogue Euthydemus, and it is Socrates speaking:

> "Then the sort of knowledge we require is that in which there happens to be a union of making and knowing how to use the thing made."

You can see the tension here: yes, making the lyre is not sufficient but you need to know how to use it. But why split the two first of all, and why privilege the user instead of the maker? Why is it that we have a position which casts a little light, but perhaps not the most interesting light, on the people who are hands-on?

There are several problems with this idea, the main problem being that the engineer, the maker, is not the one who knows the artefact better than anyone else. However, the argument was not intended to be against the makers, but was really meant against the imitators, the painters, those who observe and hear and account. Despite this, the argument was interpreted as an argument that put the lyre maker in second position, and it established a long tradition in which the maker was put in second place. This created an unnatural separation between the user and the maker as if, when you are building something, you are not testing it and you do not know the requirements or specifications of the artefact that you are building. That did not help the epistemological debate in the following years. There is also a disregard for technicians, the people who actually get their hands dirty with the real stuff. The assumption was that if you have not made something but just use it, that is so much better.

You can see how this was going to generate some problems when you consider that, later on, the Christian viewpoint would look at God as the engineer. If you take God as the engineer – the guy who makes the stuff – and us as the users and follow the Platonic argument, then we conclude that we know the world better than God. In contrast, there is a line of thinking which analyses God's omniscience on the basis of maker's knowledge. The argument was that God knows everything, because he made it – a conclusion that goes against straightforwardly against the Platonic view.

Of course, for those of us who are used to treating information or artefacts as packages that we download from Google, it is not controversial to claim that we know better than the people who are actually behind the production of that information. So you find a whole main trend in epistemology talking about who knows better – not the engineer but the guy who knows how to use the artefacts in question.

The maker's knowledge tradition: constructionism

The opposite view, namely that the maker knows better, needs a label. I call it constructionism, but any other label will do, so do not worry too much about that. The basic point – and, once again, I am simplifying and abstracting – is that if you have genuine knowledge, or knowledge of the intrinsic nature of the object that you know, i.e. its ultimate nature, then we as epistemic agents can only know what we make. That is the real point. I do not know anything about a TV even though I use it. If I need to repair it, I do not have a clue where to start from. Someone who knows the object in question is someone who can put his hands on it and change or modify it. I will look at this further in later slides.

The constructionist's anti-Platonic view is that knowledge is acquired only through the construction of the semantic artefacts, the information modelling, or hands-on by someone who would be the maker. In this context, I am happy to please the audience and say it would be the engineer.

Epistemic research and the research for knowledge and information modelling are two sides of the same coin and cannot be separated. If you know how to build it, how to put it together and how to dismantle and repair it, then you know it. If you do not, then you are on the 'imitator' or 'user' side. The whole field in which I have been working for some time, in epistemology, is almost obsessed with the vocabulary of constructionism. Ideas come from there, and there is a whole frame built on that particular orientation.

Hobbes on construction

One of the major figures in the field – Hobbes – made the point quite clearly. Hobbes is considered to be an empiricist and so that is on the right side of the battle. I shall give a long quote from Hobbes here:

> "Of Arts, some are demonstrable, others indemonstrable; and demonstrable are those the construction of the Subject whereof is in the power of the Artist himself; who in his demonstration, does no more than deduce the Consequences of his own operation."

Here is why:

> "…the Science of every Subject is derived from a precognition of the Causes, Generation, and Construction of the same"

You engineered the thing –

> "and consequently where the Causes are known, there is place for Demonstration; but not where the Causes are to seek for. Geometry therefore is demonstrable; for the Lines and Figures from which we reason are drawn and described by our selves; and Civill Philosophy"

- politics, sociology and so on –

> "is demonstrable because, we make the Commonwealth our selves."

If I make it, I can know it, but if I do not make it then I am an outsider – it is a black box.

> "But because of Naturall Bodies we know not the Construction, but seek it from the Effects, there lyes no demonstration of what the Causes be we seek for, but onely of what they may be." (Six Lessons, epistle: EW VII.193-4)

So one point is that, if you build it, you know what it is – but if you have not built it, you are never quite sure. You can model it and you can test it, but you will always be an outsider. So that is the general picture.

I am giving you this in black and white and it is rather stark but in the Platonic tradition it is the user who knows and not the maker. This is counterbalanced by this Hobbesian view that, if you can construct it and build it, then you know it. However, he was not the only one who held this view, and I will give you three quick names in the history of philosophy from this tradition.

Constructionism: three consequences

The first is Francis Bacon, with his view that knowledge is power. Bacon has this wonderful view which I will summarise in a few lines, which is that we can improve our knowledge through the improvement of the techniques by which we investigate reality – in other words, if you can do better with your technology, then you get closer and closer to the stuff that you are dealing with. If you can, for example, build something from its DNA, then the kind of knowledge that you have of that thing is better than that of someone who is just observing from the outside.

There was another Italian philosopher named Vico, who had two famous quotes. He said that what is true and what is made are interchangeable, *verum ipsum factum*, or that they are two sides of the same coin, *verum et factum convertuntur*. This was the beginning of sociology with attempts to reach a comprehension of the world as if you were God – as if you had built the world. So it is better to concentrate our attention on those sciences or subjects created by man. So you have these two views: improve your technology because, if you can build it, you know it; and, insofar as you cannot build something, concentrate your knowledge on the things that you can – for example, society. This is reminiscent of the Hobbesian view about politics.

The third view comes from Kant, whom I would like to recruit as a constructionist, concerning the limits of knowledge. Because the ultimate nature of reality remains unknowable, you may do much better to concentrate on the conditions that make your knowledge possible. What kind of artificial or natural epistemic agents are we, and what kind of constraints, limits and affordances can we deal with, so that the picture that we get of the world is one that we find satisfactory? Look at the conditions of possibility, rather than looking at some unreachable ultimate description of the world as it is in itself, because there is no such thing as that. Those are his views, and it is a lesson from the constructionist approach.

To sum up, here are the people on the chessboard: Plato, in white, and Hobbes, Francis Bacon, Vico and Kant playing in black. You can read a great deal of the history of philosophy of knowledge as shifting between one side and the other, sometimes trying to mix the two, and sometimes favouring one side or the other.

Knowledge engineering and computer science

That is the general picture, but is there any context in which this maker's knowledge idea has some convincing currency? Some time ago I started working with computer scientists in Oxford. I discovered that it was much easier to talk about constructionism to a computer scientist than to a philosopher because the computer scientist is used to building things from scratch, from plans and specifications. The philosopher is more accustomed to saying, 'The world is out there. How can I possibly come to some knowledge of it, how do I have access to it?' The philosopher plays defence, whereas the computer scientist plays attack.

From the non-Platonic constructionist perspective, what we are dealing with when we are dealing with objects in the world are black boxes, in the computer science sense. They are systems for which we ignore the internal structure, and we try to understand from their behaviour on the outside what the internal workings might be. But because we did not make these things, then we cannot have this kind of engineering knowledge of them. Things in themselves – and here I am using some sort of Kantian vocabulary – are like black boxes to us. We are just epistemic agents with our capacities and constraints; we can never know the intrinsic nature of things in the world because we did not make them.

That is the extreme anti-Platonic view. If knowing means making it, then if you did not make it, you do not know it. The more you can put your hands on something, the more you can modify it, then the better your understanding will be, but the ultimate knowledge that an engineer would have is an ideal that you approximate but which you never reach.

I have argued in the past that philosophy, and especially this constructionist perspective, can definitely learn a few lessons from computer science. I will now show you what kind of lessons we have been trying to learn with my research group back in Oxford.

Six constructionist principles

These are six general ideas about what constitutes a constructionist epistemology of 'knowledge by making it'.

First is the principle that I have already spent some time discussing, that we can know only what we make.

Second is the principle of constructability. This is the principle that knowledge is gained by theoretical or practical simulations, and the better the simulation then the better your knowledge, but what you know is only of the simulation.

Third is controllability. A good simulation should be something that you can control, otherwise you know less than you knew before, or nothing at all.

The principle of confirmation: any confirmation or refutation of the original hypothesis which may be built in the simulation concerns the simulation and not the simulated. This often happens to be a problem and sometimes, because the simulation works in a certain way, people extrapolate and decide that that is the way the world is. As far as I am concerned, I would be cautious about any God's eye perspective. I know how the simulation works: I built it and I can predict its behaviour. I hope it captures something of the simulated, but I must try not to project such features onto the simulated objects.

Non-descriptivism: here is where my philosophy colleagues are normally not very happy, so you should be aware of the controversial nature of this. We acknowledge the stuff out there and our human knowledge of it. It is not in a photocopy relationship, but more of a resource to product relationship. The stuff out there gives me the elements with which I build something concrete. I am not a relativist and I am not a social constructionist, but I just think that what you do with photons going through your eyes is up to you. We are happy with the world in the way that we model it and, as far as we know, there could be nothing else and the model could be the perfect copy or not. We should not ask, 'is what I know really, really, really the way things are in themselves?' The unexciting answer is that we should not care, because that is not really what it is about.

Finally, the principle of economy. In the field of conceptual resources, we use the essential entities; we want to be economical in our modelling of the world.

Construction and control

And now for the explanation of these six principles. As I said, construction and control come more naturally in computer science and, once again, I thought that a quick quotation could help. This comes from Newell and Simon (1976) from *Computer Science as Empirical Enquiry: Symbols and Search*:

> "Neither machines nor programs are black boxes; they are artefacts that have been designed, both hardware and software, and we can open them up and look inside."

That is the criterion for being the maker of an object – being able to open it and look inside. The whole debate about how realistic your knowledge of the world can be could be framed in terms of people who believe that you can open up reality, although not nature itself, and look inside, and people who say that they are modelling the external behaviour alone.

Controllability

A system is controllable if it is modifiable, compositional, oriented towards a particular goal (teleological), and predictable. I will briefly explain these four features. It is modifiable if we are able to change its internal structures. The outside behaviour provides data, and we can transform the data into information by attaching meaning or sense to it. Our models are modifiable insofar as we are able into account new data and so on. It must also be compositional.

It has a goal – the system that we have built in order to understand something has been made with the maker's goals. We look at the model, at the simulation, and we try to understand what the system is doing, what for, and towards what particular goals.

It is predictable – with a lot of salt and not just a pinch! Hopefully, there is something in the system which we have modelled which allows us to talk about the model as sufficiently predictable, otherwise it would not be a good model. This means that we know the rules of the system and we can know its behaviour and use it to predict the behaviour of another system, or the natural system that our simulation is trying to model. Basically, we hope that what we build will give us some idea about how it will evolve through, for example, state transitions in the future.

Confirmation, non-descriptivism and economy

The principle of confirmation operates from generalising the working hypothesis as if the simulation were the internal structure of the simulated. For example, I developed a simulation or reconstruction of the working of the legs of an ant. As far as I know, it is a pretty good one. There certainly might be features of that model which are just features of that particular model, which have nothing to do with the ant in itself. I cannot not extrapolate from there.

There is some context-dependency. Any connection between the simulated and the simulation is local and not global. You change the parameters and that one-to-one relationship will completely change. The idea is not to make a photocopy but just to make sure that some things are simulated by other things and, if you change the level of abstraction, as I shall say in a moment, the simulation may no longer work.

Non-descriptivism, or the reality of resource over knowledge: the idea is to provide effective methods to work with available affordances and constraints, which are the data that you get from the environment. However, whether those data have anything to do with the real, actual outcome and with nature is not our concern.

Economy refers to the careful management of resources. Constructivism seeks minimal ontological commitment, ie not concluding that things must exist in reality as in the model. We should be noncommittal about it. If it works, fine; if it does not, try harder. In terms of what there is outside in the world, the less you can commit yourself, the better.

The Turing Test revisited

I have given you an initial black-and-white picture. According to Plato, if you use it, you know it better, but according to the other philosophers discussed, Hobbes included and all the way down to Kant, it is the making that makes a difference. It is a matter of hands versus eyes. I then gave you some general principles which lie behind this constructionist view. This position is not uncontroversial and it is not the dominant view in philosophy, so far as I know. The history of philosophy has been written by the Platonists.

Here is a specific case to show you how this constructionist view and the principles that I have been mentioning might work in application. The Turing Test describes a kind of imitation game. Turing was not addressing the question of whether machines can think but he was trying to address a question which would then cast some light on the interaction between machines and humans. He wanted to understand whether something that we attribute to humans could also be attributed to machines in the long run.

The Turing test supposes that we have a person, a machine and an interrogator, in three different places. The object of the game is for the interrogator to determine which of the other two is the person. Turing was writing in about 1950 and, according to him, in about 50 years – round about now – it would be possible to programme computers to make them play the imitation game so that an average interrogator would not have greater than 30 per cent chance of making the right identification after five minutes of questioning.

He was being optimistic. This has not happened and, so far, we do not know when that might happen. Nevertheless, the test has become quite famous, and rightly so, as an interesting way of going about asking whether there is any sense in which a machine can do what we do, semantically speaking.

What is interesting about the test, from the constructionist perspective that I have sketched so far? The Turing Test respects the minimalist criterion. You try to assume as little as possible – you just have two agents, and you ask questions of them. It uses a specific level of abstraction, which means fixing the observables, which would then allow sensible questions to be asked. You cannot ask, for example, what the measure of something is without defining the observables – the sort of measures, scale and so on. So the level of abstraction is fixed and is done in terms of asking questions and getting answers without, for example, any discrimination in terms of bodily implementation.

It is also constructionist in the following sense. Turing refused to provide an answer to the question, 'Can a machine think?' He considered it not well-defined because it involves some vague concepts such as machines and thinking, and he was right. He proposed to replace it with an efficient game, which we have just seen, because it is more manageable and less demanding from the minimalist's point of view. In other words, you build the test and you know what to expect from it.

The level of abstraction in the question means that, if you change the rules, you can devise different tests. This has been done repeatedly in the past 50 years. You can have different Turing tests, depending on whether, for example, you are wondering whether the two agents out there can be creative. It is not about questions and answers but it is based on how creative they can be in producing a bigger text, and so on.

So Turing refuses to define what intelligence 'really' is, in itself. It is not an attempt to set necessary and sufficient conditions for intelligence. The maker's knowledge idea is that you do not try to go for the thing in itself, which – at this point – either you think is a pointless exercise or, in any case, you do not mind. You make an hypothesis and devise a system by which we can evaluate that hypothesis. That is the idea of the constructable simulation and that is why you have a test.

The system is fully controllable – we have built it and we know what to expect; we know the variables and we know the level of abstraction and the observables. We are in control – with respect to this particular test, we are in God's shoes. That is the idea behind the controllability principle and you can even adapt and change whatever you like. The fact that a machine passes or fails a test implies the fact only that the machine can or cannot be considered intelligent at that particular level.

There is a step before the answers, such as what kind of questions does it make sense to ask at this particular level of abstraction. Turing does not consider those: it is as minimalist as you can possibly make it. So you try to make it simple, constructible, controllable – whatever you end up with is something that you can consider applicable to the simulation, and not to the stuff that is there. You are not trying to define, in this case, intelligence in itself, independently of instruments, measurements, precision, circumstances or any other things that interact in this particular case.

Conclusion

It is definitely true that some questions cannot be answered by Google. Certainly, the Turing Test has failed and there are many reasons why this is not the topic of this particular talk. However, I could conclude with a quick review of constructionism, the Platonic view, the anti-Platonic view, the tradition, the principles behind it and the example provided, by giving you a quick summary.

In this quick tour, I have tried to make you friendly towards the view that there is a good deal more contribution in our knowledge of the world than sometimes some philosophers like to acknowledge. We are not empty baskets waiting for input from the world – we contribute a good deal in our construction of the world.

On this view, to know, or to have knowledge of reality, is to be able to wonder, to ask questions and to deal with semantic artefacts. In other words, it is to ask questions and provide answers that can satisfactorily and successfully address your wonders.

Let me come full circle. Plato's Cratylus is someone who is defined as the man who knows how to ask and answer questions said that I was being unfair to Plato and I felt I had to pay my debt towards the end.

Part 2: Engineering, ethics and the environment

2.1 Engineering sustainability: Synergies between science, precaution and participation

Professor Andy Stirling

Andy Stirling is Research Director at SPRU (Science and Technology Policy Research Unit) and a co-Director of the ESRC 'STEPS' Centre (social, technological and environmental pathways to sustainability) at the University of Sussex. He has a background in natural science, a Masters in Archaeology and Social Anthropology (Edinburgh) and a D.Phil in science and technology policy (Sussex). His own work focuses on issues of uncertainty, precaution, participation, sustainability and diversity in science and innovation strategies. He has served on a number of UK and EC advisory committees, including the UK Advisory Committee on Toxic Substances, GM Science Review Panel and Science Advisory Council and the EC Expert Groups on Energy Policy, Science in Society and Science and Governance.

I am at an interdisciplinary unit called SPRU (Science and Technology Policy Research Unit) which is involved in all sorts of areas that relate to policy-making on technology and science. My interests are in uncertainty and processes of public engagement in the governance of science and technology. I cannot think of any area more intensely evocative of ethical issues than questions over how to make decisions under uncertainty; and how to go about (and in what manner and under what conditions) engaging different groups in society in long term, high stakes processes of choice over the directions of science and technology.

Ethics is right at the heart of what I want to talk about. Being neither a professional ethicist nor an engineer, I will not try to take positions in respect of those disciplines in what I say. I will instead outline a series of challenges which I believe are relevant to the topic of environmental ethics in engineering.

I shall say a few words at the beginning about the big picture, about connections between precaution, science and technological progress. What are the links between these things? They are often held to be in some tension. I will argue the case that the synergies between science, precaution, participation and progress are actually more important than the tensions. These are things which perhaps many feel are in contention with one another.

I will begin by setting out some of the less well-recognised implications and then look at the problems of so-called 'science-based decision making' – using science as a basis of decision-making on technology and engineering choices. At the end – and perhaps this is unusual in a discussion about philosophy and ethics – I will look at some practical lessons for technology policy that come out of what I say. In trying to convey the practical lessons, I hope that will make clear the real, key features of what I am talking about.

The challenge of precaution

'Precaution' is in my title, so let us take a look at what precaution is all about. I am sure that many readers are very familiar with Principle 15 of the 1992 Rio Declaration – which is the grandaddy of all the many later versions of the precautionary principle which pop up in different regulations around the world:

> "Where there are threats of serious of irreversible damage, lack of full scientific certainty shall not be used as a reason for taking measures to prevent harm."

This ethical principle is very ambiguous as a decision rule. There is an extensive literature dedicated to interpreting the principle. What do we mean by threat? How serious is serious? Exactly what constitutes irreversibility in any given context? And, oh my goodness look, 'full scientific certainty' appears there! There is a well-established school of thought here which says that this is unworkable and irresponsible and it undermines the prospects of real technological progress. To the extent that that leads to all sorts of adverse effects, this would also be unethical.

However, these are not just ethical debates but they are very high profile political debates which are exercised in all sorts of fora such as, for instance, the Kyoto process; in Bali; in the World Trade Organisation in negotiations about food safety and major high-profile risk issues on chemicals; and, of course, in discussions on genetic modification. In all these areas, it is the ambiguity of this formulation of precaution, and the ambiguity that seems to attend recognising uncertainty in decision-making, that is held to be unfavourable by comparison with tried and tested, well-established, sound scientific approaches to making decisions.

There are approaches like risk assessment, cost-benefit analysis, life cycle analysis, environmental impact assessment and so on. These are held to provide this clear, crisp, determining firm basis for decision-making. I shall return to those issues in a little while.

Science and technological progress

Before I do that, however, I want to look at what goes on in mainstream political debates, when people are confronted with big, set-piece decisions on new technology. One example is that of genetic engineering. There have been few better examples in recent years of this kind of dilemma for society about new technologies, than with genetically modified food. Here, we see Tony Blair when besieged in Parliament to justify his intended policy in 2003, saying: "This Government's approach is to make decisions on the basis of sound science."

More recently, when justifying the government's approach to the energy debate and, in particular, the position on nuclear power, the then energy minister Malcolm Wicks said: "Now is the right time for a cool-headed, evidence-based assessment of the options open to us. I want to sweep away historic prejudice and put in its place evidence and science." The counterposing of evidence and science and historic prejudice, irrational, anxiety-ridden public sensibilities is very much the background to these ways of talking.

In these examples, we are seeing essentially political decisions about new technologies which are presented as not being a matter for ethics, but a matter for science and science-based decision-making. On a larger canvas, when we speak about technology in general, throughout Europe, North America, and of course other industrialised and industrialising countries, we constantly see reference to the need for pro-innovation policy. There are references to the public being 'anti-technology', and that we need to restore pro-technology viewpoints. In these discussions, innovation and technology are treated as homogenous. There are no distinctions and there are no real alternatives and therefore there is no politics and it is not about choice. That is what is going on in this type of debate, when this sort of language is used.

The nature of technological progress

The underlying view of technology here is that, despite the fact that those who profess such views would, I am sure, disavow this picture, there is a single, linear path that technology takes, which is determined by science. Lord Broers' Reith Lectures had the great virtue of being really upfront about these issues. We hear from Lord Broers that "technology will determine the future of the human race" – and interestingly, not the other way around, that "history is a race to advance technology"; that "the aim of politics is to strive to stay in the race"; and that "the role of the public is to recognise and give technology the profile and status it deserves." These quotes illustrate my point, I think, about the deeper implications of this type of language.

Lord Mark Malloch-Brown, at the time when he was UN Deputy Secretary-General, encapsulated quite nicely what mainstream debates do with this sort of picture. He refers to "anti-technology protestors" as if they were in blanket terms against all technology. They are "members of the flat earth society, opposed to modern economics, modern technology, modern science and modern life itself." So the whole thing is reduced to being for or against modernity; for or against progress; for or against the future. This is the way that these debates over technology become polarised, even by very influential diplomatic voices, such as that of Lord Malloch-Brown.

That means, under this kind of view – and you may agree that this is pretty much the established way of talking about these things in high-level policy-making, for instance under the Lisbon agenda in Europe – the politics of technology reduces to questions simply about how much technology (not what sort, but how much), how fast, and how efficient and, if we are really radical, how fair? That is to circumscribe the main ethical choices and ethical decisions about technology.

The interesting point about this is that, for science-based or evidence-based language, it is remarkable because this is simply not how it is. This is not how technology is. Every single discipline, be it economics, history, social studies, philosophy, which looks at the processes through which society generates technologies, agree that that is not how technology is. Retaining that axis of time, I will try to show you a picture of the common denominator of these different disciplines.

If you imagine the space of technological possibilities, the received wisdom is to determine if the initial multiple potentials converge on an inevitable optimal form for the technology over a period of time. This may be in the transport sector, or within the energy sector, materials and so forth. That is the received wisdom. Actually, however – and this is a celebration of technology and not a criticism of any of the sectors I have mentioned – there are many different and equally viable possible technological pathways. The great truth about innovation is not that it is circumscribed and stultified so that it comes up with one inevitable solution, but it is fertile with creativity and social agencies, and constantly open to surprise.

Technology as social choice

Just to rehearse some very well-worn examples yielded by these different disciplines, time and again we see processes of path-dependence and how things could have been different but were not, due to contingencies, chance events in the early process of innovation, momentum, lock-in and all sorts of other mechanisms. Audio/video formats are one example, with the battle between VHS and Betamax. Windows software is an example of lock-in: I am using Windows software and I will blame any error I make on this very point about lock-in to suboptimal technology! Narrow gauge railways, urban transport and the internal combustion engine; there are well-established historical accounts relating to all of these which show that, without necessarily being pejorative about the configuration of the technologies now, it is possible to see all sorts of different pathways through which they might have evolved, sometimes to very different endpoints. The QWERTY keyboard, which we still use, is a classic example of these processes of lock-in.

With light water reactors – without being pro- or anti-nuclear – there are all sorts of arguments about whether Rickover and the very high-pressure (if you excuse the pun) environment of developing technology for submarine propulsion were actually giving us basic design traditions that were optimised for civilian power production. And in industry as such, once these crucial design philosophical decisions have been made, it is very difficult to escape from them, even if we can see the possibility of different radical designs that might have yielded better outcomes for civilian power production.

I hope you get the point from this that we face these different pathways but not all possibilities can be fully realised, especially in globalised markets where we have increasingly integrated concepts around the world. We face these kinds of decisions, looking forward, in choices between centralised thermal power, distributed energy and renewable energy. There are major choices about how we run electricity grids, without taking positions on set-piece arguments for and against nuclear or renewables. There are huge issues around the configuration of grids that are being played out at the moment.

Industrial, input-intensive farming, versus more locally-based, locally-marketed, low-input farming, is a very high technology, very demanding area and poses choices. In chlorinated plastics, recycled materials, energy recovery, choices are being made between those pathways. There is fossil fuel in urban transport, and choices between different propulsion systems, primary resources and different sorts of motors and so forth are being made. You can imagine a number of these actually proving viable in the long run, if they benefit from long-term support by markets. If they do not, they will not cut it, while the ones that benefit from that support will.

I hope I have made the point with sufficient practical flavour that we face constant choice in the technology path to follow. The outcome of these large technological choices is based partly on necessity. I do not want to try to make a relativist argument that 'anything goes' because there are all sorts of things that will not work – and nobody knows that better than engineers. The point is that, just because there are plenty of things that will not work, that does not mean that there is only one thing that will work a given sector.

Behind this rhetoric of science-based decisions on risk or on technology policy, there is a hidden politics. Expressions like being 'anti-science' or 'pro-technology' are as invalid as if we heard, in health or education debates or debates on criminal justice or the military, people saying when they hear a dissenting voice, 'Well, you're just being anti-policy. I can't deal with people who are anti-policy.' It is nonsense to say that someone is anti-policy, because the whole point is that there are different types of policy. Yet, when it comes to debates about technology, we hear just that language: 'You're anti-technology', when the point is, which technology?

The limits to 'sound science'

Going back to precaution now, one might still hear that, despite the arguments above, at least these science-based approaches offer a robust basis for starting these decision discussions and being rigorous about it. At least we have these clear, rigorous pictures of what the various pros and cons are. But is that so?

To illustrate this point, I will now take a look at a field where I have done some work. I believe the energy sector is an area where it is arguably the case that none of these techniques – life cycle analysis, risk assessment, cost/benefit analysis – have been more rigorously or comprehensively applied to informing choices between technologies. They are very rigorously applied there and I do not mean to denigrate the work that has gone on. However, when you look at that picture in the energy sector – and this would apply even more in the transport, chemical or biotechnology sectors – individual studies typically represent their results with a fairly fine degree of precision. The studies I am interested in here express the risk due to different technologies in a very broad sense in terms of monetary externalities – money values per unit of electricity production – coming from the different energy technologies, electricity supply technologies on the vertical axis (Figure 8).

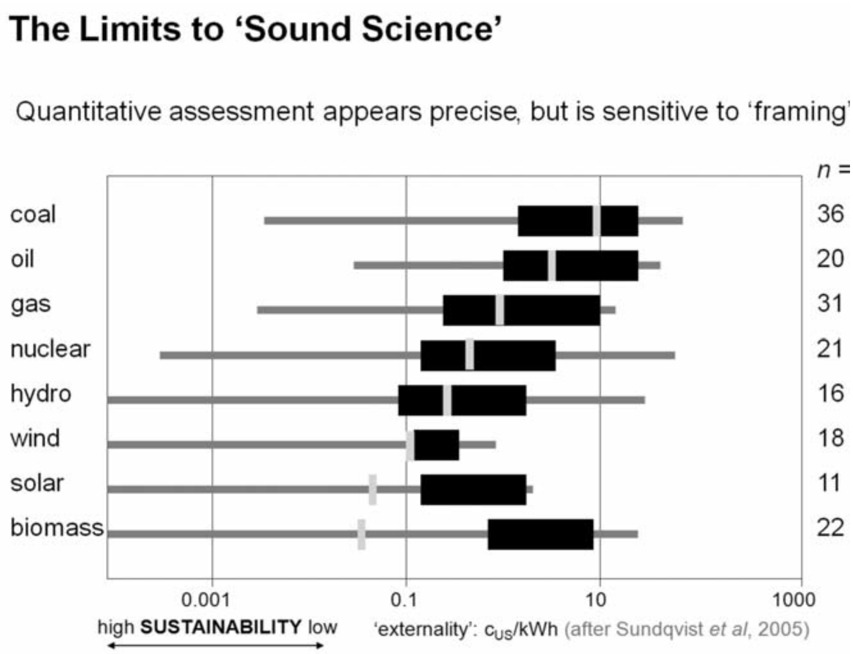

Figure 8. Diagram showing the distribution of money values per unit of electricity production for different energy technologies 'n' is the number of studies used to calculate each distribution.

This particular study, shown above, has the values for each technology, defined by the intervals. For a selection of 36 studies, all looking at coal power, this is the distribution. The differences between these relatively high and relatively low estimates of the risks due to modern coal power – and modern coal power, not just the whole range of stations actually – across these 36 government and industry studies, is more than five orders of magnitude. If you look at the underlying risk assessment literature, or energy input analysis literature, or land use analysis, you see an essentially similar picture.

Here is the distribution for the fossil fuels, nuclear and hydro, and the renewable technologies, again based on the number of studies in the right-hand picture. Individual studies typically show a really admirably precise picture but the literature as a whole gives an enormously varied and, I would say, ambiguous picture. What is going on here is the framing of the analysis, which is as important as the answers delivered from it.

The nature of 'framing'

By framing, I mean a whole series of different issues: setting agendas, the system boundaries, the baselines, the particular methods, discount rates and so on. These conspire together to yield that enormous variability. The principles and methodological rigour in any of these fields tell us what to do within a set of framing assumptions but they do not give us unequivocally a completely rigorous set of framing assumptions. Even if they did, there would be no guarantee that particular conventions adopted in the literature reflect the value judgments extant in the wider society where, for instance, one of the parameters might be how important do we regard it to avoid adverse effects on children compared to older people or compared to future generations or to animals? That is just one parameter, much beloved of ethical discussion. Exactly what is a sound, scientific answer to that? It is very difficult to justify that there is such a thing.

A simple way of putting this is in the current debates about nanotechnology regulation, for instance, though the same is true in chemical or biotechnology regulation generally. Are we asking the question: Is this safe? Or is this safe enough? Or what is the safest option? Or what would be the best option? You can see examples of these kinds of questions underpinning regulation in all sorts of different areas. They yield different answers. Sound science, or evidence-based decision-making, cannot tell you which is the right question, but it will certainly give you different answers, depending on the questions that you ask. Typically, when we hear Tony Blair or the other protagonists I mentioned in the beginning talking about these decisions, you do not hear any of this on the surface at all.

Thus, all analysis requires framing. This is not a criticism of analysis but it is just the way it is. And all framing involves value judgment.

Risk and responsibility

What are the implications for the ethics of making decisions under uncertainty, where we do not have complete knowledge? For engineers, I hope it is useful to start with the framework under which conventional approaches to risk begin (Figure 9). There is the distinction between things called likelihoods and those called outcomes. The outcomes are the things that might happen while the likelihoods are the probabilities with which these things might happen.

Our knowledge, on either of these two dimensions, can be relatively problematic, or actually relatively good.

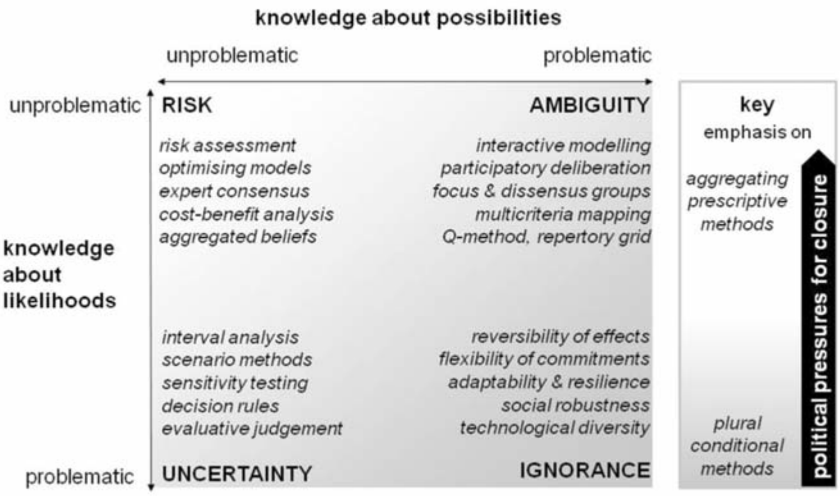

Figure 9. A framework for likelihoods and outcomes when considering issues around which there is uncertainty.

Where we feel that we are in a position of having relatively unproblematic knowledge on both axes, then we face the classic condition, which is the rigorous definition of the term 'risk', which is used very loosely in all sorts of ways in wider society and even in science. There are many very powerful techniques for dealing with this, with which we are familiar. Situations like engineering failure, the epidemiology of known pathogens, transport safety issues where they are looking at well-established systems, and flood under normal conditions, are examples where we might say that this is the dominant condition that we find ourselves in.

When we find ourselves in a position where the outcomes are relatively well-defined but the probabilities really are not so well-defined because we are facing new types of technology, new situations, new contexts, dynamic shifting environments, complexity and so on, then we see the definition of uncertainty coming in. This situation arises in areas like unassessed chemicals and new pathogens that may be becoming widely known, like avian flu, for instance. Climate change impact in particular areas also falls into this category since, even if we think we understand the overall picture at a global level, in terms of the implications for specific regions it becomes much more a matter of uncertainty.

That is as far as the debate quite often goes. It is quite controversial enough to appear to compartmentalise the applicability of probability methods in this way. However, the problem becomes worse than that – before I return to the practical implications – because there is the other axis, where not just the likelihoods but the outcomes themselves are a matter of dispute, or are problematic in some way. There are problems of apples and oranges, about which John O'Neill has written very powerfully in past work, where we are trying to compare different things and where even the principles of rational choice break down. We have Nobel Prize winning work showing that actually not only is it difficult to make unequivocally rational decisions in comparisons between apples and oranges, different things as it were, but it is actually meaningless in principle even to claim to do it. This is because it is difficult even to say what a rational decision would look like where we face these types of different plural values. There are issues of trust and issues of perception, meaning, exclusion, distribution of issues and so on, and these all come into this – in fact, the really high-profile ethical bases of decision.

Then you will see this melodramatic lower right-hand corner. Unfortunately, the term used is rarely discussed but, where it is discussed, it is called ignorance. This is the formal definition of ignorance: it is where we face surprise – Donald Rumsfeld's famous "unknown unknowns". However, in retrospect, this condition has been the dominant state in many of the most important environmental issues that we have had to contend with. BSE, CFCs, endocrine-disrupting chemicals – these are not examples of issues where, at the beginning, at the advent of the issue, those concerned with managing them were saying, 'yes, it is possible, of course, that we might see pathogens transmissible through this new industrial process that will actually affect brain function, but we think it is very unlikely.' Or, 'It is possible that these apparently benign, inert chemicals will, under certain conditions, start being violently reactive, or it is possible that there is an entirely new mechanism for toxic interference in biological systems which is not conventionally toxicological but it is a more mechanical coordination problem. These things are possible but unlikely.' The very possibilities themselves were not contemplated at the beginning of each of those stories. They were a matter of ignorance and not uncertainty or ambiguity, and certainly not risk.

Before I get too depressed about it all – social scientists are always coming over in a rather bleak fashion – ignorance is not just a source of bad surprises. It is a source of innovation. There are lasers and high temperature super-conductors, fullerenes, and so on. There are all sorts of amazing surprises which have all sorts of positive applications that we come across over the years, but they are a matter of ignorance.

Conventional approaches to these conditions are basically to say, yes, thank you for the philosophy, but we have a hammer and every problem is going to be a nail. We have risk assessment techniques that are formally only applicable to restricted cases but we will use those methods generally.

Towards more responsible governance of science and technology

The point is that, as with Cinderella, there are plenty of other methods that are not invited to the party but which are applicable under these different circumstances. So where you have uncertainty, there are very sophisticated forms of interval analysis, sensitivity analysis or scenario analysis. They are very well tried and tested but they do not deliver unequivocal final answers, although they give you a very much more robust picture of the problem you face. Economists hate them, but engineers are pretty good at sensitivity analysis, although they become less good as they get closer to policy debates.

Ambiguity: there are all sorts of social science approaches for eliciting and deliberating over competing value judgments and different framings. Even in cases of ignorance, where we do not know what we do not know, there are all sorts of things we can do. We can be more adaptive and have adaptive management, or horizon scanning. We can put a premium on research and monitoring because often these suffer at the expense of modelling – but modelling only tells us, to some extent, what the end location is. Research and monitoring are good ways of finding out what we do not know.

There is diversity – not putting all our eggs in one basket, and flexibility and resilience. We may not know that something is the best option for sure, but we are choosing between two engineering options. One of them, if we learn that it is actually ill-advised, can be retreated from more easily than from the other. If we have full confidence in our risk assessment, we are not concerned with this ability to retreat. It is only when we grasp the importance of uncertainty and ignorance that we start to value that property of reversibility. I am not talking about reversibility of effects here, but about reversibility in terms of how easy it is to withdraw from a commitment if it should turn out to be bad. That is the implication, for instance, for modular, small-scale renewable technology compared to large scale commitments in, for instance, the nuclear infrastructure.

The depths of incertitude

This is where we come across precaution because classical prevention strategies presuppose that we can characterise risk. Precaution is simply a recognition of these more intractable states of incertitude. Forgive the terminology here: incertitude is an English word and I use it because, often, the loose usage of words like risk or uncertainty, as if they cover the whole ball park, is taken as an excuse for then thinking that, if a method is applicable to risk, it must be applicable to all these situations. I would like to try to get away from that, although that involves trying to be a little more pedantic about language.

There is nothing about precaution that implies being anti-scientific: precaution is being rigorous about these other states of incertitude, recognising that sound science should be on tap and that it is essential, but that it cannot deliver answers under those conditions.

Lessons for precautionary appraisal

There is a series of quite practical lessons that arise from this about how we might make more robust decisions in the face of these types of dilemmas. These rely on us moving away from a formulation of precaution like the one that I began with, where it simply purports to be a decision rule – pretending, if you like, that you can identify what would be the right decision in an engineering context. Instead, it focuses on the process to be followed – the processes that ensue from realising that uncertainty poses a series of challenges that are not actually addressed in risk assessment.

First of all, therefore, you extend the scope of the appraisal beyond the issues. There is a tendency, in toxicology, for instance, to address those aspects that are best dealt with by the toxicology data that we have. Basically, they say that these are proxies for the other things – we can measure these things and so we will focus on them.

Instead, however, if you start to say that you are moving beyond where probabilistic claims apply, you start looking more broadly. You can look at synergistic effects, cumulative effects, life-cycle issues, compliance issues – much broader issues about the technologies in question, which tend to be missed out. By the way, here I am referring to a study by the European Environment Agency, called Late Lessons from Early Warnings, which goes through a series of different examples of where these lessons came home in retrospect.

For instance, in the story of chlorofluorocarbons, we were not looking at a whole system but we were simply looking at the rather proximate kind of issues in the technologies concerned. We were not looking at the wider environmental distribution once the chemicals had been released and got into the upper atmosphere, because there was in fact Nobel Prize-winning work which showed that, in the upper atmosphere, chemical processes would lead to degradation of stratospheric ozone. However, this was not part of the regulatory picture, which was about toxicology. So look more broadly and do not be confined to where you think you have the best datasets. Be more humble about what science can tell you – that is where sensitivity analysis comes in.

It is remarkable that engineers, for instance, are actually very well-practised when characterising a structure made of steel or concrete, in thermal conductivity for instance, and using sensitivity analysis to see how this relatively simple parameter will in fact be very complex in a structure like that. However, when it comes to making decisions on the implications of genetically modified technologies, or energy technologies in society at large, the far more complex systems that we are considering are routinely not dealt with by sensitivity analysis. So be more humble about what the science tells us. Be more active about focusing research and surveillance on the issues in question. Stories like BSE were characterised for more than a decade by persistent refusal, which the Phillips inquiry picked up, to actually commission research which would answer the questions that were being posed. We did not actually commission that research because there was a fear that the answers would lead to undue public anxieties and therefore we did not want to know those answers.

Be more deliberate about arguments: do not treat these things as being delivered. Issues like the level of proof; or what weight of evidence do we accept as being persuasive; or who has the onus of persuasion: is it those who wish to introduce a technology or those who have concerns about it? These things are not handed to us on tablets of stone but they are ethical issues which need to be deliberated upon, and they will have different answers in different contexts.

For instance, there was a very interesting committee in the 1960s, the Swann committee, on the use of antimicrobials in animal feed which deliberated very closely on these issues. The committee came up with a rather cautious response which, over the period of the 1970s, was progressively eroded. Now we find ourselves, in some countries more quickly than others, desperately tracking back on those decisions that were in fact made by the British Swann committee back in the 1960s.

Look at different options: in matters of regulation of technology, do not try to decide whether this or that option is satisfactory in its own right, but compare them across the board. I have served on a number of policy advisory committees and I have seen time and again that the game is, 'Let us think. Here, we have a particular chemical for the workplace or a particular product for the agricultural sector. Let us wrack our brains to see whether we can think of anything currently on the market that is actually worse than this. If we find one, then it is fine.' That sounds a little unfair but, if I had more time, I could go into many examples like that. That is basically the received practice. Instead, however, we should be looking at contending options and say, 'OK, we are uncertain about these things but, on balance, which looks preferable?'

Interdisciplinarity: often the stories on these issues, as they are captured by the particular disciplines, and engaging the public – this is not because of some romantic idea that the public are more in a position to make dispassionate judgments but quite the reverse. When we face a problem where we have different ways of framing the issue, with different value judgments and different interests in play, and they give different answers, we should be as rigorous about validating the framings as about validating the data.

Conclusions: precautions, participation and technology

To conclude, I do not think there are any ethically irrevocable resolutions for these types of dilemma. The answer has to lie in greater reflection, more open deliberation, more readiness to acknowledge these types of dilemma, and to avoid certain pathological behaviours.

First of all, there is the false dichotomy that we currently have between 'rational' sound science and 'emotive' precaution and participation. That is not a valid dichotomy. The two things are inextricably intertwined and we need to be more grown up about the ways in which that is the case.

We should avoid 'scientistic' language. This is not to criticise the power of science, but invoking science, as though it necessarily gives us an answer (which scientists rather rarely do but which politicians do a great deal), should be avoided because it obscures the fact that we face different directions for choice and because it undermines the rigour on uncertainty and on accountability.

In precaution, the rather simplistic formula that is sometimes propounded by environmentalists – and I say this as an environmentalist – only gets you so far. The real value is in the process. Public engagement is not about education or about fostering trust for the decisions that we know are good, or about political correctness, but it is about rigour, in the ways that I have tried to outline.

There are many different ways to implement this in policy-making and I have tried to show the range of different methods that tend to be obscured at the moment. We have a number of practical approaches to broaden out and open up policy advice, make it more plural and acknowledge the way that there are different possible answers that you can get from the science. We therefore need to be more political, not less. In other words, we need to be more accountable and more transparent about the value judgments which, in the end, hold sway.

In the end, that is a remarkable conclusion because, against this rather bleak and complex background, rather than being in contention, imperatives of scientific rigour and of democratic accountability – the two big set-piece ethical considerations in this debate – are pointing in the same way. They are both saying that we should deal with the issues in this more precautionary way, in this more participatory way, but without throwing the baby out with the bathwater. We should be more mature about the way the scientific techniques inter-relate with these different issues of values and framings.

2.2 Engineering and environmental values

Professor John O'Neill

John O'Neill is Hallsworth Professor of Political Economy at Manchester University. He has written widely on philosophy and political economy, political theory, environmental policy and philosophy, ethics and the philosophy of science. His books include Markets, Deliberation and Environment (Routledge, 2007), The Market: Ethics, Knowledge and Politics (Routledge, 1998) and Ecology, Policy and Politics: Human Well-Being and the Natural World (Routledge, 1993). He co-authored Environmental Values (Routledge, 2008) with Alan Holland and Andrew Light. He is co-editor with Tim Hayward of Justice, Property and the Environment: Social and Legal Perspectives (Ashgate, 1997) and with Ian Bateman and Kerry Turner of Environmental Ethics and Philosophy (Edward Elgar, 2001).

Engineering and the natural world

The use of the concept of engineering in the biological sphere often elicits very strong negative ethical responses. Why do people have such ambivalent feelings about engineering when it comes to the environment? Given that engineering and technological innovation will be central to the solution of environmental problems, what is the source of the deeply ambivalent attitude many take towards engineering in the biological sphere? The very juxtaposition of the concepts of 'biological' and 'engineering' in talk of 'bioengineering' or 'genetic engineering' appears to invoke worries that are not apparent in other spheres. The concept of 'engineering' appears to carry particular negative connotations. Consider for example the following from, The 2005 Euro-barometer:

> "The terms 'biotechnology' and 'genetic engineering', as in previous years, appear to have different connotations for the public. 8 per cent more Europeans see 'biotechnology' as likely to improve their way of life in the future than those asked the same question about 'genetic engineering'. In 1999 the difference was 8 per cent, in 2002 it was 5 per cent. The more positive connotation of the term 'bio', perhaps a result of the association with healthy and natural things, contrasting with 'engineering,' with its connotations of manipulating or tampering, holds across much of Europe with the exception of Spain, Italy and Malta. Perhaps most striking is the 'lead' of biotechnology over genetic engineering in some countries: it is more than 20 per cent in Belgium, Denmark, Germany, Finland and Austria." (Euro-barometer 2005 p.11)

Standard articulations of worries about 'engineering' in the biological context includes claims that bio-engineering is 'unnatural', that it involves 'playing God' or it is hubristic, failing to acknowledge the limits of human knowledge and capacities to control the natural world.

In this paper I examine some of those objections to see whether there is anything in them. My purpose is not to endorse them. I think any answer to questions of the justifiability of different forms of bioengineering or genetic engineering needs to engage with the empirical detail. However, what I do hope to do is to articulate what might be at stake in these objections. In doing so, I suggest that these objections tend to be dismissed too quickly in the debates by both philosophers and policymakers.

This paper is in three parts. First, I discuss the unnaturalness objection to biological technologies. I consider the standard responses to the unnaturalness objection to be found in the literature and examine weaknesses in at least some of these responses. I attempt to give what I take to be the strongest formulation of unnaturalness objections and suggest that there is more to it than is often supposed. However, I will leave it open at to whether the objection, in the end, can be sustained.

Second I consider some of the considerations around knowledge and control. However, I will come at this from a slightly skewed angle, by looking at the parallels between objections to engineering in the biological sphere and objections to talk of engineering in the social sphere, in terms of social engineering. Talk such as that of Stalin about artists and writers being 'engineers of the soul', brings people up sharp, eliciting a similar kind of response to biological engineering. The question is, what kind of arguments are going on there? I run through Hayek's arguments against social engineering and see what parallels can be drawn.

I finish by looking at the idea of engineering that runs through those debates. How is engineering conceived of and is engineering being done a disservice in the way that it is actually understood? I suggest that one problem with the debates is the conception of engineering that both sides often share. In particular there are forms of precaution that are implicit in the practice of engineering itself that are missed on both sides of the debate.

1. Unnaturalness objection

Euro-barometer responses suggest that the unnaturalness objection is one of the most common objections to genetic engineering (Euro-barometer, 2000). There are two standard responses to the unnaturalness objection to biotechnologies such as genetic engineering. The first is that everything is natural, and hence no technology or impact on the natural world can be described as 'unnatural'. Consider, for example, Richard North, writing in *The Independent*:

> "Man can never damage nature because nature is a set of scientific facts by which man can live or die. Man never has dominion over natural laws." (North, 1991)

The second is that there are no essential differences between old and new methods of crop developments. If you look at the traditional techniques for the development of new varieties of crops by processes of artificial selection, they are 'unnatural', so what is the difference between the new genetic technologies and the old technologies? This was the position articulated in Nuffield Council reports on genetic engineering (Nuffield Council 1999; 2004, 3.7-3.17). In the following, I will consider these two responses to the unnaturalness objection and argue they are not as clear-cut as they first appear.

'Everything is natural'

Consider first, North's 'everything is natural' objection. The objection confuses different senses of 'natural' and 'unnatural' that are raised in the debate.

Hume, in *A Treatise of Human Nature*, says of the word 'nature' that 'there is none more ambiguous and equivocal' and he points out that 'nature' is used in a series of contrasts. One contrast is 'nature' versus 'the miraculous', and a second is 'nature' versus the 'artificial'. Third, there is 'nature' versus the 'civil' or the 'cultural', and finally there is 'nature' versus the 'rare' or 'unusual'. It is the first two which really matter.

Taking first of all 'nature' versus the 'miraculous', or the 'super-natural', Hume says:

> "If *nature* be opposed to miracles, not only the distinction betwixt vice and virtue is natural, but also every event, which has ever happened in the world, *excepting those miracles, on which our religion is founded*." (Hume, 1739 III.i.ii)

A point to note about Hume in the background here is that Hume is the great sceptic about miracles: Hume is defending a version of naturalism. Once you reject the idea of the supernatural, everything is natural in this sense. So it is true that, in that particular sense, everything is natural.

The problem for the 'everything is natural' response, is that it is simply not the sense that the critic of bioengineering is appealing to. He is appealing to one of the more specific senses of natural and not this general sense of the natural opposed to the supernatural.

The specific sense that we are concerned with is the contrast between the natural and the artificial. The person who develops this further is John Stuart Mill in this essay on nature:

> "It thus appears that we must recognise at least two principal meanings of the word 'nature'. In one sense, it means all the powers existing in either the outer or the inner world and everything that takes place by means of those powers... (Mill, 1874, pp.8-9)."

This is natural in the sense appealed to in the idea that everything must be natural.

> "In another sense, it means, not everything which happens, but only what takes place without the agency, or without the voluntary and intentional agency, of man. This distinction is far from exhausting the ambiguities of the word; but it is the key to most of those on which important consequences depend." (Mill, 1874, pp.8-9).

Hume's claim that the word 'nature' is equivocal is significant for a number of debates and there are other distinctions to be made between what is natural and what it not. However, in terms of the debate on bioengineering, the contrast of the natural and the artificial is the important distinction.

The concept of the artificial

How is the contrast between the natural and the artificial to be understood? Mill characterises the artificial in terms of what is the result of "the voluntary and intentional agency of man". This is clearly a necessary condition for something being artificial but it is not sufficient. The existence of children is the consequence of the voluntary and intentional agency or men and women. However, we do not want to call my children for that reason 'artifices'. They are the product of my voluntary and intentional agency and yet they are natural beings, so I do not think Mill's account yet captures what we mean by the artificial.

In our recent book, *Environmental Values*, my co-authors Alan Holland, Andrew Light and I make a further distinction between intentions, which bring something into existence, such as having children, and intentions which determine and shape the nature of a thing and which give it the properties that make it what it is (O'Neill et al). There is a difference, for example, between us bringing a child into existence, and us evolving various kinds of bioengineering, which would allow us to actually create a certain kind of child with certain properties and qualities.

On this view something is artificial if, and only if, *it is what it is* at least partly as the result of deliberate and intentional action, usually involving the application of art or skill. So being artificial is a matter of having a nature which is determined by deliberate and intentional action. This is the concept that I want to use for the remainder of this paper.

One important use of this contrast between the natural and the artificial is that of Darwin, who makes the distinction between natural and artificial selection. In artificial selection, you are artificially selecting to determine the nature of something and you are trying to create a certain kind of being via certain means of artificial selection.

However, Darwin's use of the distinction invites the unnaturalness objection to genetic engineering, the idea of which is developed in the Nuffield Council reports. There is no essential difference between modern bioengineering and traditional techniques for the development of crops and animals by a process of artificial selection. Artificial selection has existed for a long time with selective crop breeding and selective animal breeding. Traditional forms of agriculture are themselves unnatural in this particular sense and so there is no difference in kind here. Unless you want to object to all traditional methods of crop selection, then there is nothing to object to in terms of unnaturalness in respect of the new agricultural techniques. That is essentially the line of the Nuffield Council, and it is the typical philosopher's line. One can clearly see why that argument is persuasive.

However, a possible reply is to say that there is a difference between mimicking natural selection and replacing natural selection. The argument runs that first generation Mendelian artificial selection mimics natural selection and that is why Darwin could evoke it in helping to explain natural selection. It involves genetic modifications that do not occur in nature for simply contingent reasons. This is different, the argument would go, from *replacing* natural selection. Contemporary techniques for genetic modification, where they cross species boundaries, involve processes of selection that could not happen in the absence of human intervention. Where previous forms of artificial selection mimic natural selection and select from a range of possible natural lineages, genetic engineering would determine what lineages are possible, thus replacing rather than mimicking natural selection. When and how far this is the case is an empirical claim which is open to discussion.

There might be something in the distinction between mimicking and replacing natural selection. When the ordinary person says there is something unnatural about genetic engineering, this might be a way of developing that thought. However given this is right and genetic engineering is unnatural in this way, why does it matter? What is the ethical significance of it and why should degrees of artificiality or naturalness matter at all?

Naturalness as a source of value

One response to these questions is the claim that naturalness itself is a source of value. Goodin articulates the thought as follows:

> "According to the distinctively [green theory of value] . . . what it is that makes natural resources valuable is their very naturalness. That is to say, what imparts value to them is not any physical attributes or properties that they might display. Rather, it is the history and process of their creation. What is crucial in making things valuable, on the green theory of value, is the fact they were created by natural processes rather than by artificial human ones. By focusing in this way on the history and process of its creation as the special feature of a naturally occurring property that imparts value, the green theory of value shows itself to be an instantiation of yet another pair of more general theories of value - a *process*-based theory of value, on the one hand, and a *history*-based theory of value, on the other…" (Goodin, 1992, pp. 26-27).

Amongst some environmentalists, naturalness has this kind of source of value. One of the things that seem to distinguish a certain kind of green thought is that naturalness seems to matter. But why should it matter?

As Goodin notes, what makes something natural in this sense is the processes that created it. It is a backward-looking understanding of the natural. The concept of naturalness, in that sense, is a historical, spatio-temporal concept. Being natural is not constituted by the state that a thing is in, and it is not a matter of some identifiable set of characteristics. There could be two objects that are identical, but one deemed natural and the other non-natural in this sense, because it is the actual processes that created them that determine whether they are natural. So, whether something is natural or not is a matter of the processes that make it what it is and not its end state property.

Given this use of the concept of natural, if naturalness is a source of value, then one is presupposing a process-or history-based account of value. As such it is prima facie in conflict with consequentialism, which is forward looking. Correspondingly, a reason that standard risk assessment analysis cannot capture naturalness is that it looks at consequences. It does not take history or process into account.

However, this still leaves unanswered the prior question of why this should matter. What reason is there to assume one should take this historical perspective seriously? One starting point to a possible answer to that question is the observation that this historical source of value isn't peculiar to natural objects. There are a variety of objects where it is their history that matters, rather than its actual existing properties and features.

One can make a similar point about everyday objects, for example a hammer. There are great hammers that you can buy, perhaps at B&Q, or there is the one that was left to me by my father. That hammer came from his grandfather and it is hopeless, it has a loose head, but that hammer matters more to me because of its history. So there are certain objects which seem to matter to us because they provide a context for our lives. The same applies to a whole set of objects and buildings. Why is it that we would not, for example, imagine replacing Stonehenge with something else which looks identical? It is because it is part of the cultural history which gives sense to our lives.

Something similar is going on with the natural in this context, in that it provides a context in which we live out our lives and carry out our projects (cf. Goodin, 1992). Just as it would matter if someone came along and said, 'We want to put a road through Stonehenge but don't worry that we are going to destroy Stonehenge because I will build you an identical one over here', it matters to people if you put a road through a copse because it is that particular copse which has value. Many environmental problems are actually historical in this sense because they are looking back at what made something the way it is. History matters for the attribution of environmental value.

I have defended this account of environmental values in other works. What is not clear is how far it gets you with genetic engineering. There are two questions here. The first concerns whether and how far this account of the value of certain objects gets a purchase in the case of bioengineering. It works clearly for why places and particular environments matter – for example, the local copse or wood. How far can it be extended elsewhere, in particular to cases such as genetic engineering? The second question is, given it can be so extended, how do these considerations around the context for people's lives play against other significant considerations, for example around the meeting of basic welfare needs for a population? I will not pursue these considerations any further here. My aim has been more modest. It is to suggest that there is an argument to be had here. Even if turns out not be decisive, the argument from unnaturalness has more going for it than is often supposed in much of the philosophical and policy literature.

2. Engineering, knowledge and control

Here is a different line of argument. Engineering applied to the social world elicits similar kinds of negative responses to engineering in the biological context. For example, Stalin's description of writers and artists as 'engineers of the soul' elicits a similar response. Indeed, I remember many years ago being a member of a class discussing what it is to be a good teacher with the students coming up with a variety of descriptions and metaphors. Then one visiting Chinese student came up with 'engineers of the soul'. The tutor busy putting up the responses on a board blanched, and did not list this suggestion along with the rest. Why does it invoke this response?

Consider another case of a metaphor that has fallen into disrepute, that of social engineering. Social engineering was a popular term, with positive connotations in the early part of the twentieth century. It is now normally used in a negative sense. What were the objections that were introduced to engineering in the social domain?

To answer that question I want to focus on the debates on social engineering in the 1930s and 1940s. In particular I will consider the debate on the one hand between Hayek, the father of modern-day economic liberalism and critic of social engineering and socialism, and Otto Neurath, a positivist socialist who was happy to defend social engineering.

For Hayek, socialism is based on a rationalist illusion that is embodied in the figure of the social engineer.

The 'characteristic mental attitude of the engineer', leads to two errors when it is extended beyond its proper scope into the social domain:

1. A rejection of pluralism of ends in social life.
2. The belief that it is possible to bring together all the information that is required to optimally achieve this particular end.

First the engineer is taken to hold that there is a single end or hierarchy of ends. The engineer's task is to ascertain which of a set of possibilities provides an optimal means that can be directed to achieve that end: "he will be concerned with a single end, control of all the efforts directed towards this end, and dispose for this purpose over a definitely given supply of resources" (Hayek, 1942-4, p.167). Secondly, the engineer is taken to assume complete knowledge of his domain to achieve this end.

> "The application of the engineering technique to the whole of society requires… that the director possess the same complete knowledge of the whole society as the engineer possesses of his limited world. Central economic planning is nothing but such an application of engineering principles to the whole of society, based on the assumption that such a complete concentration of all relevant knowledge is possible." (Hayek, 1942-4, p.173)

Social engineering is taken then to require a single end or hierarchy of ends, and complete knowledge to achieve those ends.

Hayek claims that in the social world neither condition holds. Ends are plural. Complete knowledge is not possible. The belief that the social engineer can possess such knowledge is based on a mistaken belief about the omnipotence of reason – Cartesian rationalism. Against this view he suggests knowledge is limited: "it may … prove to be far the most difficult and not the least important task for human reason rationally to comprehend its own limitations" (Hayek, 1942-4, p.162).

One source of this ignorance in the social case lies in the division of knowledge in society, in particular the dispersal of practical knowledge embodied in skills and know-how, and knowledge of particulars local to time and place which cannot be articulated or vocalised in a general propositional form that could be passed on to a central planning body. Another source is the unpredictability of future wants and needs. These are, in principle, unpredictable, because they rely on future inventions. Here the engineer is herself or himself part of the unpredictability. Given that future innovation is not predictable – because if it was predictable, we would have it – then it follows that you will not know the future of what human needs and wants will be.

For those reasons, he thinks that you cannot collect all of that knowledge in one place. In consequence, the socialist dream of bringing the whole social world under conscious human control is a mistake. He says, in his Nobel Prize lecture, that the gardener rather than the engineer should be the model of social intervention:

> "If man is not to do more harm than good in his efforts to improve the social order, he will have to…use what knowledge he can achieve, not to shape the results as a craftsman shapes his handiwork, but rather to cultivate growth by providing the appropriate environment, as a gardener does for his plants." (Hayek, 1974)

The passage might appear to assume similar limits to control that natural world. Indeed applying these thoughts to environmental problems, it might look to entail a form of environmentalism.

However, a feature of Hayek's thought is that this acknowledgement of limits of knowledge and control applied to society is combined with an absence of any acknowledgement of similar limits applied to the natural world. As Gamble notes:

> "[Hayek] never extended to natural science and technology his critique of constructivist rationalism in social science. Although rationalism has retreated in the social sphere, it still has few restraints in its quest to master and control the natural world, posing increasingly serious questions for the civilization that Hayek so valued" (Gamble, 2006, p.130).

The reason he fails to extend his argument to the natural sciences is because he takes the problematic knowledge to be peculiar to the social. The interesting point in the background here is the image of engineering that Hayek is employing in criticising social engineering.

3. The image of engineering

Are Hayek's claims about engineering true? First, is it true that engineers are always concerned with pursuing a single end? Secondly, is it true that the engineer possesses 'complete knowledge' of his 'limited world'? Or is there a more modest account of the nature of engineering? And what implications would there be for applying engineering to the natural world?

This is where the 'bad guy', the positivist Otto Neurath, is actually a much more interesting thinker than the standard picture of positivism would suggest. Neurath defends social engineering but he does so based on a much more modest idea of what engineering involves.

First he claims that social engineering needs to be pluralist. It needs to recognise that there are multiple ends and not just a single end. Correspondingly, the job of the engineer in a democracy is to present multiple possibilities..

Neurath is also a pluralist in another cognitive sense. He thinks that there is normally a plurality of rationally acceptable scientific theories given some empirical basis. Given that, he thinks that there is an essential uncertainty and unpredictability in our predictions about both the natural and social worlds:

> "Even before the first world war I realized that acknowledging a kind of primary 'pluralism' in our scientific approach has also its consequences for our daily life. If science enables us to make more than one sound prediction, how may we use science as a means of action? We can never avoid a 'decision', because no account would be able to show us one action as 'the best', no computation would present us with any 'optimum', whatever actions have to be discussed. Therefore 'decision' plays its part in any kind of scientific research as well as in our daily life. That is the reason why I stressed the 'unpredictability' as an essential element of empiricism thus repudiating all attempts to use unequivocal historical predictions as the basis of social actions." (Neurath, 1946, p.80)

Neurath's picture of the engineer is more modest than that of Hayek. He rejects the idea that the engineer could collect all the knowledge of his particular limited world, to use Hayek's phrase. He wants to reject that picture of the engineer as someone who has complete knowledge of a domain.

He calls that view pseudorationalism. His criticism of pseudorationalism is that the pseudorationalist thinks that you can come up with a set of rules of rationality which determine for you an optimal answer to any decision, or rules of scientific method which uniquely determine which of a set of theories is true. A rationalist who believes in reason must recognise the boundaries to the power of reason in arriving at decisions: "Rationalism sees its chief triumph in the clear recognition of the limits of actual insight." (Neurath, 1913, p.8) Given this pluralism and this rejection of the possibility of complete knowledge, he thought there were no purely technical solutions to social choices. They involve uncertain knowledge and plural values that require the use of ethical and political decisions and judgement.

You therefore have the positivist here sounding very much more like the modern participatory critic of what is considered to be positivism. "Our life is connected more and more with experts, but on the other hand, we are less prepared to accept other people's judgements, when making decisions" (Neurath, 1945c, p.251). Democracy is "the continual struggle between the expert… and the common man" (Neurath, 1945c, p.251). On the one hand, citizens have to trust experts but, on the other hand, citizens have good reasons not to. Politics is partly about establishing the conditions of trust. At the same time he sets democracy as the continual struggle between the expert and the common man, and there is something right about that picture of democracy. That is why certain types of democratic experiment are important in bringing expertise to democratic accountability.

Genetic engineering, the limits of knowledge and precaution

Neurath's points are important to understanding responses to bioengineering. Consider again the debates around genetic engineering. With respect to the plurality of values that bear on the issues, there are important questions about the framing of the debate. There is a tendency for technical experts to explicitly import particular assumptions about what values are at stake in something like genetic engineering. If you examine the deliberative panels, such as citizens' juries that have been run on genetically modified organisms, expert framings tend to focus on risk to health and welfare. In contrast, members of deliberative panels often do not talk about these questions at all. They are more concerned about the political economy of GM technologies – for example, what happens if you increase the power of large commercial enterprises who can therefore control seeds (O'Neill et al 2008b). It seems to me that there is a perfectly important set of values and considerations to be brought to the debate, which fall outside the typical expert framing of the issue. The question of which values shape the framing of the debates matter.

Secondly, there are questions about the limits of conscious control and knowledge, and there is ignorance and uncertainty about what the consequences of introducing GM crops will be. It is a matter of introducing organisms into nature which are quasi-natural – that is, that they will have their own reproductive patterns and they will undergo a whole series of natural processes which are independent of our will. At this point, it is important to employ the more modest idea of what engineering involves.

This modesty is indeed implicit in the practice of good engineering itself. In contrast to the theoretical sciences, engineering is a practice that deals with the limits of human control. The art of getting things to work carries with it the awareness of the variety of ways that things can go wrong. Engineering, as a result, has a version of precaution and is presupposed by good design. Consider the following from the engineer Lev Zetlin:

> "Engineers should be slightly paranoiac during the design stage. They should consider and imagine that the impossible could happen. They should not be complacent and secure in the mere realisation if all the requirements of design handbooks and manuals have been satisfied, the structure is safe and sound." (Cited in Petroski 1994)

He continues:

> "I look at everything and try to imagine disaster. I am always scared. Imagination and fear are among the best engineering tools for preventing disaster." (Cited in Petroski 1994)

If that is right, and I think there is a kind of precaution that is built into good engineering.

This precautionary approach is particularly important in genetic engineering because it is particularly important in conditions of human vulnerabilities and human dependency. There is a presumptive precaution with certain environmental conditions which form the background of natural processes on which our lives depend. Human beings are biological organisms. We are natural organisms - needy beings dependent on natural processes. David Wiggins captures the particular importance of precautionary considerations in this context: "Less and less are we in a position to risk that which is (or will be) necessary to human life in the hope of acquiring the superfluous ie that of which we have no vital need" (Wiggins, 2006, 50).

The concerns again often emerge in deliberative panels in contrast to expert panels. Members of panels will ask why we need GMOs. The question – do we need it? – is an important one. It is what the aims and values are that matters and, once again, that is picked up in a number of those citizens' panels and informs worries about certain kinds of GMOs: you do not just want a redder, nicer tomato, if it threatens human needs. This is not to preclude debate.
There might perhaps be good answers to that question in certain circumstances. Vital needs might be at stake in certain conditions, for example where there exist severe risks to people's health and safety, where there is a certain kind of crop spraying and so on. However, the question of the value framing again matters.

Conclusions

At the start of this paper I noted that there was a discernible public scepticism about engineering in the environmental sphere. In this paper I have tried to articulate what is at stake when concerns are raised about unnaturalness and human hubris. My central aim here has not been to defend those lines of criticism, but to show that there are questions they raise that do need to be addressed. There are ways of stating those concerns which raise important issues about the contexts of human lives and the limits of human knowledge and control. However, at the same time I hope to have shown that part of the problem here is a particular image of engineering itself that permeates these debates and older debates on social engineering. The picture of the engineer as necessarily pursuing a single end in the possession of complete knowledge of a domain is false. Engineering is potentially a more pluralist enterprise than this image suggests. It is also an enterprise which deals with the limits of abilities to control and manipulate the world. Precautionary considerations are part of the practice of good engineering.

References

Euro-barometer (2000) *The Europeans and biotechnology* Brussels: European Commission Directorate-General for Research **http://ec.europa.eu/research/pdf/eurobarometer-en.pdf**

Euro-barometer (2005) *Europeans and Biotechnology in 2005: Patterns and Trends Brussels:* European Commission Directorate-General for Research **http://www.goldenrice.org/PDFs/Eurobarometer_2005.pdf**

Gamble, A. 2006 'Hayek on Knowledge, Economics and Society' in E. Feser ed. *The Cambridge Companion to Hayek* Cambridge: Cambridge University Press

Goodin R. 1992 *Green Political Theory*. Cambridge: Polity Press.

Hayek, F. A. (1942-4) 'Scientism and the Study of Society' *Economica*, 9, 1942, pp. 267-291; 10, 1943, pp. 34-63; 11, 1944, pp.27-29 reprinted with revisions in F. Hayek *The Counter-Revolution of Science* (Indianapolis: Liberty, 1979).

Hayek, F. A 1974 *Nobel Prize Lecture* **http://nobelprize.org/nobel_prizes/economics/laureates/1974/hayek-lecture.html**

Hume, D. (1978) [1739] *A Treatise of Human Nature* Oxford: Oxford University Press

Mill, J. S. 1874 'Nature', in *Three Essays on Religion* London: Longmans

Neurath, O., (1913) 'The Lost Wanderers of Descartes and the Auxiliary Motive' in Cohen R.S. and Neurath M. eds. (1983) *Otto Neurath: Philosophical Papers* Dordrecht: Reidel

Neurath, O. (1945) 'Visual Education' in E. Nemeth and F. Stadler eds. *Encyclopedia and Utopia: The Life and work of Otto Neurath* (1882-1945) Dordrecht: Kluwer, 1996

Neurath, O. (1946) 'After Six Years' *Synthese*, 5, 1946, pp.77-82

North, R. 1990. 'Why man can never kill nature' *Independent*, February 1

Nuffield Council 1999 *Genetically Modified Crops: The Ethical and Social Issues*
Nuffield Council: London **http://www.nuffieldbioethics.org/sites/default/files/GM%20crops%20-%20full%20report.pdf**

Nuffield Council 2004 *The Use of Genetically Modified Crops in Developing Countries* Nuffield Council: London **http://www.nuffieldbioethics.org/sites/default/files/GM%20Crops%20Discussion%20Paper%202004.pdf**

O'Neill, J. Holland, A. and Light A. 2008a *Environmental Values* London: Routledge

O'Neill J. C. Carter, K. Ekeli and W. Kenyon 2008b *Representing Diversity* PATH Policy Research Brief 1. EC Coordination Action, 'Participatory Approaches to Science and Technology' Aberdeen: Macaulay Institute available at: **http://www.macaulay.ac.uk/socioeconomics/research/path/PATH%20Policy%20Brief.pdf**

Petroski, H. 1994 *Design Paradigms: Case Histories of Error and Judgment in Engineering*, Cambridge: Cambridge University Press
Wiggins, D. 2006 'An Idea We Cannot Do Without' in Reader. S. ed. *The Philosophy of Need*. Cambridge: Cambridge University Press

2.3 Children of Martha: On being an engineer in the environment

Professor Roland Clift, CBE FREng

Roland Clift is Distinguished Professor of Environmental Technology and Founding Director of the Centre for Environmental Strategy (CES) at the University of Surrey. His research specialisation is in the broad field of Environmental System Analysis, including Life Cycle Assessment, Industrial Ecology and Sustainable Energy Systems. Professor Clift is also a Visiting Professor in Environmental System Analysis at Chalmers University, Gothenburg, Sweden and a Director of the Merrill Lynch New Energy Technologies investment trust. In 2003, he was awarded the Sir Frank Whittle medal of the Royal Academy of Engineering "in recognition of an outstanding and sustained engineering achievement contributing to the wellbeing of the nation". In 2005 he completed a nine-year term as a member of the Royal Commission on Environmental Pollution, and acted as Expert Adviser to an inquiry by the House of Lords Science and Technology Committee on Energy Efficiency. In 2006, he was appointed a member of the Science Advisory Council of the Department of the Environment, Food and Rural Affairs (Defra). Professor Clift was honoured in the 2006 New Year's Honours as a Commander of the British Empire (CBE) for services to the environment.

Introduction

In what follows I will be addressing questions of ethics, as an engineer. I am concerned with things like the Statement of Ethical Principles for Engineers[1]. When I look at those principles and think about how, for example, some of my students might use them in practice, I think that they will have a problem. The Academy has had debates on the subject of ethics over the last few years and - perhaps I am being harsh by saying this - concluded that 'ethics are a good thing and we need more of them'. There is actually more to it than that, and that is what I shall try to explore here. To apply something like a code of ethics to a real engineering problem, there are issues which have to be resolved clearly first. The particular issue that I shall address here is the difference between consequential and deontological ethical traditions. These really are different and the difference matters.

Along the line, I will be examining, not just what it means to apply philosophy or ethics in engineering, but what it means to be an engineer. To do this I shall be introducing some of the thinking of Martin Heidegger.

Custom and principle

I recommend the book *Environmental Ethics* by Joseph R. Des Jardins [1]. According to Des Jardins,

> "ethics refers to the general beliefs, attitudes or standards that guide customary behaviour."

That word 'customary' will turn out to be important; customary behaviour in the context of the engineering profession, for example. On page 3 of *Philosophy for Beginners*, Robinson and Graves [2] say:

> "…philosophers are obliged to provide some kind of explanation, proof or evidence for their ideas. And this obligation marks the one obvious difference between philosophy and religion".

I include this statement because I still encounter confusion on the matter. Some academics, with whom I occasionally debate, think that ethics is a branch of religious studies. This view is expressed by some very significant figures, particularly in the USA. I want to state clearly that ethics is a branch of philosophy, not a branch of religion.

If you do not agree with me on an intellectual level, then let us look at it on an entirely pragmatic level. When I am trying to teach ethics to a group of engineering graduate students, the class will typically contain eight religious traditions plus some agnostics and atheists. Therefore, if I tried to cover ethics from a religious point of view, we would get absolutely nowhere. As far as I am concerned, the problem about ethics lies in how to develop ethical principles which do not depend on religious belief or received wisdom. That is why I say that ethics is a branch of philosophy and not of religious studies. By the way, as far as I am concerned, 'deep ecology' is in the same category as religion.

Turning to Des Jardins once more, the book I recommended includes the remark: "Ethics as a branch of philosophy seeks a reasoned examination of what custom tells us about how we ought to live." The meaning of that sentence is that ethics is essentially normative. One of my conclusions is that engineering as a profession needs to be normative. That is a little uncomfortable, if you were brought up in engineering in the way that I was.

---

[1] *http://www.raeng.org.uk/societygov/engineeringethics/pdf/Statement_of_Ethical_Principles.pdf*

Ethical traditions

I shall be talking about two different ethical traditions, so I will begin by explaining what they are. Consequentialism states that actions are judged by whether their consequences are right. At first sight, you might think that this is the way that engineers go about their business. In a way it is, because a particular form of consequential ethics is utilitarianism, associated with people like Jeremy Bentham and John Stuart Mill. The idea behind utilitarianism is that you should maximise the overall good or produce the greatest good for the greatest number by your actions – and 'good' is interpreted pretty widely.

Utilitarianism is the basis of methods like cost-benefit analysis. Engineers are used to deploying cost-benefit analysis. However, when you search around in the intellectual undergrowth, as I do sometimes, you find that cost-benefit analysis hides some pretty nasty little beasties, one of which is the Kaldor-Hicks compensation test. I will not go into that here, but it is something which is built into the economic concept of cost-benefit analysis and, when you realise it and dig it up, it smells. As far as I am concerned, there are actually serious ethical problems in applying cost-benefit analysis as an example of a utilitarian ethical approach.

So what is the alternative? The alternative is the deontological approach, which means acting on principles rather than being guided by consequences. The idea develops mostly from the work of Immanuel Kant, in his idea of the categorical imperative. He stated that, to be moral, you should act in such a way that it would be acceptable for everyone to act in accordance with the same rule that guides your action. These are really quite distinct ethical approaches.

Albert Flores [3, p.207] makes the distinction clear:

> "Whereas according to the deontologist an action is right or wrong insofar as it conforms with our moral duty, according to (the) consequentialist approach the rightness or wrongness of an action is a function of the goodness or badness of its consequences".

Why is the difference important? On the few occasions when I have tried to teach ethics, usually to engineering students, I have done it in terms of case studies which I use rather in the manner of *koans*, from the Zen tradition. I give them an ethical problem, not with the objective of solving the problem but of analysing how to approach it.

Let me give you a couple of examples. One of them is a case study which is developed in a book which I had a hand in editing, called *Sustainable Development in Practice* [4]. The case study centred on a decision on whether or not to exploit a rich uranium ore deposit in Northern Australia, which is on land held sacred by a small group of aborigines. You reach entirely different conclusions if you approach that problem from a consequentialist or a deontological starting point.

Another real problem, geographically closer to home, was Brent Spar. A way of understanding what happened over Brent Spar was that Shell took a consequentialist point of view and looked at the cost-benefit analysis of ways of disposing of an old oil platform. The NGOs, and Greenpeace in particular, managed to persuade the public that this should be treated as a deontological problem: what is the *right* way, in principle, of disposing of oil platforms?

I hope these concrete examples illustrate why I am so taken up by this issue of consequentialist versus deontological ethical systems. It matters, because you reach completely different conclusions according to which system you apply. Simply having a two-page code of ethics in front of you does not help you with that, unless you know which approach to take.

The existential problem of engineering

In order to try to draw some conclusions about whether we should be consequentialist or deontological, I want to consider something completely different. This is where I will go into the existential problem of engineering: what it means to be an engineer.

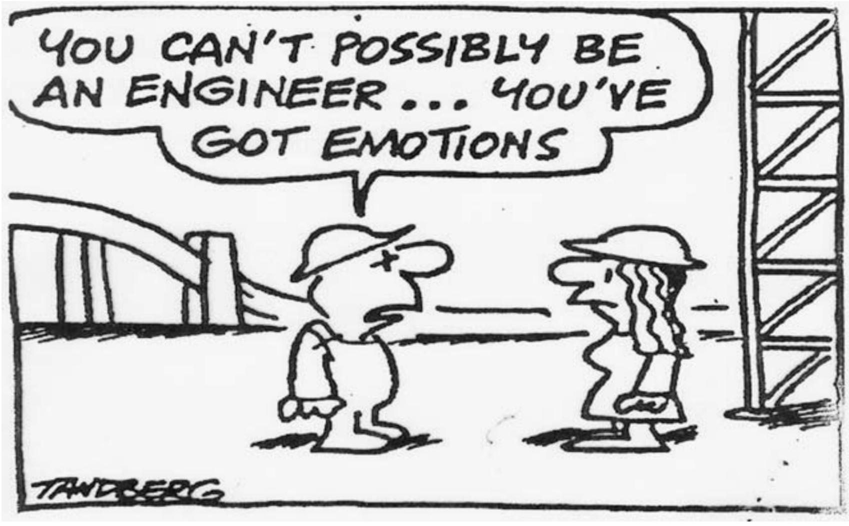

Figure 10. Cartoon showing the stereotypical view of what it is to be an engineer

I got this cartoon from an Australian student magazine (Figure 10), which introduces the problem. I will maintain that the idea that engineers have to be technological automata is fundamentally wrong and that actually having emotions is an advantage to practising as an engineer.

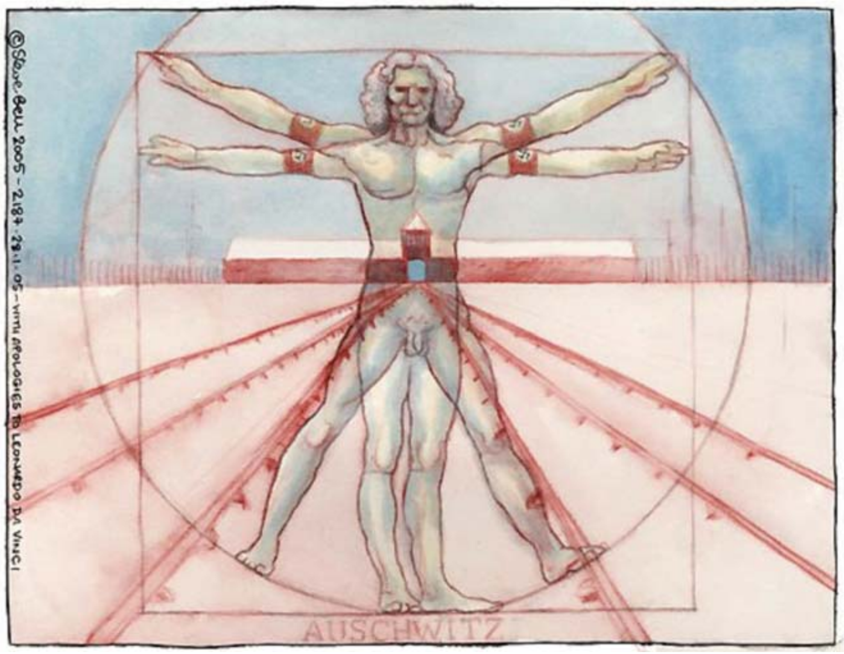

Figure 11. Cartoon by Steve Bell produced on the 60th anniversary of the liberation of Auschwitz

Here is an example of why. This is a cartoon by Steve Bell which appeared in *The Guardian* on the 60th anniversary of the liberation of the concentration camp at Auschwitz (Figure 11) – it actually shows the extermination camp at Buchenwald. I imagine - in fact I sincerely hope - that most people here would say that it would have been wrong to work on developing that particular piece of infrastructure – the 'camps' at Auschwitz and Buchenwald – but it was a great feat of engineering. The logistics worked very well and the processing was very rapid. The use of chemical agents was very effective. Purely as an engineering project, it was a big success. If you want an example of why engineering ethics are essential, surely this is one. I do not think it helps with the distinction between consequential and deontological ethics because, however you look at it, you should not do it. Nevertheless, it is an example of why engineers cannot act purely as technical automata.

Engineers as the children of Martha

I did my graduate studies in Canada and, because I graduated with my doctoral degree from McGill University in Montreal, I was entitled to go to the ceremony called the Calling of the Engineer. In Canada, when engineers graduate and are allowed to practise the profession, they are given an iron ring to wear on their little finger. Originally, the iron rings were made from the steel of a bridge over the St Lawrence River, which collapsed – so it served as a reminder of the responsibility of the engineer.

Part of the theme of the ceremony of the Calling of the Engineer is that engineers are the children of Martha. This refers to a passage in *Luke*, Chapter 10. Jesus goes into a village and goes into the house of a woman named Martha, who had a sister called Mary. Mary sits at Jesus's feet, listening to his wise words, while Martha gets lunch, 'cumbered about much serving'. Mary, who should be helping her to get the lunch, is sitting, talking to Jesus. So Martha says, in effect, 'Why don't you tell Mary to come and help me get the lunch', but Jesus says, 'You're far too bothered about these things', and he goes on to say that Mary is right not to get the lunch but to listen to him and debate with him.

Kipling introduced this into the ceremony of the Calling of the Engineer, to say that engineers are the children of Martha – we are the people who get the lunch, not the ones who engage in lofty debate. I think that is wrong, actually: if engineers see themselves solely as children of Martha, then we could easily end up designing the transport system for Auschwitz or gas chambers for Buchenwald. We have to move away from Kipling's idea and not see ourselves as children of Martha, but as necessarily involved also in the lofty debate that engaged Mary.

And so to Martin Heidegger

This leads me to that difficult question of what it means to be an engineer, which is where I invoke Heidegger. Heidegger was a German philosopher who was born in 1889 and died in 1976. He was a professor at, amongst other places, Freiburg and he is regarded as one of the key figures in the development of existentialism, although he maintained that he himself was not an existentialist.

Heidegger's best known book is *Being and Time*. One of the questions that Heidegger tries to address is, what is *being*? Not 'what does it mean to be', but what *is* this condition, being, which enables things to come into existence? A key word in Heidegger's writing is *Dasein*. It means 'existence', as in the existence of everyday people, but it also means - if you deconstruct it slightly differently - *da sein*, being there. For Heidegger, *Dasein* means a human being but it also means human existence. We only become Dasein when we are socialised into a particular group of people with common attributes. He is really saying that we only exist in social contexts and it does not make sense to try to separate existence from that context.

Heidegger uses the example of a hammer as what he calls a 'ready to hand' – something you can use automatically. We use a hammer to bang in a nail: we do not really think about it; we just do it. This is an example of *Dasein*: we are thrown into a world where one uses a hammer to hammer things. If you had never seen a hammer before, perhaps you would do it differently, but it is idle to pretend that we have to think very hard about how to use a hammer – we just do it.

Furthermore, Heidegger says:

> "Even when we flee from the crowd, we do so in the way one does in the society we live in."

That resonates with me. Even when I am being critical of the engineering profession, I do so in the way that an engineer criticises things – and actually I cannot escape from that. In that sense at least, Heidegger has got it right: I am an engineer, even when I am mixing with philosophers and social scientists, but that does not necessarily mean that I am just a child of Martha.

Many people do not understand the difference between engineers and scientists. They do not understand that scientists basically study things, while engineers *do* things. Our natural role, as engineers, is *action*, not analysis. I do not want to hide that, or move away from it: it is what engineers do in our culture and it is the nature of being an engineer. There are quite a number of good books on the social problems of science. One that I particularly like is *Scientific Knowledge and its Social Problems* by Jerry Ravetz, published in the early 1970s. However, no one has yet written about engineering and its social problems and I think this is a literature that needs to be developed.

The existential pleasures of engineering

There are, however, some books recognising that there is this difference between engineering and science. In *The Existential Pleasures of Engineering*, Samuel Florman produces the well-known definition of engineering "…engineering is 'the art … of making practical application of the knowledge of pure sciences'". [5]

He asks what it is like to be an engineer at the moment when the profession has achieved unprecedented success but is coming under increasing criticism. He wrote the book in 1976 and that question has become even more critical since. What we are possibly looking at now is not just the ruin of our civilisation but, conceivably, even the collapse of our existence on the planet. It is sometimes held – wrongly – that engineering has had some role in that. I do not think that is a fair criticism at all but, nevertheless, there is a sense of anxiety, a palpable uncertainty about being an engineer. We have to face up to the fact that there is something of an existential crisis in engineering.

Florman goes on, and I agree with him again:

> "Engineers are not mechanics and nor are they technicians. They are members of a profession."

He, like me, was interested in what it means to be an engineer.

Heidegger and authenticity

What does it mean to be an engineer? Let us go back to Heidegger, as discussed by Hubert Dreyfus, in the BBC documentary series *The Great Philosophers* with Brian MacGee. He says: "… there is no reason one has to do things the way one does. God has not ordered us to do things this way, nor does human nature require it". The way we do things – the way we use a hammer, for example, without thinking about it – are not based on particular reasons but on habit. Heidegger explores the nature of *angst*, which you could perhaps translate as "uncomfortable unsettledness". This is shown by people when they realise that they are acting in a way which they had not adequately considered. This resonates with Socrates: "the unconsidered life is not worth living."

One way of fleeing that unsettledness is simply to conform; to display the kind of behaviour which Heidegger terms 'inauthenticity' – at least that is the usual English translation. This behaviour basically disowns what it means to be existing in this society. The alternative is to try to engage in 'authentic' activity. As Dreyfus put it, you can own up to what it means to be *Dasein* and, rather than trying to flee the anxiety, actually embrace it. The existential idea is that accepting your groundlessness is what is liberating.

This means that you go on, quite probably, doing the same thing but how and why you do it change radically. You accept your role and condition – I accept my role as an engineer, for example – but you constantly test, question and evaluate what you are doing. You do not just treat it as you would a hammer, as a ready-to-hand. The tools of engineering are not to be treated as ready-to-hand but they are to be thought about.

I want now to explore this idea of 'authenticity'. Dreyfus describes authenticity as "doing the sort of thing that one does in a way that allows a response to the unique situation without concern for respectability and conformity." It is doing the sort of thing that one does in one's culture, but in a way which is considered and therefore authentic, which means continually examining, reflecting and analysing.

Let us go back to Florman, who was very much on this agenda. He talked about the existential pleasures of engineering but in my view he made two mistakes. He said:

> "The first and most obvious existential gratification felt by the engineer stems from his desire to change the world he sees before him."

I do not think that is right, actually. It may have been right at the time but not any longer and I will shortly explain why. He then makes this remark, which of course I will not agree with:

> "Civil engineering is the main trunk from which all branches of the profession have sprung."

Civil engineering may have been the first non-military engineering, but I am a chemical engineer and so I come from a different intellectual tradition, as do electrical engineers. The branches of engineering which have their origins in new science do not stem from the same trunk as civil engineering and are therefore different.

Florman sees the engineer as someone who is obsessed with changing things, but how does that sit with the present situation in which we find ourselves, where often the role of the engineer is to try to save things, including trying to save bits of the natural world from being changed? It is still action, but it is a different kind of action. This is where, as a chemical engineer, I have a different perspective because I deal with flows of materials and energy, whereas a civil engineer is usually concerned with building things.

The new model engineer

At a lecture I gave at the Royal Society in 1997 [6], I talked about the way in which the role of the engineer has changed. If you go back to the 19th century with Brunel as an example, heroic materialism was the theme. The Mark II engineer, which I identify with the earlier years of this century, saw himself not as an heroic materialist but rather as someone who met human needs. I actually used my wife's father as an example there, Dermot Manning, a founding fellow of The Royal Academy of Engineering, because he developed polyethylene. When he was developing polyethylene, what he really envisaged was developing polytunnels to aid agriculture in arid regions, rather than consumers using plastic bags.

We have now moved on further, to the Mark III engineer. I called the kind of engineer who is necessary now, rather vaguely, the 'steward of the global commons'. In the Hartley lecture, I said that this necessarily includes:

> "representing the implications of technological choices in public deliberation, as an impartial expert rather than as an advocate of his or her preferred option" [6].

What I am saying here is that engineers are not just problem solvers; they have to become engaged in framing the problems which are to be solved. This is something to which the profession perhaps is not terribly accustomed, and I have summed this role up as being that of the 'honest broker'.

Here is another quotation from something I wrote with two Australian friends, Cynthia Mitchell and Anna Carew – it is a chapter in the book, *Sustainable Development in Practice*, where we tried to explore what it meant to be an honest broker:

> "The new approach we envisage asks the expert to take a longer-term, broader-scale, systems perspective on environmental impacts; to engage a range of stakeholders in consultative decision processes; and to practise personal quality assurance through self-critical and reflective practice."
>
> "… The Honest Broker paradigm expands the bounds of… engineering… in moving beyond problem solution towards problem formulation."
>
> "Problem formulation asks that the specialist consider the problem in terms of its broader context – environmental, social and economic."[7]

At the time, I did not conceive of the 'New Model' engineer in the context of authenticity but, coming back to it for the purposes of this lecture, it seemed to me that I was actually saying things about what it meant to act authentically, in Heidegger's sense, as an engineer. In particular, "practising personal quality assurance, through self-critical and reflective practice" to me is authentic, rather than inauthentic, engineering practice – always questioning, always examining. For example, it is not just accepting a design brief, but recognising that there are several dimensions to the problem, not just economic.

Deontological or consequentialist?

To sum that up, what I think I am saying is that the sustainability agenda does not change the tools the engineer has to use. A hammer is still a hammer and it is still 'ready to hand'. However, what it means to be authentic as an engineer changes, so the way you approach the problem of 'doing engineering' has to be different. It has to be much more self-aware and much more reflexive – in other words, much more authentic.

We have come rather a long way from the idea that engineers are children of Martha. We combine features of Martha and Mary. Now I need to get back to the problem which I posed at the outset. Are we going to use deontological or consequentialist, utilitarian ethical approaches?

If we are to take on this authentic role of engaging in framing problems, then we have to recognise that the outcomes of our actions are uncertain. But if you are looking at guiding your actions by outcomes which are very uncertain, then you have to abandon the false sense of security that comes from things like cost-benefit analysis. Andy Stirling gave a concrete example in his paper, where predictions of outcomes differed by several orders of magnitude. In such circumstances you cannot put any weight or precision on predicting the outcomes of your action. By the way, that is not inconsistent with doing a lifecycle assessment, lifecycle assessment being a key part of my toolbox, because I always argue that you should present a lifecycle assessment as an inclusive approach to decision-making in the recognition of uncertainty.

I think we have to recognise that uncertainty is inevitable. This actually reads across to the area which is sometimes called 'post normal science'. I won't go into this now, but it is reassuring to me to find internal consistency in what I am saying and try to practice. If we have to abandon the pseudo-certainty of utilitarianism, that says to me that our ethical code as engineers needs to be deontological. We need to be guided by principles of conduct, not by expectations of outcome. Returning to the two specific examples which I introduced earlier, we have to be prepared to say things like, 'actually Greenpeace were right over Brent Spar' and, 'we shouldn't extract uranium ore in Jabiluka'.

At the outset I raised the question, 'which ethical system should engineers use?' and pointed out that unless that question is answered, a simple code of ethics is of little use. Invoking the very obscure thinking of a German philosopher (who was a very unpleasant character, an unfortunate fact which I will not hide), I conclude that when we are implementing a code of ethics, we need to see it as a set of principles to be implemented in a deontological framework.

Coda

I want to make two further points before I finish. It might seem a bit odd to cite Heidegger in this talk because Heidegger was very critical of technology. However, I do not think he meant technology in the sense of the set of skills which is deployed by engineers. Rather, I think he was criticising something else – the economistic approach rather than engineering. He was critical of the approach which equates 'the world out there' with stuff which has economic value – the purely instrumental conception of the world. His word was *Bestand* – we turn things into *Bestand*, which in English is usually rendered as 'standing reserve' or 'stock' although I personally think 'capital' is a better translation. When you read some economists who talk about natural capital and invoke the nonsensical concept of 'weak sustainability' and the idea that you can trade natural capital off against manmade capital, they are making the mistake that Heidegger was ranting against. His target was a way of thinking, rather than technology in the sense of the skills which engineers use. He calls his approach *Gestell*, which is sometimes translated as 'framing', although it is not framing in the sense in which I and others have used it in these talks.

So I am invoking a German philosopher who is sometimes interpreted as being critical of technology, although my interpretation is that he was criticising the economistic view of the world. I do not think there is a problem there. However, there is another very real problem with Heidegger, and I will not try to hide from it. Heidegger was an active Nazi supporter and a raving anti-semite and, to his death, he never apologised or recanted. I do not know what he thought of the engineers who built Auschwitz and Buchenwald.

References

1. Des Jardins, Joseph R (1993) *Environmental Ethics: an Introduction to Environmental Philosophy*, Wadsworh
2. Robinson, David and Groves, Judy (1998) *Philosophy for Beginners*. Cambridge: Icon Books.
3. Flores, Albert (1982): *The Philosophical Basis of Engineering Codes of Ethics*. Reproduced in Veselind, P.A. and Gunn, A.S. (1998) Engineering, Ethics and the Environment, Cambridge University Press.
4. Azapagic, Adisa; Clift, Roland and Perdan, Slobodan (2004) *Sustainable Development in Practice: Case Studies for Engineers and Scientists*, Wiley-Blackwell
5. Florman, Samuel (1995) *The Existential Pleasures of Engineering* (2nd. Ed.), Souvenir.
6. Clift, R. (1998) Engineering for the Environment: The New Model Engineer and Her Role. *Trans.IChemE*, **76 B**, pp.151-160.
7. Mitchell, C.A.,Carew, A.L. and Clift, R. (2004) 'The Role of Engineers and Scientists in Sustainable Development', Chapter 2 in [4]

The Royal Academy of Engineering

As the UK's national academy for engineering, we bring together the most successful and talented engineers from across the engineering sectors for a shared purpose: to advance and promote excellence in engineering. We provide analysis and policy support to promote the UK's role as a great place from which to do business. We take a lead on engineering education and we invest in the UK's world class research base to underpin innovation. We work to improve public awareness and understanding of engineering. We are a national academy with a global outlook and use our international partnerships to ensure that the UK benefits from international networks, expertise and investment.

The Academy's work programmes are driven by four strategic challenges, each of which provides a key contribution to a strong and vibrant engineering sector and to the health and wealth of society.

Drive faster and more balanced economic growth

The strategic challenge is to improve the capacity of UK entrepreneurs and enterprises to create innovative products and services, increase wealth and employment and rebalance the economy in favour of productive industry.

Lead the profession

The strategic challenge is to harness the collective expertise, energy and capacity of the engineering profession to enhance the UK's economic and social development .

Foster better education and skills

The strategic challenge is to create a system of engineering education and training that satisfies the aspirations of young people while delivering the high calibre engineers and technicians that businesses need.

Promote engineering at the heart of society

The strategic challenge is to improve public understanding of engineering, increase awareness of how engineering impacts on lives and increase public recognition for our most talented engineers.

The Royal Academy of Engineering promotes excellence in the science, art and practice of engineering.

Registered charity number 293074

The Royal Academy of Engineering
3 Carlton House Terrace, London SW1Y 5DG

Tel: 020 7766 0600 Fax: 020 7930 1549
www.raeng.org.uk