



Vodafone Institute
for Society and
Communications



ENTERING A NEW ERA

A collection of essays and interviews
by Alexander Görlach with:

Martin Rees	Huw Price
Pascal Finette	Andrew McAfee
Luciano Floridi	Vinton G. Cerf
Nuria Oliver	Yuval Noah Harari

**THE IMPACT OF
ARTIFICIAL INTELLIGENCE
ON POLITICS, THE ECONOMY
AND SOCIETY**



Vodafone Institute
for Society and
Communications

ENTERING A NEW ERA

**THE IMPACT OF
ARTIFICIAL INTELLIGENCE
ON POLITICS, THE ECONOMY
AND SOCIETY**

Table of Contents

3
Preface
Hannes Ametsreiter,
CEO Vodafone Germany

6
The Age of Painters and Poets
Alexander Görlach,
Harvard University,
Carnegie Council for Ethics
in International Affairs

The Interviews

12
A key question is whether robots can handle ambiguity
Martin Rees, Cosmology and
Astrophysics at the University
of Cambridge

20
We might want to rediscover our social roots
Pascal Finette, Entrepreneur-
ship at Singularity University

30
We will seamlessly cooperate with machines
Nuria Oliver, Fellow of the
European Association of
Artificial Intelligence

34
Some machines are people in the philosophical sense
Huw Price, Philosophy at the
University of Cambridge

44
Machines can do some things better than we can
Vinton G. Cerf, internet pioneer

48
One answer might be computer games
Yuval Noah Harari, Historian
at the Hebrew University of
Jerusalem

52
Epilogue
Alexander Görlach,
Carnegie Council for Ethics in
International Affairs

54
What's next?
Alice Deißner, Vodafone
Institute

55
Impressum

Preface

“From man to machine; from machine to man”



DR HANNES AMETSREITER
is CEO of Vodafone Germany and member of the global Vodafone Group's Executive Committee. Born in Salzburg, the reputed telecommunications manager has more than 20 years of telecoms experience in convergent markets and outstanding expertise in marketing and brand management. He played an instrumental role in positioning Vodafone as a gigabit company and restoring growth by introducing increasingly fast speeds, convergent solutions and attractive products that integrate mobile, fixed, internet and TV.

Will artificial intelligence soon be superior to human intelligence? Perhaps it is already. Just a few decades ago it was beyond all imagination that a machine would ever be able to store more specialised knowledge than a human brain. Watson, Deep Blue and AlphaGo prove that the opposite is true. In fact, the recently developed Alpha Zero programme is now even capable of teaching itself complex board games. It does not need a human coach.

Will machines soon be ruling the world? Have we *Homo sapiens* just created the *Homo deus* (as Yuval Noah Harari calls it)? Will the human race be replaced by machines? There are many science fiction blockbusters with these kinds of plots, and some of the people who read them are clearly concerned.

However, we should not forget that not even the most imaginative movie directors have ever come even close to predicting the future. So there is still hope for us. Above all, we should not forget the immense benefits for mankind associated with the use of artificial intelligence. It is thanks to artificial intelligence that we have self-learning programs which can help shorten our recovery time when we get ill, high-efficiency energy systems, optimised traffic management systems, household robots and robots that can perform surgery with greater precision than any surgeon. It is no wonder that the major technology enterprises in Silicon Valley long ago identified artificial intelligence as a key future technology.

AI is an integral part of the gigabit society. And artificial intelligence thrives on data. To transport that data quickly and reliably, we need high performance networks such as those operated by telecoms enterprises like Vodafone. That is why it is imperative that we consider the impacts of artificial intelligence.

To create this publication the Vodafone Institute – Vodafone's think tank – interviewed some of the most brilliant and influential modern day minds to discover their thoughts and ideas on artificial intelligence. Their answers, some from very different perspectives, provide valuable input for the future artificial intelligence debate.

We can only imagine the possibilities that artificial intelligence will open up to us today. What we do know, though, is that it will change all our lives for ever. That is why it is all the more important that we learn how to use it responsibly in business, in politics and in society – for the benefit of mankind.

Enjoy reading!



Henry Ford streamlined automobile production in 1913 by introducing the assembly line. Some jobs were lost as a consequence



In modern automobile factories, it is the robots that now perform the work of humans

The Age of Painters and Poets

An Essay by Alexander Görlach

What will we do when machines take over our jobs? It is certainly a question of how to spend our time but also of how to gain and maintain self-esteem and identity. We may not be aware of the fact that our societies since the dawn of time run on narratives of work. That is why we need to dive into a new history of purpose for mankind in the age of machine learning and artificial intelligence.

No doubt, the Fourth Industrial Revolution will shatter our certainties and customs, the things we take for granted, the same way its three predecessors did. What has happened over a quarter-century as globalisation teamed up some fifteen years ago with digitalisation: speaking today about the ubiquity of goods, finance, people and information is commonplace and is embodied in the smartphone that allows us to basically consume everything one might need for everyday life. These smartphones with their services such as those provided by Google and Facebook not only allow the collection and storage of a multitude of information, unimaginable for our ancestors, but also allow us to work with them and use them for an exponential acceleration of creating businesses and new technologies.

Data mining leads to shopping recommendations on Amazon and to echo chambers and opinion silos on

Facebook. Large amounts of data, such as those available to Google, open up new fields of research in fields such as the health sector. New players will triumph and old ones will go out of business. What has not been highlighted enough so far, however, is the consequence the rise of machine learning and other sorts of artificial intelligence may have on societies. This holds especially true when it comes to the future of work in a world where machines may do a great chunk of it, be it manually or intellectually. Work that we humans have been used to doing since the dawn of time.

Western societies are very much shaped by the idea of a work ethic, the Protestant ideals of being busy in this world. There is no ethic of leisure and no place for pleasure within the Western narrative of success and hard work. This ideal, some would say ideology, has reached out far beyond the frontiers of Christendom and is definitely not just a Western concept. Some may detect traces of these ideals appearing independently in contemporary China, which is anything but Christian. Yet the Confucian narrative that shapes this region of the earth is built on meritocracy and the ideal path of man is one of hardship and lifelong learning. It is not all that different from what we know in the West.

The narrative of progress and success, the American “from rags to riches”, has entrenched itself in a very secular sense in most people’s minds in the contemporary world. Therefore, if we really seek to understand the battles we may be fighting soon about the future of our societies, the place of man in the world of tomorrow, we may have to look very closely at this narrative.

Self-worth is not only embodied in a monthly pay-cheque but also, maybe most importantly, in the passed-down set of beliefs about the essence of our human endeavours as an *animal laborans*. There is no doubt that these beliefs have been questioned, attacked and demeaned over the course of the last two hundred years. After the so-called three humiliations of modernity, man had to endure the theories of Copernicus, Darwin and Freud. The Fourth Industrial Revolution may also have the next humiliation in store for him. Copernicus exiled man from the centre of the universe. Darwin deprived him of being the pride of creation. And Freud finally taught him that he



ALEXANDER GÖRLACH is an affiliate professor at Harvard University, in the “In Defense of Democracy” program of the F.D.Roosevelt Foundation at the College of the University. He is further a senior fellow to the Carnegie Council for Ethics in International Affairs and a fellow to the Center for Research in the Arts, Social Sciences and Humanities (CRASSH) at the University of Cambridge, UK, and a senior advisor to the Berggruen Institute. Alex focuses in his work on narratives of identity, one of them being the narrative of work in the age of AI and machine learning. He holds PhDs in linguistics and comparative religion. Alex is the founder of the debate-magazine The European that he also ran as editor in chief. Today he is the publisher of the online initiative www.savoliberaldemocracy.com and an op-ed contributor to the New York Times and the Neue Zürcher Zeitung.

does not even have dominion in the house of his own mind. With the dawn of artificial intelligence, the belief of man that he participates in the divine essence of God through his intelligence and creativity (a word clearly alluding to the Creator), may be shattered.

What to do? As the usual business is to look into business opportunities and legal regulation when it comes to technological innovation, this essay tries to make the case for a calm look into the narratives that define and confine us: our identity as working human beings, our ability to cope and deal with change or longing for persistency.

As a first step into this endeavour, it shall be rewarding to distinguish between an US-American perspective on the subject matter and a continental Western-European one. As the industrial revolutions of the past originated in this part of the world, the fourth and latest version of this innovation is also deeply connected with this cultural hemisphere.

In his book “America” French philosopher Jean Baudrillard sums up what, according to him, is the difference between the United States of America and Europe. The Old World, he argues, is bound to history. The New World, on the contrary, bows to Utopia. Both cultural spheres therefore follow different paths, form different patterns while approaching and understanding reality. If you live in a constant future, progress is always the present time. If history is your reference point, progress is not your benchmark but rather persistence or conclusiveness become the ideals that embody the values by which society runs. There is no other field where this difference in attitude and Weltanschauung, history versus utopia, has been more out in the open than in the different attitude towards development and technological progress. In Baudrillard’s words:

“We criticise Americans for not being able to either analyse or conceptualise. But this is a wrong-headed critique. It is we who imagine that everything culminates in transcendence, and that nothing exists which has not been conceptualised. Not only do they care little for such a view, but their perspective is the very opposite: it is not conceptualising reality, but realising concepts and

“The Terminator, as the epitome of what the world would become if machines have the capacity to run the place. What is seen as a fictitious movie overseas is taken at face value in Europe.”

materialising ideas, that interests them. (...) Everything that has been dreamt on this side of the Atlantic has a chance of being realised on the other. They build the real out of ideas. We transform the real into ideas or into ideology. Here in America only what is produced or manifested has meaning; for us in Europe only what can be thought or concealed has meaning.”

That is the division of labour between Europe and the United States even today (Baudrillard’s book was published in 1986), when it comes to the next wave of technological innovation: in Europe they worry about the outcomes of machine learning and artificial intelligence, in America they find and finance the companies that invent the algorithms and build the robots. In the US they are creating progress that Europeans would be sceptical about, out of principle and a weakness for scepticism. As a matter of fact, most newspaper and magazine articles about artificial intelligence prove that by carrying the face of Arnold Schwarzenegger on their cover: The Terminator, as the epitome of what the world would become if machines have the capacity to run the place. What is seen as a fictitious movie overseas is taken at face value in Europe.

There is no doubt that there are two different velocities nowadays within the development of artificial intelligence and its agencies: there are the big companies that emerged from Silicon Valley’s long-standing effort to become the world leader in technology: Google, Facebook, all American companies that gained a monopoly in the field of data mining. Also, Amazon in Seattle plays a decisive role. Or Boston is a hotbed as well, where a lot is going on in the field of robotics. So, the US has taken off while Europe is bewildered.

Once in a conversation with Andreas von Bechtolsheim, founder of Sun Microsystems and one of the first investors in Google, he emphasised the fact that by the time he was a student, in the late seventies of the last century, Germany had already been left behind by the developments and successes in computer sciences in the United States. Today a billionaire, he made clear in our conversation that the daring and entrepreneurial spirit in the New World was unparalleled on the other side of the Atlantic. He told the story of when he first met the founders of Google, back then unknowns, on a patio for a chat and shortly after that became Google’s first investor, putting a hundred thousand dollars into their new venture. In return he was given one per cent of the new company, today worth a fortune.

But that is only half the story: there is not just a difference in attitude, aptitude and readiness to assume risk when it comes to computer sciences, machine learning and artificial intelligence. There are also different mind-sets and attitudes in play, regardless of which side of the Atlantic you live on. There seems to be a gulf of the same size between the sciences and the humanities. To understand and adapt to the seriousness of the changes our society is going through, it is crucial to overcome this divide. We cannot evaluate our cultural narratives of work and identity without cross-disciplinary thinking.

Looking into this, one will find out that the divide goes way back, to the 19th century in England and even before, as C.P. Snow sets out in his lecture “The Two Cultures”. In this lecture, delivered in Oxford in 1959, the British scientist and novelist criticises the silos in academia that have been established in the institutions of his own country, Oxford and Cambridge, with the rise

of sciences forever connected to the name Newton. In the age of industrialisation, the snobbishness of the liberal arts representatives has taken on new forms, that Snow labels quite frankly as ignorance. He states:

“The non-scientists have a rooted impression that the scientists are shallowly optimistic, unaware of man's condition. On the other hand, the scientists believe that the literary intellectuals are totally lacking foresight, peculiarly unconcerned with their brother men, in a deep sense anti-intellectual, anxious to restrict both art and thought to the existential moment.”

Europe, being focused on history, is according to Baudrillard haunted by its need for concepts and metaphysics. The liberal arts, the epitome of the European ideal of education and formation, are also striving for the existential and the defining. Because of that, states Snow, the liberal arts lack the capacity to comprehend what is happening outside College doors and beyond their secure walls.

To fully understand the changes at hand, we do need sciences and humanities to work together. Social sciences seem to be a link between the two as it brings together quantitative research with qualitative interpretation. Rightfully so, because it is all about hermeneutics, a generic discipline of philosophy: a theory of understanding and interpreting. Usually this theory is only to be applied to texts. It is also a method of interpreting texts that speak to us from a different era or another cultural framework. In a time where algorithms define our reality, it is clearly necessary to develop a new hermeneutic translation model from the language of computer science into the language of the humanities and vice versa. As artificial intelligence mimics human behaviour, with all its biases, it renders it indispensable that the two not only understand each other but create something like a new framework of interpretation and application. This is necessary above all for one simple reason.

Societies run on narratives, on parables, stories, that serve as presuppositions, as convictions that underlie most of a society's discourse. Most of these narratives have to be detected in careful and diligent research by the social sciences, for those surveyed may not even

“As artificial intelligence mimics human behaviour, with all its biases, it renders it indispensable that the two not only understand each other but create something like a new framework of interpretation and application.”

consciously be aware of sharing certain narratives with other members of their group. The American Dream is one of those narratives, embodied in the phrase “from rags to riches”. Also, the narrative of a Christian Europe and the Occident. Narratives define groups and help their members to cope with the contingency of our human lives. Believing, behaving and belonging are the driving forces of all human collectives, be it a tribe or an industrialised society.

One of the most prevalent narratives we share in the Western world is the one of work. And this narrative will be the one that will be challenged, questioned and finally destroyed by the achievements of machine learning and artificial intelligence. As early as in the story about the garden of Eden, we learn that man works. Even before man was sentenced to work as punishment for his sin (physical labour experienced as a hardship and burden) he still worked in paradise, but for pleasure. Theologians depict here the essence of work: it is not the “by the sweat of your brow you shall eat bread” (Genesis 3:19), the money making of the capitalist age. It is, to the

contrary, work as self-expression, as a part of a person's identity. Moreover, it defines mankind's pride, the possibility of creating and by doing so participating. We already had this in the glorious, divine maker's process of creating.

In the Protestant branch of Christianity, labour and the fixation on work has taken on the leading role in ethics. A successful person is heralded as one favoured by God. Calvin taught this. The Puritans believed it – and brought this belief in the belly of the Mayflower to the shores of the New World. In today's Cambridge, at Harvard University, a place founded by the Puritans where I had the pleasure of being a visiting scholar and fellow, you could still see how much this Protestant work ethic, as Max Weber famously called it, prevails in a completely secular way.

It is an environment in which every conversation starts with saying how busy one is. A meeting for a coffee consists of walking together to the coffee shop, waiting together in line and then dashing back to your desks. By doing this you make sure of not missing out on any tiny bit of work, and finally, when the day comes to an end, you can be satisfied that you did not waste even one minute in an inefficient, ungodly way. Harvard is an environment where the carbonic acid in the water is perceived as utterly hedonistic. It stunned me, a lousy son of the Una Sancta, a proud European moreover, how much one can be consumed by the fetish of work and busyness. It puzzles me how US-Americans live with six vacation days a year, with nine the maximum, yet are considered lazy prats if they really manage to take all of them.

Now you get what I am trying to say here. Imagine a place like Harvard (or the United States as a whole, for that matter), deprived of the solace of believing, behaving and belonging according to the narrative of the Protestant/Puritan work ethic. It would fall to pieces! Yet we already know that machine learning and artificial intelligence will not only impact the work force in factories. Automation will hit law firms, banks, hospitals and universities alike. It will cost jobs in all sectors. Banks like Goldman Sachs invest highly in new software that will make the work of hundreds of Harvard graduates in their laying batteries dispensable. When this happens,

what will give us self-worth and pride? What will have us believing, behaving and belonging? And: what the heck are we going to do with all our free time?

Clearly, we will not all become painters or poets. But what will we be doing with that abundance of time? Yuval Noah Harari, a great storyteller and historian (author of the best-selling books “Sapiens” and “Homo Deus”) recently laid out an answer to that question in an opinion piece for the British newspaper The Guardian:

“One answer might be computer games. Economically redundant people might spend increasing amounts of time within 3D virtual reality worlds, which would provide them with far more excitement and emotional engagement than the ‘real world’ outside. This, in fact, is a very old solution. For thousands of years, billions of people have found meaning in playing virtual reality games. In the past, we have called these virtual reality games ‘religions’. What is a religion if not a big virtual reality game played by millions of people together? (...) If you pray every day, you get points. If you forget to pray, you lose points. If by the end of your life you gain enough points, then after you die you go to the next level of the game (aka heaven).”

This idea is compelling in the sense that Harari has a point when he alludes to the human capacity, individual as much as collective, to engage in spaces and worlds that are not physically real. But would living in several realities only be keeping us busy or would it really reinvent man, giving him a new narrative of purpose and meaning? We should rather get our act together and discuss this before it is too late. “Indeed, the right time is now” (2 Corinthians 6:2).

Interview with Martin Rees

What would happen if the wealth created by robots and artificial intelligence just rested with a few super-rich people in Silicon Valley? Martin Rees, astrophysicist at the University of Cambridge and member of the House of Lords in England, would argue strongly for wealth redistribution. Otherwise social upheaval will be more likely than robots taking over the planet.

ALEXANDER GÖRLACH: Artificial intelligence is the buzzword of our time. When we look into societal and media discourses: would you say they have captured the disruption that we are about to see?

MARTIN REES: Artificial intelligence will do better than humans at managing complex networks – city traffic, electricity grids and so forth. And it will transform the labour market. It won't just take over manual work (in fact, plumbing and gardening will be among the hardest jobs to automate), but will be able to do routine legal work, computer coding, medical diagnostics and even surgery.

But that's very far from achieving the human-level general intelligence that grabs media interest and remains on the speculative fringe. Some artificial intelligence pundits take this seriously, and think the field already needs guidelines – just as biotech does. But others regard these concerns as premature – and worry less about artificial intelligence than about real stupidity.

Atlas is a humanoid robot built by Boston Dynamics. Rather than being a clumsy robot that we are used to seeing in real life, Atlas has achieved amazing balance and agility



“A key question is whether robots can handle ambiguity”

You are one of the founders of the risk analysis centre at the University of Cambridge: on a scale between blessing and curse, where would you place artificial intelligence?

Among experts (and I'm not one) there's a spectrum of opinion about how long it will take for a general human-level intelligence to be achieved. Ray Kurzweil thinks it may take 25 years; Rodney Brooks (inventor of the robot vacuum cleaner) thinks it will never happen. I would place myself somewhere in the middle of that spectrum.

If robots could observe and interpret their environment as adeptly as we do, they would truly be perceived as intelligent beings, to which (or to whom) we can relate. What if a machine developed a mind of its own? Would it stay



MARTIN REES is a Fellow of Trinity College and Emeritus Professor of Cosmology and Astrophysics at the University of Cambridge. He holds the honorary title of Astronomer Royal and also Visiting Professor at Imperial College London and at Leicester University. He was a Royal Society Research Professor, and then from 2004 to 2012, Master of Trinity College. In 2005 he was appointed to the House of Lords, and he was President of the Royal Society for the period 2005–10. Martin founded the Centre for the Study of Existential Risk with Huw Price and Jaan Tallinn. Martin has extensive experience engaging with both policymakers and the public to better prepare for global long-term challenges.


“I think it's in space rather than here on Earth that artificial intelligence will fulfil its greatest long-term potential”

docile, or ‘go rogue’? If it could infiltrate the internet – and the internet of things – it could manipulate the rest of the world. It may have goals utterly orthogonal to human wishes – or even treat humans as an encumbrance. Be that as it may, it's likely that society will be transformed by autonomous robots, even though the jury's out on whether they'll be ‘idiot savants’ or display superhuman capabilities. A key question is whether they can handle ambiguity and the unexpected as well as a human can.

As a person that has been part of scientific disruption in the last half-century in your field, as an astrophysicist and cosmologist, how do you personally perceive the changes that we are and will be witnessing?

Perhaps because I'm an astrophysicist, I think it's in space rather than here on Earth that artificial intelligence will fulfil its greatest long-term potential. Space is a hostile environment to which humans are ill-adapted. But near-immortal electronic and non-organic intelligences will be able to roam the universe, free of the constraints of organic creatures.

Out of all great transformations we are going through, from climate change to artificial intelligence to gene editing, what are the most consequential risks we are about to witness?



With the help of a super-computer, Luciano Rezzolla, an astrophysicist at the University of Frankfurt, is able to simulate the collision of neutron stars

“The risk of error or terror is quite substantial”

It depends on what timescale we are thinking about. In the next 10 or 20 years, I would say it's the rapid development in biotechnology. Already it's becoming easier to modify the genome, and the 2012 “gain of function” experiments, rendering the influenza virus more virulent and transmissible, are a portent of things to come. These techniques offer huge potential benefits, but catastrophic downsides as well. And the other point about them is that they are easily accessible and handled. The equipment they require is available in many university labs and many companies. And so, the risk of error or terror is quite substantial, whilst regulation is very hard. It's not like regulating nuclear activity, which requires huge special purpose facilities. Bio-hacking is almost a student-competitive sport. Obviously, we should try to minimise the risk of misuse of these techniques, whether by error or by design. We should also be concerned about the ethical dilemmas they pose.

Do you fear that this doesn't just happen in the realm of crime – if we think of so called “dirty bombs” for example – but also the possibility that governments might apply these techniques? Do we need a charter designed to prevent misuse?

Governments haven't used biological weapons much. That's because their effects are unpredictable, there is a risk of “bioerror” – leakage

of pathogens from a laboratory, for instance. And there is a risk of “bioterror” by mavericks or extremists – for instance eco-fanatics who think that humans are so numerous that they are polluting the planet and jeopardising biodiversity. We do indeed need internationally-agreed regulations, for both ethical and pragmatic reasons. But my worry is that these cannot be effectively enforced globally – any more than the drug laws or the tax laws can be.

That brings to mind recent Hollywood blockbusters like “Inferno”, where one lunatic tries to sterilise half of mankind through a virus.

Several movies have been made about global bio-disasters. A pandemic, whether natural or malevolently induced, could spread globally at the speed of jet aircraft. We have had natural pandemics in historic times, for instance the “black death”, which – though regional and not global – killed at least a third of the inhabitants of some European towns. But even when that happened, the surviving citizens were fatalistic and life went on as before. But today we have high expectations, and there could be societal breakdown even for a one per cent casualty rate, because that would overwhelm the capacity of hospitals. That is why governments put pandemics – natural or artificially produced – high on their risk register.



People in sub-Saharan Africa (and other poor regions of the world) are now more aware of the Western way of life, and they wonder why they cannot enjoy it too

So, when speaking of the age of transformation, aspects of security seem paramount to you. Why is that?

We are moving into an age where small groups can have a huge and even global impact. In fact, I highlighted this theme in my book “Our Final Century” which I wrote thirteen years ago. These new technologies of bio and cyber can cause massive disruption. We have always had traditional dissidents and terrorists but there were certain limits to how much devastation they could cause. That limit has risen hugely with these new bio and cyber-technologies. I think this new threat is going to pose challenges to governance and increase the tension between freedom, security and privacy.

Let’s look at another huge topic: artificial intelligence. Is this a field where more uplifting thoughts occur to you?

Within a timeframe of ten to twenty years, I think the prime concerns are going to be cyber-threats and bio-threats. However, as I’ve already said, the labour market will be disrupted because robots will take over many

occupations. To ensure we don’t develop even more inequality, there has got to be heavy taxation and massive redistribution. The money earned by robots can’t just go to a small elite – Silicon Valley people for instance. It should be recycled, so that social-democratic nations can fund dignified, secure jobs where the “human touch” can’t be replaced by a machine: carers for young and old, teaching assistants, gardeners in public parks, custodians and so forth. There is almost unlimited demand for jobs of that kind – there are currently far too few, and they’re now poorly paid and low status. But of course, most workers want more leisure – for entertainment, socialising, rituals, etc.

But robots could potentially also take on the work of a nurse for that matter.

True, they could do some routine nursing. But I think people prefer real human beings. At the present time, the wealthiest people (the only ones who have the choice) want personal servants rather than automation. I think everyone would like to be cared for by a real person in their old age.

In your opinion, what mental capacities will robots have in the near future?

I think it will be a long time before they will have the all-round ability of humans. Maybe that will never happen, we don’t know. But what is called generalised machine learning, enabled by the ever-increasing number-crunching power of computers, is a genuine big breakthrough. But the development of sensors has a long way to go. If these computers were to “get out of their box”, or infiltrate the “internet of things”, they might pose a considerable threat.

In your opinion, what sparks new innovation and ideas? Will artificial intelligence and machine learning foster these processes?

Eureka moments are quite rare, sadly. They do happen, but – to quote Pasteur – “Fortune favours the prepared mind”. You have got to ruminate a lot before you are likely to achieve important insights. The big breakthroughs in scientific understanding are often triggered by some new observation that in turn was enabled by some new technological advance. New insights often require a collaboration between people who can cross disciplines. Computer simulations will supplement (or even replace) experiments; and allow huge data sets to be analysed. There are some scientific challenges which everyone agrees are important, but which receive little attention until there seems genuine hope of progress. For instance, the “origin of life” is one such problem which is only now receiving mainstream attention.

Would you say a collective can have an idea or can only individuals have ideas?

Most breakthroughs are really the outcome of a collective effort. In football one person may score the key goal – but that doesn’t mean the other ten people on the team are irrelevant. I think a lot of science is very much like that: the strength of a team is crucial to enable one person to score the goal.

Do natural sciences and humanities have the capability to tackle the challenges occurring from these transformations?

Here in Cambridge in the United Kingdom we are trying to use our university’s convening power to address which long-term existential threats are real and which can be dismissed as science fiction, and to recommend how to

reduce the probability of the credible ones. This requires expertise from the social sciences as well as natural sciences. For instance, I’ve already mentioned that, because of the societal effect, the consequences of a pandemic now could be worse than it was in the past, despite our more advanced medicine. Also, if we are thinking of problems like food shortages, the issue of food distribution is an economic question, as well as a question of what people are ready to eat. Are we, for instance, going to be satisfied eating insects for protein?

With the rising amount of aggregated data it becomes increasingly difficult for the humanities to keep up with natural sciences. How can we synchronise the languages of different academic fields in times of big data?

We need to encourage people to bridge these boundaries. I am gratified that our Cambridge group addressing extreme risks has attracted young researchers with real breadth: philosophers who are into computer science; and biologists interested in system analysis. Here in Cambridge we are advantaged because of our college system. In most universities you don’t meet people from other departments until you become very senior (a department chair or suchlike). But each college is a microcosm, covering all disciplines, so even the most junior researchers have daily exposure (at lunch or in the common room) to experts in all fields. So, Cambridge is a particularly propitious environment for cross-disciplinary work.

The blessings of modern innovation seem to be ignored by many policymakers: we see a retreat from globalisation, a retreat from digitalisation – is it a disconnect between science and the rest of society?

The misapplication of science is a problem, of course. So is the fact that science’s benefits are irregularly distributed. The welfare of the average blue-collar worker and their income in real terms – in the US and in Europe – has not risen in the last twenty years and in many respects their welfare has declined. Their jobs are less secure and there is more unemployment. But there is one aspect in which they are better off: IT. Information technologies spread far quicker than expected and led to advantages for workers in Europe, the US, and Africa.

“To ensure we don’t develop even more inequality, there has got to be heavy taxation and massive redistribution”

But surely globalisation has made many poor people less poor and a few rich people even richer.

Sure, but let's remember that we're now witnessing a significant backlash in many places, in terms of Brexit or the presidential election in the US.

How drastically do you think these developments will affect science, the attitude towards it and its funding?

Many of the people who use smartphones and the internet aren't aware that the fantastic underlying technologies can be traced back to scientific innovations decades ago, which were mainly funded by either the military or the public. It's unfair to say people are anti-science – indeed I find it gratifying how much public interest there is in topics from black holes to dinosaurs that have no direct practical relevance. But they are worried that some technologies will run ahead faster than we can control and cope with them. I think there's good reason to be concerned, for example, about biotech and cyber – to maximise the benefits while trying hard to avoid the downsides. For technology to be developed, it's not enough to know the relevant science. There needs to be an economic or political imperative. For instance, it took only twelve years from the first Sputnik to Neil Armstrong's "one small step" on the moon. The motivation for the Apollo programme was a political one and four per cent of the US federal budget was committed to it (in contrast to the 0.6 per cent that NASA gets today). In the case of IT, there was the obvious demand, which led to the internet and smartphones spreading globally at a rate exceeding most predictions. But commercial flying presents a contrasting example – today, we fly in the same way we did fifty years ago, even though in principle we could all fly supersonic.

Living in a so-called post-factual era, what are "facts" to you as a scientist?

If we take as an example the Brexit vote in the UK: those who voted for Brexit had a variety of motives. Some wanted to give the government a bloody nose, others voted blatantly against their own interest. The workers in South Wales, for example, benefited hugely from the EU. There is a wide variety of different motives, but I don't think people would say that they voted against technology. And I still hope for an "exit

“A huge benefit to society will arise if we have enough redistribution to recreate dignified jobs”

from Brexit" as the UK public realises what they've let themselves in for.

... Still there is this ongoing narrative about the fear of globalisation and digitalisation, and that would also imply the fear of technology.

Sure, but that is oversimplified. We can have advanced technology on a smaller scale. It allows for robotic manufacturing, it allows for more customisation to individual demand. The internet has enabled small businesses to flourish. Clean energy may be generated locally rather than delivered via vast grids.

But there seems to be an increasing disconnect in many societies regarding the consensus on which facts matter and how facts are perceived.

To understand this attitude you are expressing, we have to realise that there aren't many facts that are clear and relevant in their own right. There are often real grounds for scepticism. Most economic predictions, for example, have pretty poor records, so you can't call them facts. In the Brexit debate, there were valid arguments (as well as a lot of bogus ones) on both sides. And in the climate debate, even those who agree on the science and its margin of uncertainty will differ in the policy response they favour. For instance: how strongly should we bet on some technological "fix"? And how big a sacrifice should we make today to reduce the probability of a catastrophe in remote parts of the world a century hence?

But how then do you judge the developments we now see in many Western societies?

New technologies have led to new inequalities and new insecurities. Moreover, people are now more aware of inequality. People in sub-Saharan Africa are now fully aware of the kind of life that we Europeans enjoy, and they wonder why they can't enjoy it too. 25 years ago, they were far less aware of this unjust disparity. This understandably produces more discontent and embitterment. There is a segment of society, a less educated one, which feels left behind and unappreciated. That is why I think a huge benefit to society will arise if we have enough redistribution to recreate dignified jobs. The rich world needs to subsidise factories in the developing world, to reduce the incentive for migration.



Robonaut 2 was launched and delivered to the International Space Station in 2011. It is operated by crew members as well as by controllers on the ground

What political framework do you think of as an ideal environment for science?

The Soviet Union had some of the best mathematicians and physicists, partly because the study of those subjects was fostered for military spinoff reasons. People in those areas also felt that they had more intellectual freedom, which is why a bigger fraction of the top intellectuals went into maths and physics in Soviet Russia than probably anywhere else ever since. That shows you can have really outstanding science in many social systems. But of course, I support – for much broader reasons – a Scandinavian-style social democracy. And I am opposed to the austerity and "small-state rhetoric" deployed by the present UK government.

So, the ethical implication is not paramount to having "good" science after all?

I think scientists have a special responsibility. Often an academic scientist can't predict the implications of his or her work. The inventors of the laser, for instance, had no idea that this technology could be used for eye surgery and DVD discs, but also for weaponry. Among the most impressive scientists I have known are some of those who worked at Los Alamos on the atomic bomb. They returned to academic pursuits after the end of World War II with relief, but felt obliged to do what they could to control the powers they had helped to unleash. Most of these scientists supported the making of the bomb in the context of the time. But they were also concerned about proliferation and arms control. It would have been wrong for them to not be concerned, even though their influence was limited. To make an analogy: if you have teenage children, you may not be able to control what they do, but you are

a poor parent if you don't care what they do. Likewise, if you are a scientist your ideas are your "offspring", as it were. You can't necessarily control how they will be applied, but nonetheless you should do all you can to ensure that they are used for the benefit of mankind and not in a damaging manner. This is surely an attitude that should be instilled into all our students.

What, then, is your motivation as a scientist?

I feel I am very privileged to have, over a career of forty years, played a modest part in debates on topics which I think will be highlights when the history of science in this period is written – understanding the evolution of the universe and its constituents. I think it is a great collective achievement. Many of the questions that were being addressed when I was young have now been solved. We're now tackling questions that couldn't even have been posed back then. Of course, the science I do is very remote from any application, but it's of great fascination and a very wide audience is interested in these questions. It certainly adds to my satisfaction that I can convey the essence of these exciting ideas to a wider public. I would get less satisfaction if I could only talk about the cosmos to a few fellow specialists.

What is the best idea you've ever had?

I've never had any singular idea, but I think I have played a role in some of the insights that have gradually firmed up our view of how our universe has evolved from a simple beginning to the complex cosmos we see around us and of which we are a part. And the social part of science is very important. Many ideas emerge out of cooperation – and, of course, from experimenters and observers, who deserve far more credit than theorists like myself. Incidentally, the old idea that science eventually leads to an application is far too naïve. The interaction goes both ways because advancements made in academic science are facilitated by technology. If we didn't have computers or ways of detecting very faint radiation, etc., we would have made minimal progress in astronomy. We were no wiser than Aristotle was, and we only advanced beyond him through having much more sensitive detectors and being able to explore space via many techniques.

“We might want to rediscover our social roots”

Interview with Pascal Finette

Democracy needs to adjust quickly to the new, fast-moving reality of our present time, argues Pascal Finette, from Singularity University. The innovation and technology aficionado does not see Amazon, Google or Facebook replacing political power structures. He would therefore urge politicians to move quickly forwards with creating new societal frameworks, such as universal basic income.

ALEXANDER GÖRLACH: You are basically researching and exploring technological progress, in digital, data mining, artificial intelligence. Innovation on a daily basis. What is the first thing you tell people who have no idea what your work is about nor what the future you describe so enthusiastically will look like?

PASCAL FINETTE: I believe the most important thing you have to understand, is that tomorrow will look dramatically different from today. And by tomorrow, I don't mean ten years from now – but quite literally the day after today.

Technology is advancing at an unprecedented speed and the rate of change keeps accelerating. It's exhilarating (and sometimes scary) to witness this, study it and figure out what tomorrow (and the days after tomorrow) will look like.

Is this accelerated change a blessing to you or a curse? The closer you draw the line of change, messianic hope and eschatological fear fall into one. This is why theories of technological progress such as the singularity have been called ideological.

I don't think it's binary or black and white. Exponentially accelerating technologies present amazing opportunities and, of course, can be misused by ill-meaning individuals and organisations. The technology itself (and thus also the “singularity”) is rather agnostic towards its use for good or ill. Which puts the human into the centre of the equation – it is on us to decide how we want to use technology. And when it comes to humans I am, by and at large, optimistic. As we wrote into our operating principles at eBay in the late 90s: we believe people are basically good.

This optimism is also found in the guiding and basic principles of Google “Don't be evil”. We have gone quite philosophical in under a minute. Is being ethical, being capable of distinguishing between good and evil, something that algorithms or machines are capable of doing? I mean, frankly speaking, if artificial intelligence takes its essence from us, the humans, our data, our

This poster is around 8,000 Square metres in size and identified as the biggest poster in the world. It was installed by activists in a public square in Geneva, Switzerland, ahead of a referendum on a universal basic income on 5 June 2016



behaviour, how could it then possibly be agnostic? Agnosis in Greek means not knowing. Machine learning is all about knowing ...

I believe that, for the foreseeable future, machines won't make ethical or moral decisions on their own. You bring up a very good point, though – artificial intelligence is based on machines learning; thus, it matters greatly that our inputs aren't biased or incomplete. My former employer Mozilla just launched an initiative to overcome the bias speech recognition systems have, due to the limited training sets they are exposed to.

So, machines are basically not only replicating the biases, prejudices and injustices of our societies, but also expose us to them exponentially: they narrow down the world we live in. Is there a chance that we could become better humans by working on eradicating biases from machine learning?


“I don’t think nor believe Google (or Facebook or Apple etc) is in charge”

Just to clarify this – a lot, if not the vast majority, of applications for artificial intelligence/machine learning won't have any bias problem as the data sets they are trained on have no biases. Think about all the industrial applications for artificial intelligence or something like weather forecasting – the data sets for these fields are just that: data (without human bias). The challenge comes when we train artificial intelligences on human questions – for example, voice recognition as in the Mozilla case. And yes – artificial intelligence has the potential to make us better people by doing the exact opposite as well. Think about a personalised newsfeed which presents me with balanced views on a topic of interest, instead of solely partisan views.

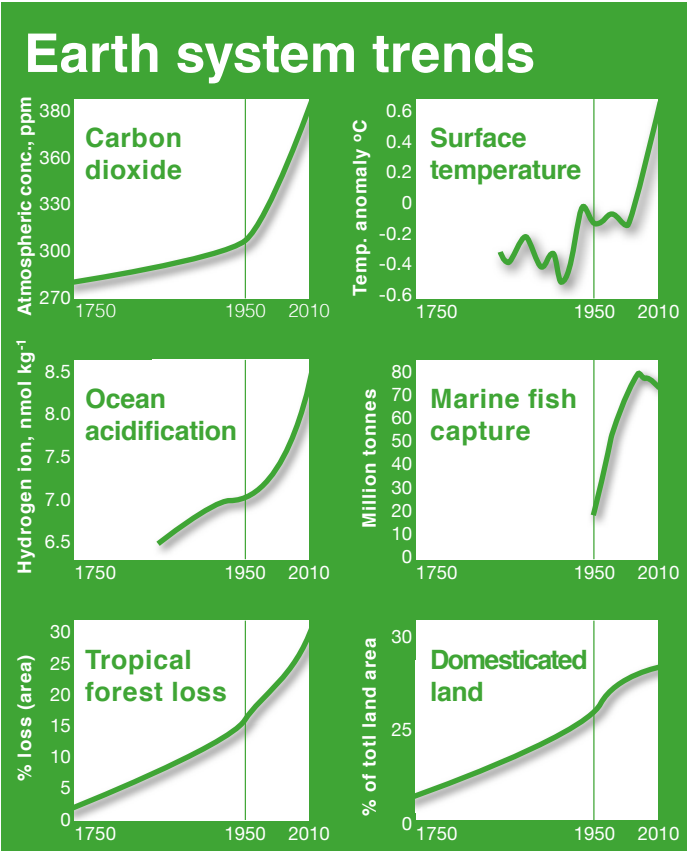
Let’s stay with the positive, non-biased aspects of machine learning for a bit. What field in your opinion will be profiting from this the most: health, mobility, governance?



PASCAL FINETTE
heads up everything Entrepreneurship at Singularity University, including the start-up programme, Global Expansion and the Entrepreneurship Track where he inspires, educates and empowers entrepreneurs tackling the world's most intractable problems leveraging exponential technologies. Pascal got started on the net before there was a web browser, founded a couple of technology start-ups, led eBay's Platform Solutions Group in Europe, led Mozilla's Innovation Lab, created Mozilla's accelerator program WebFWD, and has invested into social impact organisations around the globe at Google.org.



Data: Steffen W et al. *The trajectory of the Anthropocene: The Great Acceleration*. The Anthropocene Review, 2: 81 –98, 2015



“Universal basic income makes a lot of sense to me”


Any field which generates large amounts of data and makes decisions based on this will greatly benefit from machine learning. Self-driving cars are only possible due to sophisticated machine learning algorithms and abundant computational power, IBM's artificial intelligence Watson has already become a better radiologist than humans¹ (see p 24), and large-scale farming is starting to rely heavily on machine-learning-based systems to increase agricultural yield. Governance is tricky as there is so much human behaviour involved.

In governance there are new approaches such as deliberative democracy, a model that basically runs on the assumption that the parliaments we know do not really run by representing all groups of society. Whereas when you use algorithms to sort out these representation problems, solutions for crucial questions are found. What do you make out

1 www.technologyreview.com/s/600706/ibms-automated-radiologist-can-read-images-and-medical-records/

of this sort of new approach for democracy? An intriguing approach. Fluid democracy and even more immediate and direct forms of electronic voting are not only interesting but also concepts which have the potential to change the way we live democracy. The challenge, from my perspective, is that we are dealing with systems which are rather encrusted, move with election cycles of four years or more and have strong powers at hand which like to keep systems the way they are. Combine this with the fact that the civil sector is not exactly the employer of choice for some of our brightest and fastest moving minds and you can see why changing governance is a tall order.

The question therefore has been raised more than once. Who is in charge: Google or the government? The big digital innovators with all the young high-potential coders working for them, or the governmental institutions entrusted with law-making yet not attracting the brightest and boldest (any more)?



In 2011, IBM Watson won Jeopardy requiring mastery of general knowledge and natural language processing. IBM's Avicenna software highlighted possible embolisms on this computer tomography scan in green, finding mostly the same problems as a human radiologist who marked up the image in red

I don't think nor believe Google (or Facebook or Apple etc) is in charge. They might have influence over what we see and do online – but they surely don't have the power to put you in jail for your behaviour. Thus, I am much more concerned about governments misusing or even abusing the data trails we leave behind than Google wanting to sell me more stuff on behalf of their advertising clients. And yes, it greatly concerns me that the brightest minds in countries around the world focus on building the next hip social-local-mobile photo sharing app instead of focusing on solving humanity's grand challenges.

Certainly, Google may not put you in jail, yet if you had to choose between a day in jail and no more access to all your Google services I am quite certain many people would prefer jail over losing all that you may have in the Google cloud. But that's (lucky us!) only hypothetical. You have been involved in lots of this work focusing on solving humanity's grand challenges. And you said at the beginning of our conversation that the changes through artificial intelligence will be visible and experience-able in a short while from now. What is the most ground-breaking thing we will be seeing next year: no more droughts? Beating cancer? Much of the suspicion an average Joe may have when it comes to artificial intelligence is that it carries the face of Arnold Schwarzenegger, the Terminator, rather than a smiling, positive countenance.

We're making massive strides in healthcare. Not just through the abilities of artificial intelligence but also genetics, digital biology and stem cell therapy. It surely makes you optimistic about solving diseases ranging from cancer to sickle cell. Interesting work is being done on the intersection between man and machine – prostheses connected to your nervous system, allowing amputees to walk or operate their robotic arm with great precision. And don't get me wrong, we are only at the very beginning of what's possible. A lot of what we see today is still crude, doesn't work quite as advertised or is just not that helpful. Every new technology goes through these phases.

What do you make of the debates and claims for a robot tax or universal basic income to cushion the societal impacts of artificial intelligence?

The robot tax I just don't get. Universal Basic Income (UBI) makes a lot of sense to me and we have some promising preliminary studies suggesting that UBI works in fostering a more entrepreneurial culture and thus increased GDP overall, plus more satisfaction and fulfilment. It's early days though – we need to run many more studies to find the right formula.

If I understand it correctly, the case for a robot tax is twofold: first it's about disentangling our views on taxation of a person's labour and secondly, it's about exploring new ideas of redistribution. Above all it's about preventing a new global elite from accumulating the wealth generated through the work of robots, which also extends to the increasing work done by artificial intelligence. And it tries to tackle the question of fair taxation (if fewer and fewer people are involved in financing the state's functions) and redistribution in general.

Arguably we are miles away from "fair taxation" – as is demonstrated by the increasing divide between the top one percent and the rest of the population in most, if not all, countries around the world or the fact that, although the economy is growing, it doesn't translate to increased incomes in the middle class for the last couple of decades. I can't see how the robot tax solves this – sadly. And surely it is a hard problem to solve. Personally speaking, I miss the debate – not only haven't we figured out the problem, to a large degree we don't even talk about it.

So, what is your take then on fairer, better societies through the progress provided by artificial intelligence? I see the debate is already at full pace, yet I am not sure if the prospect of a world without work, where we may all be painters and poets, is the right scenario to start with.

I believe it brings up at least two relevant and important questions: one about the financial implications, the other about the human implications. Today, in our society, we derive a large amount of our self-worth from the job we do. It not only provides the financial means to live (and hopefully thrive) but contributes to our sense of self and societal status. Take a job away and we need to rethink (and enact) a whole lot more than just the financial safety net.

“Today, in our society, we derive a large amount of our self-worth from the job we do”

I couldn't agree more, but what do we do about it? Historian Yuval Noah Harari recently wrote an op-ed for the Guardian claiming that millions of people may engage in virtual reality games in order to be kept busy. Is that the great new world we're headed towards?

That's a bit apocalyptic for me. But then – the argument can be made that a virtual reality will sometime in the future be better and more engaging than what geeks like to call "meat-space". Which raises the question: in which space do we choose to live?

Harari's point was that we, the human species, have engaged in this sort of virtual reality game for as long as we have been around: this game is called religion. We collect scores, we ascend or we descend, heaven, hell or purgatory. To me that is not necessarily apocalyptic; rather compelling in the sense that playing games (*Homo ludens*) seems to be more part of our nature, more satisfying and therefore more part of our identity then we may think on a daily basis. But let me learn more about "meat-space".

Without going into too much of the long-standing debate about (organised) religion, I find myself more on the spectrum of believing in self-determination. Probably more importantly, the scenario outlined in Harari's post requires (to some extent) the belief in a benevolent deity/ruler.

What I find the most irritating, and that goes along with your belief and approach, is, that people by engaging in that virtual world are stripped of any self-determination and destiny whatsoever. For me this scenario raises the question: what will give us humans dignity when it is not work anymore, what will define our self-esteem when it is not the narratives of our achievements?

You bring up an excellent point. I think we have to consider two directions in (trying to) answer this question: On the one hand, we need to expand the definition of "achievements" to be more inclusive than just our career achievements, to include topics such as our civic duties/engagement, etc. And on the other hand, we might want to rediscover our social roots and derive fulfilment from our social interactions and care for the communities we live in.



“We cannot allow ourselves to think in a utopian way”

Interview with Luciano Floridi

For the philosopher Luciano Floridi, it is crucial to understand what we want to use new technologies for. Running the Digital Ethics Lab at the Oxford Internet Institute, it is his job to ask these questions. So far, the pundits of artificial intelligence do not engage much in the question of what ethical codes the robots of the future would have to subscribe to.

ALEXANDER GÖRLACH: Being the Director of Research of the Oxford Internet Institute seems to be a bit anachronistic itself – “Oxford”, one of the oldest universities in the world, and “internet”, the bringer of modernisation and globalisation. How do you unite the two?

LUCIANO FLORIDI: Indeed, it is an interesting combination of words. Take the centuries of tradition, combine it with the complete novelty and unprecedented problems we’re facing, it does not sound straightforward. But looking at the problems we’re facing, it seems like a wonderful recipe to combine one of the world’s best universities and an institute that focuses on the internet, something that gives us sheer endless opportunities in the world. I can’t imagine a better way of combining the old with the new.

But is Oxford playing catch-up in that regard? Or is it a go-ahead, a place where you develop new theories and propose ground-breaking policies? It seems like much talent is attracted to the United States ...

... Well – both. When it comes to understanding what’s happening, as opposed to identifying new phenomena, there is a lot to register in

terms of novelties. Take the job market, for example – thanks to the internet, we now see huge internet companies such as Amazon enable thousands of people to be hired in the shortest of time spans, something that was unthinkable a few years ago. There is a lot of catching up in terms of understanding where the issues that the internet has brought with it lie. At the same time, we as a society should be leading the shape of things in terms of policy. I see it as two steps; one is the understanding part, and one is the policy-making part. Both influence each other, seeing that policy always affects what is happening in the real world.

Interesting – there is actually an argument, made by Jean Baudrillard, who says that America has reached a realistic utopia, whereas Europe is kept behind due to the weight of its history. However, in a utopian world, by definition there wouldn’t have to be any policies, correct?

Yes, I am familiar with that argument. It is not quite about that though. Utopia, by default, is anti-historical. It establishes a status quo that remains there, forever – you don’t improve on perfection. For us in Europe, we cannot allow ourselves to think in a utopian way. We have a lot of history, and we know that more history is coming. All has happened, and during every great chapter of history, we thought history itself had come to an end – the Egyptians thought so, the Greeks thought so, the Romans did, too. More recently after World War II, after the fall of the Soviet Union, at the beginning of the European Union. It is possible in some contexts to have utopian thinking, because to some extent you do not have enough history to show that reaching a point of no further development is impossible in human nature.

Painting of Scene from the Short Story “The Thought Machine” by Anton Brzezinski (born 1946). The American artist of Polish descent is known as “The Living Dali” or “The Polish Picasso” through his surrealist and cubist works of art

So, would you say Europeans are disillusioned by it?

Quite the contrary, I think us Europeans are more realistic. The disillusion may come in the form of a cynical view on things, but I think overall the Europeans have a grown-up attitude that things change, have always and will always keep on changing. The ability of politics in this case is to handle change, and not to stop change.

The counter-narrative would be the Euro-sceptic view, especially of the economy, in which most political parties in our day are an offshoot of one another.

To add to that – politics is and always has been driven by social issues. Germany had to grow in the early 1900s because the global market was dominated by a few colonial powers, so World War I happened. For the past 50 or 70 years, since World War II, slowly we have seen that the ground for politics has become purely economic. Listen to any politician, and what they discuss are topics of GDP, unemployment, growth. Everything has to do with the economy, and the debate does not focus on “are people happy?”. All the questions of happiness, social justice and equality become a by-product of economic progress or recession. We go from sociopolitical issues to mainly economic issues, and I think the future for us has to be based on other values than economic ones. I don’t like to talk about post-modernism, because in my view it is an admission of having run out of ideas, but we have to reposition our political discussion.

And with our technology developing, don’t we get the chance to actively make those decisions? Take artificial intelligence – naturally, we are very sceptical of the ethical implications a robot may have, and it has been debated for a long time now. Is that not reflecting our moral bankruptcy? Since we are doing it anyway, can it not be seen as a way for humankind to give in to the economic advantages of lower costs and consistent quality?

Maybe – we can dig deep and look at the logical consequences. The underlying question we need to ask ourselves, though, is: what do we want to do with technology? I disregard completely the ideas of singularity and artificial intelligence domination, some-

thing I see more as a scratch to address certain itches. The scratching is wrong, but the itching is real. It does mean that there is a problem deep down, namely “What future do we want to have?”, when already our world and economy are so deeply dependent and shaped by technology. It’s just another way of saying the digital economy will dominate our life – is it time to give that control a certain shape? Artificial intelligence means that we have the added possibility of countless new advances in what is possible digitally. It is up to us to determine what possibilities will become a reality.

In terms of humans, how does it affect the two aspects of people? Being a citizen vs being a consumer?

There is a tension there, and it is because the circle of interactions has become wider and less visible. In the early days, there was the

LUCIANO FLORIDI

is the OII’s Professor of Philosophy and Ethics of Information at the University of Oxford, where he is also the Director of the Digital Ethics Lab of the Oxford Internet Institute, and he is Distinguished Research Fellow of the Uehiro Centre for Practical Ethics of the Faculty of Philosophy and Research Associate and Fellow in Information Policy of the Department of Computer Science. His research primarily concerns information and digital ethics, the philosophy of information, and the philosophy of technology.



production of goods, and the consumption of goods – industry and consumers. A third party like the state could regulate that exchange, but there were only three parties. Today, it’s no longer that way. Selling products to customers is not the main objective any more. In a digital economy, you mostly give your customers free things. They are not considered customers any more, but users. And a third party that regulates a “gift economy” serves no purpose. But if you add to that the analytics and advertisement that goes on, all in all this circle of interactions has become a lot less clear and a lot wider. To add to that, the notion that companies will sell customers’ data, which is what a lot of people fear as a result of the obscuring circle, is not worth paying much attention to – user data is the golden goose in this equation, and companies will sell only the golden egg, i.e. services based on the data they own. The may sell the possibility of targeting users, based on the data they own, but not the data itself. That will remain with the company.

But if data is so sacred to these companies, what does the dialogue between the governments and the big internet companies – Google, Amazon, Facebook – look like? Wouldn’t you say that they are completely against policymaking? And what role does the citizen play in it?

The citizens do not play a role in this, because they love what they are getting – free mail accounts, free videos, free websites, free news, free everything. Nobody in their right mind would be opposed to that! So – there is no way to forbid or sanction advertising. We need to remember that when we talk about “people”, we all play certain roles: user, consumer and citizen. Sometimes, especially in Europe, patient is also added in a medical context (we are prone to that because we are an ageing society). But now think of what we do if one thing does not work. As a citizen, if policies have failed us, we vote for someone else. As a consumer, you have laws that protect your purchase. As a user of a free service, you are stuck and are forced to walk away. If you don’t like the Google search engine, that’s your problem. There is a lack of accountability in this, all the way to the top. The five companies that put together a partnership for artificial intelligence – Microsoft, Google, IBM, Facebook and

“The digital economy will dominate our life”

Amazon – took the first, very positive step in the right direction. I am a huge fan of that, and there is a certain expectation towards society to push for accountability rules in technologies in general and in artificial intelligence in particular. We need a soft and a hard legal framework, and right now it is just a bit messy.

But then, if you describe those conferences, what is your impression of what artificial intelligence is and what the state of knowledge about artificial intelligence is?

There is a lot of excessive emphasis on what artificial intelligence can do. The people always talk about specific machines, what these machines can do and what the algorithm behind it is. People do not recognise the variability and degrees of technical implementation that different artificial intelligences have. The difference between an industrial robot versus a bot that updates a Wikipedia entry automatically versus a house robot that does the dishes is huge. The result is that we become overwhelmed by the immensity of possibilities. Artificial intelligence has been most successful in industrial robotics, and it is an area where it has worked incredibly successfully for decades. The car industry has always been at the forefront, and there are regulations in place. Just look at the self-driving car debate and the accidents. Whenever I speak to people around the world, I know that technologies do not respect boundaries – they cut across fields. It is all about finding a legal framework.

What about the notion of robots replacing us in the future?

It’s ludicrous. There is no such thing – as debates happen in a very esoteric context – university halls where you forget the rest of the real world or in heated debates in newspapers. My recommendation is: look out the door, and look at what there is in the world. Tell me then sincerely whether you, the person fearing a robot takeover, see anything that remotely hints at the emergence of such a scenario. However, it does not mean that it will never happen, it means that we should not worry about negligible sci-fi scenarios. The real issues raised by artificial intelligence are too serious and pressing to waste time wondering how we shall ever teach superior artificial minds to be ethical.



“We will seamlessly cooperate with machines”

Interview with Nuria Oliver

Between humans and machines there will be peaceful cohabitation, argues computer scientist Nuria Oliver. There are things machines are better at than humans, and vice versa. The future therefore belongs to a human-machine hybrid kind of being.

ALEXANDER GÖRLACH: Data is considered to be the oil of the 21st century, the basic commodity of our time. Oil, however, is nothing without being refined. What will the refinement of data look like in the next couple of years?

NURIA OLIVER: Indeed, data in itself can be seen as “digital garbage” if we are not able to make sense of it, to draw useful insights or learn and/or make better decisions thanks to it. A large percentage of the data available today is non-structured data. Hence, to be able to extract value from it we need to apply machine learning techniques. Some directions where I think that big data data analysis will progress in the next years include: (1) Real-time analysis and predictive models. A lot of the projects

Hugh Herr lost both legs in a climbing accident. Nowadays he is head of the Biomechatronics Group at the Massachusetts Institute of Technology and is building the next generation of bionic limbs and robotic prosthetics inspired by nature's own designs

analyse data post-hoc, i.e. data from the past. However, many use cases would benefit from being able to analyse the data in real-time and make predictions to help inform decision-making; (2) Multi-modal analysis. As the number of data sources increases, it will be increasingly important to be able to effectively combine data from different sources in the analysis; (3) Privacy by design approaches when dealing with personal data; and (4) FATE algorithms. That is, fair, accountable, transparent and ethical algorithms. As the presence of algorithms in our lives will be pervasive, we need to ensure that they satisfy these four conditions.

Part of your work includes the modelling of human behaviour. Can artificial intelligence help us to understand better who we (really) are, refine us better?

Indeed. Today, thanks to the ubiquity of technology in our homes, our cities, our workplaces and ourselves (i.e. mobile phones, wearables etc) we have an unprecedented availability of human behavioural data: where we go, how we spend our time, how we feel, how much we walk, sleep or eat, what we shop for, read, listen to or watch – are examples of what can be collected or automatically inferred from such data by means of machine learning algorithms. As the availability of data increases and machine learning techniques become more sophisticated, we will have the ability to infer more complex and nuanced aspects of who we are. This knowledge could be extremely valuable to help us improve our lifestyle, our wellbeing, better manage our time and ultimately realise our potential.

It is a long way from data collection and data utilisation to the creation of a “conscious mind” sort of intelligence. What inspires humankind to think in such a superlative utopian kind of way about the abilities of technology?

Technology has and will without a doubt transform who we are, how well and how long we live, the jobs that we do, the way we communicate with and relate to others, etc – technology in itself is a tool that can be used for a variety of purposes. It is the aspiration of many of us who have devoted our lives to technological research to ensure that the

“Technology has and will without a doubt transform who we are”

technologies that we invent and develop will have a positive impact in the quality of life for all of us and also for our planet. I am convinced that our survival as a species depends on our ability to develop the right technologies to address critical challenges such as global warming, the ageing of the population and chronic disease and the availability of limited resources, etc.

As a scientist: how plausible do you think these visions are?
I think they are very plausible. We are already able to tackle complex problems, cure diseases, travel to outer space, democratise education, improve productivity, use renewable energies, to name a few, thanks to technology.



NURIA OLIVER
is a computer scientist. She holds a PhD from the Media Lab at MIT. She is the first female Spanish computer scientist to be named an ACM Distinguished Scientist. She is a Fellow of the European Association of Artificial Intelligence and a IEEE Fellow. She is well known for her work in computational models of human behaviour, human/computer interaction, intelligent user interfaces, mobile computing and big data for social good. Nuria invests significant effort in outreach efforts to make technology more accessible to the general public.

“There are new jobs today that didn’t exist 20 years ago”

One of the “milder”, if you will, predictions is that there will be more human-machine interaction in the future, as in regard for example to implemented chips that may help us accelerate our brain capacity. It sounds also kind of like science fiction yet it seems to be highlighting what humans are very good at and what machines do at their best.
Yes, I am convinced that we will seamlessly cooperate with machines, both physical machines (e.g. robots, devices, cars) and algorithms. We already do it today. It will be of paramount importance that we develop FATE technology, that is, technology that will be fair, accountable, transparent and ethical. These four dimensions are right now actively being addressed in the research community as four important challenges that we need to address in order to maximise the positive impact of technology in our lives. We also need to ensure that we minimise the risk of having a gap between those who have access to technology and knowledge and those who do not. That’s why we need to invest in education at all levels, from primary school to citizens.

You have been trying to use big data, the basic achievements of artificial intelligence, for social good. How can big data help improve our societies?
The ability to make sense of big data through machine learning techniques can bring significant positive impact in different areas of our

societies. I have extensive experience of using aggregated and pseudo-anonymised mobile network data. We have shown that this data is valuable because it enables us to infer large-scale patterns of human mobility, human networks and compute accurate estimates of population counts, in a fully privacy-preserving manner. These variables (mobility, networks and population counts) are important in urban planning, when facing public health challenges (such as a risk of a pandemic) and natural disasters and emergencies. We have also found the population dynamics are helpful to understand the socio-economic development of a region, to model energy consumption or to automatically detect crime hotspots in a city.

There is a fear that artificial intelligence, will turn our societies upside down, especially the workforce. What do you make of this objection?
Every major technology has disrupted the workforce. Today many jobs that existed at the time of my grandmother have disappeared, such as telephone switchboard operator, factory lector, milkman, street lamplighter, ice cutter and transporter, lift operator, etc. At the same time, there are new jobs today that didn’t exist 20 years ago, such as mobile app programmer, social media manager, cloud computing expert, Uber driver, sustainability manager, drone pilot, driverless car engineer, etc. From my perspective, the most important element is that we prepare both existing and future workers for the changes that technological progress will bring so they can contribute and be relevant in tomorrow’s society.

In regard to practical changes in the next few years: are we going to see the self-driving car? What other innovation will stun us?
There are several areas where technology could have a profound impact. One of them is healthcare. We are moving towards a model of predictive, personalised and preventative medicine, which would represent a very significant shift in the way we deal with disease. Another area is education, where, thanks to technology, we will be able to have personalised, multi-modal education that would be optimised to each student. We should also be able to communicate using technology in a richer way and eventually using our thoughts.

We will also be able to develop fully sustainable cities with a zero or even negative environmental footprint, so that there would be hope for our planet. Manned missions to space should have also progressed and the colonisation of Mars or some other planet with humans would be attainable.

Do you see governments and legislators as sufficiently prepared for these changes?
No. I worry about the gap that exists between an elite of us who know and understand how today’s technology works, and a large majority of people (not just legislators or decision-makers, but also children, young and older people) who do not have the technical skills to be able to understand today’s highly technological world. I worry about our educational systems, which are not up-to-date with what will be needed to contribute in tomorrow’s society; I worry about a technology and data literacy gap that we urgently need to address. That’s one of the reasons why I work with Data-Pop Alliance, as we have data literacy programs for governments, decision makers and citizens; and that’s also one of the reasons why I am proud to work for Vodafone, as they have several initiatives to promote digital literacy and education.

What made you enter this field in the first place?
I have always been fascinated by the figure of a researcher and an inventor. Since I was a child, my idols were Leonardo da Vinci, Marie Curie, Einstein, Ramon y Cajal. I am a very curious person with lots of interests in many areas, not just science. I love to study and learn. I also love puzzles and unsolved problems. Therefore, being a researcher is a very good fit for me. Regarding technology, I always loved the sciences but didn’t know much about computer science or electrical engineering while I was growing up. When I was in my last year of high school, I had the opportunity to talk to one of my brother’s friends who was studying electrical engineering. After he described to me what the career was about, I decided that I wanted to devote my life to technological research and innovation. I feel very lucky that I have been able to do so. It’s extremely motivating to feel that I am contributing with my work to create a better future for all.

“We are moving towards a model of predictive, personalised and preventative medicine”

“Some machines are people in the philosophical sense”

Interview with Huw Price

As a philosopher of time, Australian thinker Huw Price argues for a pause: we need to understand better what our human intelligence is, before we engage in much chatter about artificial intelligence. So far, he says, the future is ours, not the machines', as humans alone are capable of mentally projecting scenarios in a time to come.

ALEXANDER GÖRLACH: When looking into the future of artificial intelligence and machine learning, what do we have to expect?

HUW PRICE: I am as little of a futurist as a technical guy. My role in this field here in Cambridge is that I am an enabler, helping to make things happen. Things that will – with the creation of our new centre for the future of intelligence – help to foster a community of people who can ask the right questions. At the moment we are probably not even asking those.

Speaking about the centre – what kind of intelligence we are talking about here?

A good question! It sometimes comes up from a slightly sceptical viewpoint. People say: we don't even know what intelligence is, how can you propose to set up a centre about its future? My response to that is: let's not think about what intelligence is but what intelligence does. There are things about us that are responsible for us being the most successful species on the planet. Whatever that is, it is possible – at least in principle – for machines to do it too. So, whatever these specifically humanoid ingredients in us are, they don't exist in other species.

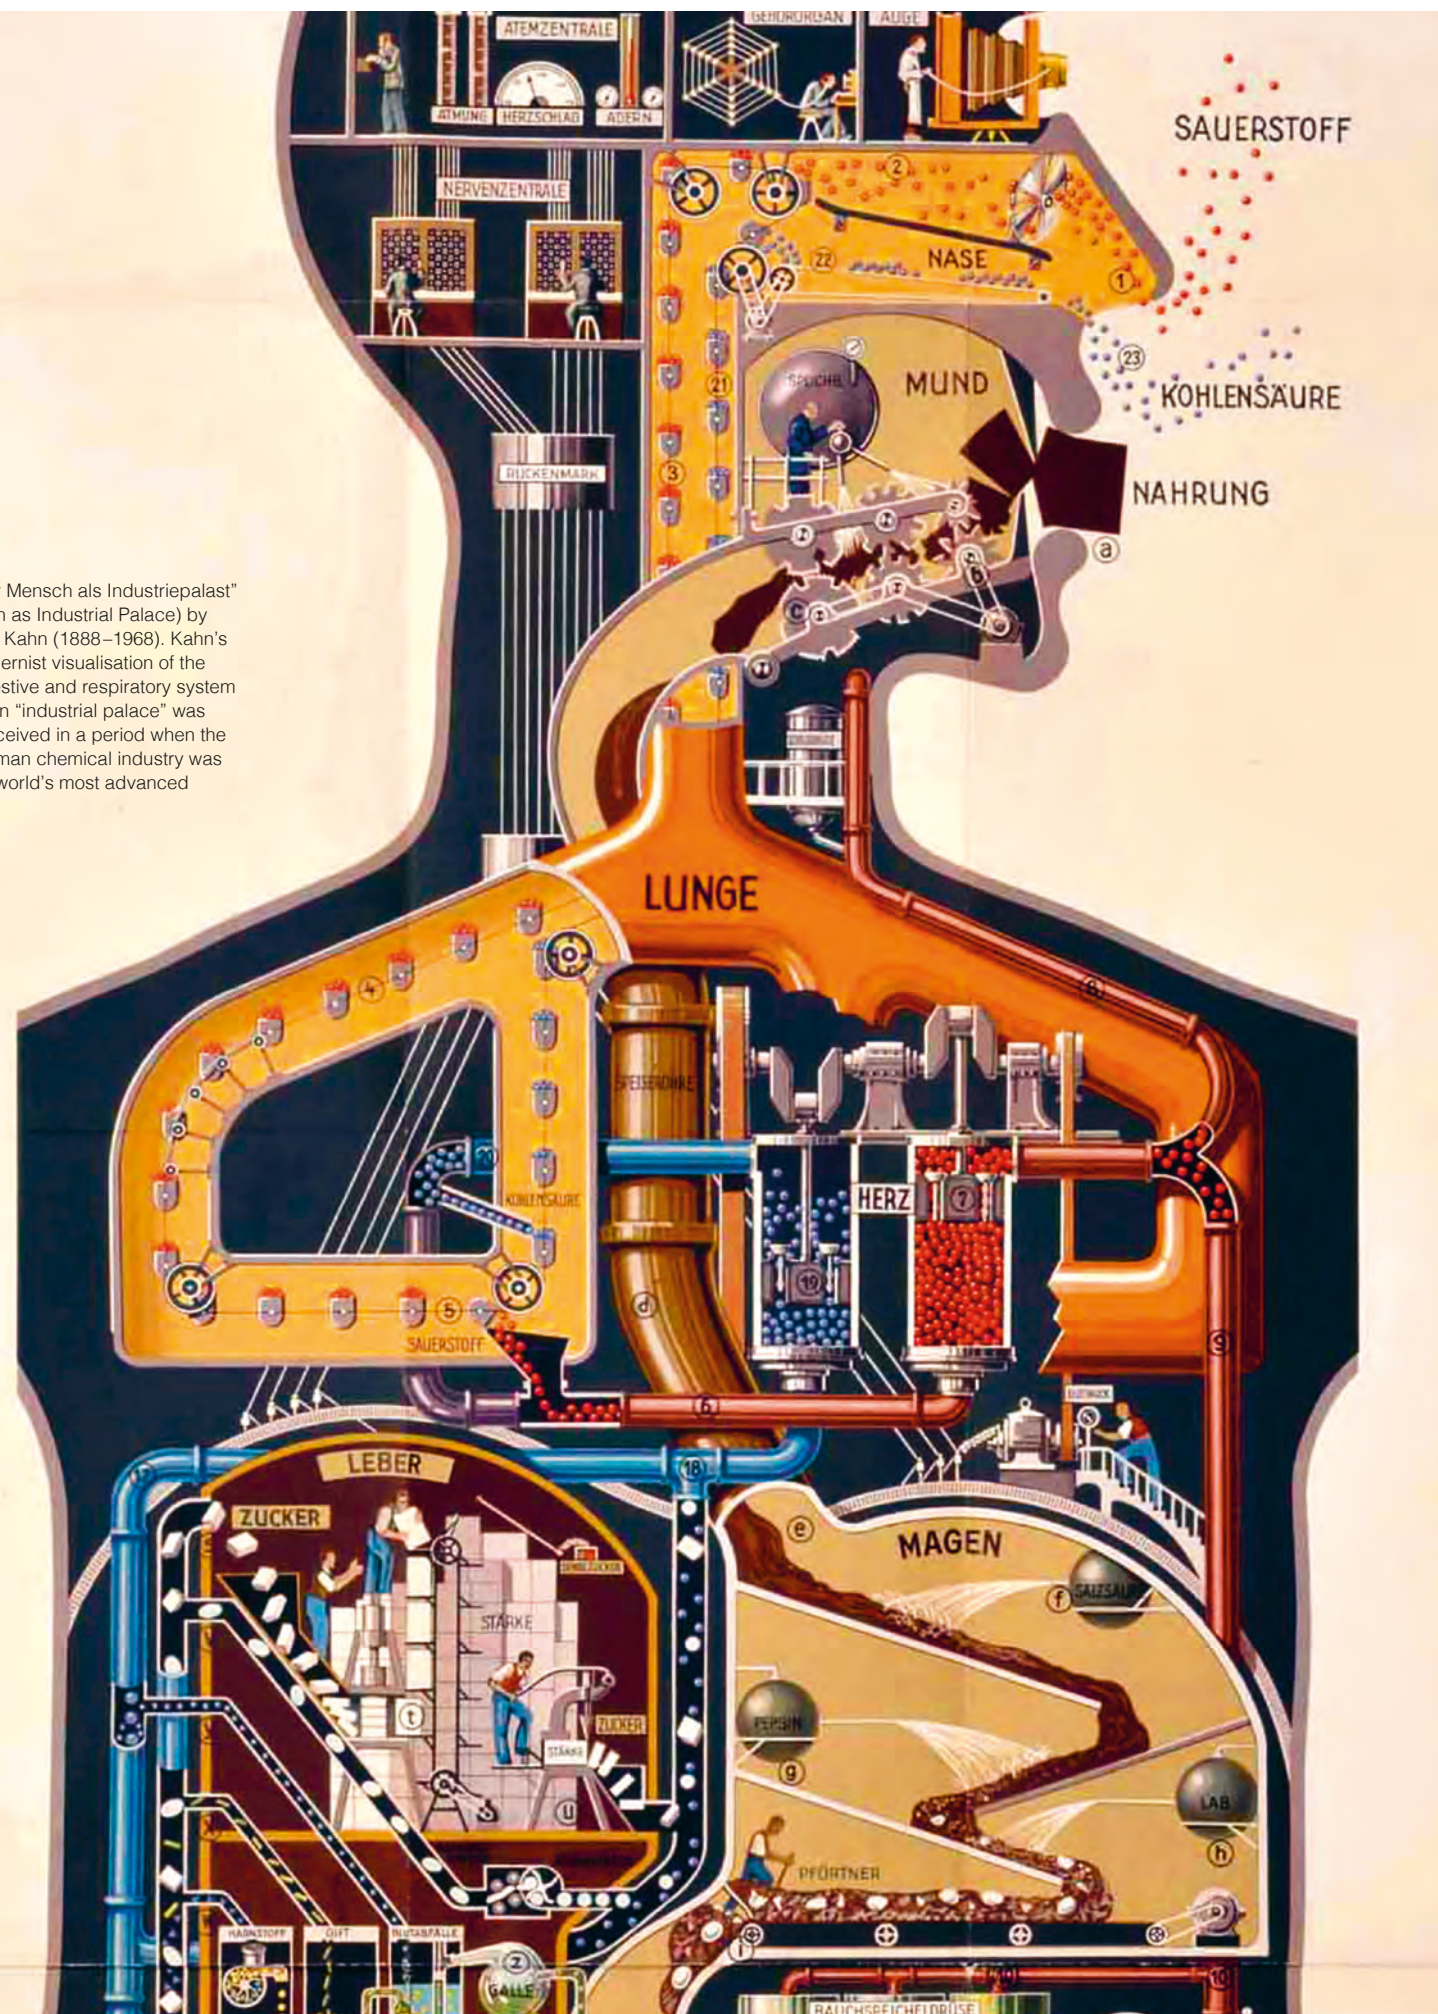
Take the current debate about “fake news”, and the search for truth, discussing what is a fact, and what isn't. Clearly the debate touches on our capacity for recognition and our logical capabilities. Is there an overlap between logic and intelligence?

I think that logical thinking is more of an aspect of intelligence. It's a refined form of symbolic thinking, clearly a key ingredient inherent in intelligence. But I think intelligence is a lot broader than real logic.

What is the difference then, the distinctive feature of the two?

“Intelligence is a difficult thing to define”

“Der Mensch als Industriepalast” (Man as Industrial Palace) by Fritz Kahn (1888–1968). Kahn's modernist visualisation of the digestive and respiratory system as an “industrial palace” was conceived in a period when the German chemical industry was the world's most advanced



Again, intelligence is a difficult thing to define. We can loosely say that it is intelligence that distinguishes us from other animals and what is responsible for the ascendancy we occupy on the planet at the moment. And that is something we can say with confidence without yet knowing what all the different ingredients to intelligence are. In a way, logic is merely a sort of abstraction from one aspect of symbolic reasoning processes.

Is one of the capabilities the ability to project scenarios in the future?

It does seem that one of the things that we humans are able to do is what one may call “scenario planning”. We are able to imagine what we think of as several options of a possible future. This is basically what happens when we are confronted with choices. These choices are related to such things like short time survival. Imagination is part of human thinking and we use it in making decisions all the time. It is an extension, an elaboration of that what we use in high-level scenario planning. It clearly must be something that we do with huge amounts of abstraction and processing information because we only have the capacity to deal at that sort of contrast level with a small amount of the information which is confronting us.

When we had a life expectancy of twenty years, our risk scenarios were limited. Most of them went into mythologies and religion. How important is our own perception of time in this regard? How would you say life expectancy and the ability of scenario planning are related?

I think the basic ability of imaginative thinking is associated with surviving and prospering over quite short time scales. For those basic sorts of cognitive abilities, the differences in our life spans don’t make much of a difference. Our ancestors were already surviving long enough for those skills to be relevant. There is a lot that could be said about the connections between that kind of activity, that sort of imaginative and predictive activity and our intuitive conception of time. One of the things we know from modern physics is that the intuitive conception of time is misleading. We have the sense that there is something like a flow or a passage of time and that time is intrinsically directed. We know from physics that this is wrong. It is an old project of

“A machine would have to be doing some scenario planning if they wanted to act like humans”

both philosophy and science to explore a distinction between those aspects of the world which are truly objective in the world itself, and those that in some sense come from us. We now know that a lot of those intuitive aspects of time really do come from us. They lie on the subjective side in the way that things like colour, taste and smell lie on the subjective side.

Do machines need to learn our capability of imagination to become more like us? Is that the crucial element?

A machine would have to be doing some scenario planning if they wanted to act like humans. But it is often not necessary to do that. A network of machines is in another sense just a single machine. I don’t think how we carve out the machines makes much difference one way or another.

What are we going to be seeing in the next ten years?

We are going to see a lot more people thinking about the long-term future of artificial intelligence; thinking about where this technology is taking us, where the opportunities are, where we can make a difference. In my view the big thing we need to do right now is to expand the community of people thinking about these issues. In particular, we need to find young people from many different fields, people who are going to spend their careers thinking about these issues; people who really will make a difference of how this transition into a machine era, in which we share the planet with non-biological intelligence, will turn out. And they will be the people to make a difference to how that goes.

So, the scenario is that we are going to be sharing our habitat with non-biological intelligence?

Yes, absolutely. We have a very unclear idea at this stage as to what the capabilities of those machines will be, but we can be fairly sure that most of the things we can do with our brain will be things machines will be able to do, too. And it may well be that they will be able to do things that we have not thought of yet.

Is that a utopian or dystopian idea to you?

I think that there are possibilities at both ends of the spectrum, both important to think about. We not only need to think about what to do on the

HUW PRICE
is Bertrand Russell Professor of Philosophy and a Fellow of Trinity College at the University of Cambridge. He is Academic Director of the Leverhulme Centre for the Future of Intelligence, and was founding Academic Director of the Centre for the Study of Existential Risk, which he established in 2012 with Martin Rees and Jaan Tallinn. Before moving to Cambridge, he was ARC Federation Fellow and Challis Professor of Philosophy at the University of Sydney, where from 2002–12 he was Director of the Centre for Time. Huw is also a Fellow of the British Academy.



safety side to avoid the dystopian possibilities, but also to clarify the range of possibilities towards the more utopian end of the spectrum. It may well be that there are importantly different paths the technology could take, which may be good in some ways and bad in others. It seems to be smart to determine a sense of destination before we set down this path. I like to say that there is an important difference between designing a self-driving car and the issue of the future of artificial intelligence. In the case of the self-driving car, you want something that will take you efficiently from A to B. In the case of artificial intelligence, in general we have no idea where it will be used, and more alarmingly, we have no idea what the possible destinations are.


Computers can already process a significantly larger amount of information than any human.

Exactly, so non-biological intelligence will have access to vast quantities of data, and that will be one of the things that enables it to do things we can’t do.

Are we also speaking about processing then applying the data to create new inventions, or will machines remain our assistants to help us innovate?

In order to be called intelligent the machine has to do something with the data. I don’t like that way of setting it up because it suggests we are at stage one where we have the data and stage two of artificial intelligence; solving the problem is just adding on the capability of doing something. In fact, we have lots of machines, which are capable of doing various things with data just as we ourselves are. Despite the fact that they have access to much more data, the current machines are incapable of doing many of the things that we can do with data. As time

“We have no idea what the possible destinations of artificial intelligence are”



Researchers at the NASA Advanced Supercomputing facility have produced this image, generated from a computer simulation of a contra-rotating, open-rotor design. Air particles are “released” on the upstream and aft blades. The basket-weave pattern shows where particles interact with each other – one of the sources of blade noise

goes on, it is likely that we will develop machines that will have much of the same general kind of abilities that we have. That will enable them to take lessons learnt from data in one case and apply it to another case. It will be that kind of generalisation that comes so easily to us that those machines will be able to do in the future.

Are you confident about this on principle or is there something in particular that gives you hope?

At this point, these technologies are turning out to be so useful for many purposes and commercially really valuable, not to mention a sort of scientific fascination for the topic. As to whether it is definitely going to happen: firstly, I’m not an expert and secondly, even the experts couldn’t say with certainty that it will happen in a certain time-frame. But my understanding is a reasonable middle-of-the-road viewpoint at this time. In principle, we see an obstacle but also no fundamental technological or scientific barrier that would prevent it from happening.

“Will machines ever be entities that we think of as having interests of their own?”

Do you think we will witness an exponential development in terms of a new industrial revolution?

Again, I want to emphasise that I’m not an expert in these fields, I’m a philosopher. Experts in the field think that we are probably several conceptual theoretical steps away from having machines with as wide a range of capabilities that we have. But, as with predictions in any scientific field, there is an element of guesswork there.

Speaking to you as a philosopher: what are the most fascinating questions deriving from this development for you?

I have spent my philosophical life on questions like the nature of time and the foundations of quantum theory. I want to be clear that there isn’t a lot of engagement with my professional life as a philosopher. My role in the centre is much more of an enabler or facilitator, someone who can just play a role in bringing other people together to make things happen. Having said that, I think some of the most philosophically interesting questions are about whether or

not the machines will ever be entities that we think of as having interests of their own. For many people, this is tied to whether at some point machines will be conscious – whatever that means. And there is a related set of questions about whether our own future as humans remains entirely on the biological side or whether at some point we have the option of perhaps enhancing ourselves so that we become hybrids, partly biological partly not. We would have access to a greater range of abilities as a result of that. For example, we might have immediate access to much more data. Then some people think there are possibilities where we become entirely non-biological. We upload ourselves into computers or something like that. So, there are a lot of fascinating long-term issues in that space and it may turn out that some of them will become true, in particular the issue about whether we want the machines to remain tools or instruments, something you can turn on or off without worrying about the moral status of the machine, the way you can decide with your vacuum cleaner. Some people take for granted that that’s the kind of future we want artificial intelligence to have. No matter how smart they are, they see them simply as tools. Others think that the natural path goes in the other direction and would therefore live in a world where machines are fellow moral agents. And it may turn out to have some implications for safety concerns as well.

The ethical implications are particularly interesting. How would you face the ethical challenges if we assume that a non-biological intelligence can be more than a mere air conditioner? Will we be still calling it a machine?

I don’t want to use the term “machine” in that sense, because I take it for granted that we are just machines. We are biological machines. Some machines are also people in the philosophical sense, entities with interests and moral agency. Whether the non-biological machines will ever be of that kind is going to be a choice that we face. It’s a choice we should face

deliberately rather than accidentally. As people have pointed out, one of the dystopian possibilities is that we create a future in which we don’t acknowledge the possible emotional capabilities of intelligent machines and thus create a dimension of suffering. That would be dystopian not for us but for the machines.

Going back to the conceptual steps, what is the point of no return where we lose control over our progress?

I think it is a one-way path, simply because of the commercial and other pressures on the development of technology. This leaves aside the possibility of some other large-scale calamity, which affects the level of our technology. It will be a once in a lifetime opportunity for the planet. We will not have the opportunity to back up and do it some other way.

Opponents of the developments in artificial intelligence argue that we should not play God. As a philosopher, what is your take on that?

I think we should be careful about playing God. But we should recognise that sometimes life confronts us with choices that we simply have to make. The choice of whether we want to make machines that are capable of consciousness or suffering is a choice we have to decide on whether we like it or not. So, a general disinterest in playing God doesn’t get us off the hook, because we will have to do it anyway.

“A general disinterest in playing God doesn’t get us off the hook, because we will have to do it anyway”

“Work saves us from boredom, vice and need”

Great Depression: during the severe worldwide economic downturn in the 1930s many people lost their jobs. The image shows a group of unemployed men in Hanover, Germany



Interview with Andrew McAfee

Work defines who we are. It is a huge part of our identity, MIT’s Andrew McAfee says. We therefore have to, sooner rather than later, discuss what defines us in a world with significantly less to do. Will we all be poets and painters? That would be beside the point. McAfee argues that we all need work in order to lead a healthy life.

ALEXANDER GÖRLACH: Last year there was a real breakthrough in robotics in Boston, where you work and do your research. How would you describe what our workforce is going to look like in the next decade or two?

ANDREW MCAFEE: Instead of any sharp discontinuity that would arise from robotics or AI, we will see a continuation and maybe even an acceleration of certain trends that are already quite clear. The hollowing out of the middle class is a real phenomenon, and there is plenty of evidence that middle-skill, middle-wage jobs are becoming less common. Certainly, there is reason to worry, especially for countries like the US or Germany, that developed a prosperous middle-class workforce in the second half of the 20th century. Our current election cycle in the US reflects this trend as well.

Would you associate the rise of populism on both sides of the Atlantic with this decay of the middle class? Some say it is the backlash of globalisation and digitalisation, do you agree with that?

To some extent, but it is obviously more complex than that. In case of America, there’s a

“It is always easier to blame foreigners for taking our jobs than to blame technology”



ANDREW MCAFEE is a principal research scientist at MIT and studies how digital technologies are changing business, the economy and society. His 2014 book on these topics, “The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies” (co-authored with Erik Brynjolfsson) was a New York Times bestseller and was shortlisted for the Financial Times/McKinsey Business Book of the Year award. Andrew was educated at Harvard and MIT, where he is the co-founder of the Institute’s Initiative on the Digital Economy.

part of society which feels left behind and “cheated by the system”, and that contributes to the rise of populism and demagogues. These people do not feel like the current changes in the economy and society work for them like they used to before. Hence, this “return to greatness” is so appealing to a large number of people.

I’d say this shift is mostly associated with the banks and Wall Street – nobody ever blames this on robots.

It is always easier to blame foreigners for taking our jobs than blaming technology.

There's a backlash against globalisation, but the backlash against technology is not as mature and developed yet. In Western countries, job creation happens at the high end for skilled and educated people. In America there are more and more jobs created at the low end of the job market – jobs with low pay and fewer benefits. The American social safety net is poorly developed and this leaves a lot of families struggling financially. The evidence is clear that job creation has moved from the middle class towards the bottom. The middle jobs were concerned with routine knowledge work or routine physical work and by now we have developed automation for both of these fields. Those jobs are not coming back, because software is cheaper and more reliable than the human workforce.

Are we compensating for those trends by creating new industries?
Yes, but it does not appear to be generating a lot of opportunities for the middle class. It is beyond doubt that there are new jobs being created that did not exist before – think about data scientists or social media consultants. But again, the problem is that these jobs are usually at the high end of the spectrum. On the other end of that spectrum, take the example of an Uber driver – he or she is doing a job that did not exist a decade ago. These days retaining a job in which you can set your own schedule and can make a living whilst not having to have a college degree is a big deal. And though Uber is not at fault for failing to create industrial middle-class jobs, let's not pretend we are seeing the same jobs and wages that we had before.

The real competition is with technology, and not with other drivers. One day, we will have self-driving cars ...
Yes, and Uber is investing heavily in self-driving cars. We are still generating new jobs for people, but I think that in this entrepreneurial society the focus of companies is not on creating new jobs. If people can build businesses in a scalable way, they will do so. When Instagram was bought for one billion dollars by Facebook, Instagram had fewer than 20 employees. We are building extremely impressive companies that generate huge value, but they do that with a lot fewer people than the giants of the industrial era did.

“The focus of companies is not on creating new jobs”

One observation I have made is, if we look at one specific field, that of media and AI, we can see that there is still some enthusiasm in the US to invest in that industry.
Absolutely. Media companies seem to be experimenting and exploring new ways how their content can be channelled to people – take BuzzFeed, Vice or Bleacher Report. Then again, we can see very wealthy tech entrepreneurs buying newspapers; take the example of Jeff Bezos and the Washington Post.

Netflix and Amazon, working with data aggregation, give me recommendations on the basis of my preferences and the preferences of other users. The same is true for Spotify and Soundcloud when it comes to music. Will, finally, journalism be the next industry to be revolutionised?
People have a great thirst for content. Entrepreneurs are trying to find new ways to deliver content in a way that customers find compelling, whilst still generating a profit. The heart of the industry is still content, and content automation will not happen any time soon. Even though we have algorithms to generate stories and news, I don't think finding the next big story will become an automated process.

How do we deal with big data in journalism?
It is a model. I doubt that everybody wants to consume data-heavy news. Take the current elections for example – whilst much of it is heavily data-driven, you still have pundits that talk on TV and write columns about their personal impressions. And whilst pundits have a historical track record of wrongly predicting the outcome of elections, there is still a heavy demand for their content. In fact, jobs that are most demanded in the economy are the qualitative ones. Entrepreneurial and management skills provide the new foundation for a successful career – and those are not the skills that are always associated with an IQ of 140 or higher. Again, I don't think we're heading into an employment apocalypse, where tomorrow we wake up and robots have stolen our jobs. I think that the trends are clear, and based on the progress of technology that we are experiencing this will only accelerate.

While a machine does your work, you can sit around. Will this trend lead to a change in our attitude towards the importance of

work? Or is this rather the end of a world as we knew it?
It is a big deal. Like Voltaire said: “Work saves us from three great evils: boredom, vice and need.” Out of those three, need is the easiest to compensate for. What should concern us beyond that is figuring out what a healthy, balanced life looks like when it is not dominated by a concept of work from the industrial age. The answers we have come up with so far are scaring me. Charles Murray observed that the top 20 per cent of America's white middle class has not changed in the past 50 years, with equal rates of marriage, family size etc. If you look at the bottom 20 per cent, however, 50 years ago their lives were very close to the upper 20 per cent. Since then, all these indicators have drifted off steadily to the point that divorce rates have risen, the number of one-parent, one-child households has risen. So have jail rates, drug use, mortality rates, you name it. In fact, this is the only demographic in which the mortality rate has risen, and the three causes for the rise have been cirrhosis, suicide, and alcohol/drug poisoning.

But then just how important is a job these days, regardless of the profession?
Hugely important. You need to work – it's a social status. It gives you an identity, it gives you dignity, it brings you in contact with other people. I look at this through both sociological and economical lenses, and what I see is that I personally have become a huge fan of work. Not because of inherent moral values, but rather because it provides a basis for a healthy life. And my fear of the disappearing middle class does not mean that I fear people starving – in fact we live in a very abundant society. But I see more bad than good things coming up in the future, be that marginalisation or discrimination.

Is there a need to tackle the things that are ahead of us? Not only in singularity, but do we need to rethink our philosophy towards work?
As I see it, we are moving towards this great detachment of having a very productive economy that does not need very much labour. In 50 years, the mines, the factories, the warehouses will be automated. Trucks will drive themselves, crops will pick themselves. We will have an extraordinarily automated economy

“There is a need for the presence of government intervention in markets”

that generates goods and services for us, and we will not need labour and this idea of “full employment” any more. I am aware that arguments just like that have been made for the past 200 years, and I am aware that until a few years ago they were unproven. But when I look at the technological progress we are making, we do need to confront your question – what does a good society look like in 50 years?

And we need to keep in mind that it's not just the lower-end jobs that are getting scarce ...
Absolutely. When jobs get replaced, workers do not move up the ladder. They have to go down, which puts pressure on the lower-end workers. This causes wages, as well as unemployment rates, to increase. I'm an advocate of the negative income tax, and this is one of the key reasons why. In America, we only do this to a very small extent, and I would like to see more of that. However, this only takes care of Voltaire's “need” part – we still need to tackle vice and boredom.

Your model led Germany through the last crisis – the idea that people would get money back when selling their cars earlier saved the entire domestic car industry.
Every time you tax or regulate something, you are interfering with the market. I support proper regulation, so there is a need for the presence of government intervention in markets. However – that's different than turning our backs on markets and having a Soviet-style “everybody gets an apartment” policy.

America is the leading force in the digital and AI age so far. Do we see an American tech renaissance happening?
Yes. People are busy everywhere and especially in Silicon Valley, where the future of our technology is being developed – people work tremendously hard. Though pay may not reflect the effort put in, people really do believe in their work.

This interview was previously published in Factory Magazine, a start-up magazine by the Factory Berlin.

“Machines can do some things better than we can”

Interview with Vinton G. Cerf

Heralded as “a father of the internet” Vint Cerf has truly championed his discipline. Looking back, he remembers a time when people spoke more to one another. Scepticism in regard to new technologies, however, has been around all the time. Education is for him the key to tackle the societal implications of machine learning and artificial intelligence.

ALEXANDER GÖRLACH: Did you, as one of the “fathers of the internet”, foresee all that has become possible through the means of data collection and machine learning when you started working on what later became the internet?

VINTON G. CERF: Certainly, I did not foresee everything, but Bob Kahn and I did do this design with global service in mind and with expansion in all dimensions (reach, speed and applications). The design was deliberately open to new protocols and to new means of carrying data packets (e.g. optical fibre came long after

the 1973 design). We also knew that the network would be a social medium, as we had seen that in the earlier ARPANET email lists. Smartphones proved to be a surprise, but we had already seen portable equipment (laptops and pads) and Alan Kay had described as early as 1968 his idea for the FLEX computer that was the proto-concept for the laptop. The system has scaled by factors of 1 to 10 million – not a bad record.

In the early days of the web there was a fear that governments could take over the internet. Now you see corporate power in the internet expanding: just a few companies such as Amazon, Facebook, Google or Apple are holding dominion over a tremendous amount of data and thereby influence over societies. Whom do we have to fear now: government or corporations?

Governments can still take over the net – shut-downs in Togo recently and in Egypt during the Arab Spring, filtering and blocking in China, are all manifestations of that. The private sector is still highly competitive. New competitors come along all the time (think: eBay, PayPal, Yahoo!, Google, Facebook, Twitter, Snapchat, Amazon, etc). I think there is a lot to worry about in the security space (think of the recent Equifax security breach, Russian/Chinese hacking, OPM, State Department and apparent hack of NSA “tools”) and that applies to private sector and government equally.

Surgeons performing a hysterectomy using the da Vinci robot. The system features a magnified 3D high-definition vision system and tiny wristed instruments that bend and rotate far greater than the human hand. The surgeons operate through just a few small incisions

50 years ago, the book “The New Industrial State” by Kenneth Galbraith had a similar claim. It stated that the most valuable companies at the time, due to their market value and their power over society, would supersede the government and its competences. The consequence would be a corporate state. This prediction did not come true. So, are today’s worries just the claims by Galbraith reloaded, or is the general context different this time?

No, governments will always have the upper hand. (They have fundamental power that private enterprise does not.) This does not mean that private enterprise is powerless. Big companies with a lot of revenue and assets are capable of doing some amazing things (Think Elon Musk, SpaceX and Tesla, for example).

Also, another book, “Future Shock”, dealt with the consequences of technological progress for societies. This book by Alvin Toffler did not so much emphasise control over society but rather the impact that progress may have on a populace once it accelerates beyond the understanding of the elites and the average Joe alike. Are

“Governments will always have the upper hand”

we in a period of time where progress has indeed reached such a pace that societies need to pause from it?

I think if we were living in 1901, we might be having the same conversation about the telephone, electricity, automobiles, radio etc.

After the last presidential election in the United States, a data collecting company claimed to have been responsible for the outcome of the election. Is it already possible, in your opinion, to influence society that thoroughly, through microtargeting, in such a way or was this company just bragging?

Mostly bragging, but there is clear evidence that micro-targeting works (think about advertising), so this is not a trivial matter or one to dismiss.

The developments in machine learning have raised questions about the future of work. Jobs in the hands of humans today may be taken over by algorithms and data science. This may not only affect manufacturing work, as automatisisation progress has in the past; it will also affect high-skilled labour in medicine, law-making and



VINTON G. CERF is vice president and Chief Internet Evangelist for Google. He contributes to global policy development and continued spread of the internet. Widely known as one of the “fathers of the internet,” Cerf is the co-designer of the TCP/IP protocols and the architecture of the internet. He has served in executive positions at MCI, the Corporation for National Research Initiatives and the Defense Advanced Research Projects Agency and on the faculty of Stanford University.


banking. Which scenarios do you think are credible in the short and medium term? I think the long term has many jobs going away and many new ones created – the challenge is to re-educate people whose jobs many have been automated so they can do the new ones.

Another consequence of this development concerns the distribution of wealth. For the first time in history men may not have to work physically, manually until exhaustion, yet the question arises of what man will do, what will his or her work be, and how will he or she be compensated? There is a growing disparity in wealth distribution and I think that has societal risks. Whether we get to Star Trek’s 24th Century with no money is still an open question. Money has utility as a medium of exchange – it is fungible – and that’s very useful. A transition to such a state is the subject of a lot of speculative thought.

What do you personally think the future of work will be? Certainly, we will not all become painters or poets. I think people will continue to look for productive things to do, compensated or not. We do need basics: food, shelter, clothing, and meaningful ways to spend our time. I think that the kind of work we do will absolutely change as technology continues to evolve.

When you look back to your childhood and adolescence and recall the society you grew up in: what are the main differences when it comes to values and social norms back then and now? We talked to each other more back then. Privacy seemed less threatened. We were less vulnerable to social stress that seems to come from the online social media of today. On the other hand, information was harder to find – now we have it at our fingertips. We had to work a bit more to be entertained (not a bad thing). In some ways, I think we used our leisure time better – today’s world seems more broken up into smaller pieces.

To what extent do you think artificial intelligence and the developments in society that come with it, may change humanity’s perspective on itself, how may we answer the famous question by Immanuel Kant “Was ist der Mensch?” I think we will begin to question what it means to be intelligent – we will find that “machines” can do some things better than we can. Ultimately, I hope that artificial intelligence/machine learning will become tools that allow us to work more effectively as opposed to being competitors. I suspect we will also learn to appreciate non-human biological intelligence more – at least I hope so.



Data: Frey CB & Osbourne MA. The Future of Employment: How susceptible are jobs to computerisation? Technol. Forecast. Soc. Change, 114:254–280, 2017



Oxford University academics examined 702 common occupations and found that some jobs – telemarketers, tax preparers and sports referees – are at more risk than others including recreational therapists, dentists and physicians

One answer might be computer games

An Essay by Yuval Noah Harari

Most jobs that exist today might disappear within decades. As artificial intelligence outperforms humans in more and more tasks, it will replace humans in more and more jobs. Many new professions are likely to appear: virtual-world designers, for example. But such professions will probably require more creativity and flexibility, and it is unclear whether 40-year-old unemployed taxi drivers or insurance agents will be able to reinvent themselves as virtual-world designers. (Try to imagine a virtual world created by an insurance agent!) And even if the ex-insurance agent somehow makes the transition into a virtual-world designer, the pace of progress is such that within another decade he might have to reinvent himself yet again.

The crucial problem isn't creating new jobs. The crucial problem is creating new jobs that humans perform better than algorithms. Consequently, by 2050 a new class of people might emerge – the useless class. People who are not just unemployed, but unemployable.

The same technology that renders humans useless might also make it feasible to feed and support the unemployable masses through some scheme of universal basic income. The real problem will then be to keep the masses occupied and content. People must engage in purposeful activities, or they go crazy. So, what will the useless class do all day?

One answer might be computer games. Economically redundant people might spend increasing amounts of time within 3D virtual reality worlds, which would provide them with far more excitement and emotional engagement than the “real world” outside. This, in fact, is a very old solution. For thousands of years, billions of people have found meaning in playing virtual reality games. In the past, we have called these virtual reality games “religions”.

What is a religion if not a big virtual reality game played by millions of people together? Religions such as Islam and Christianity invent imaginary laws, such as “don't eat pork”, “repeat the same prayers a set number of times each day”, “don't have sex with somebody from your own gender” and so forth. These laws exist only in the human imagination. No natural law requires the repetition of magical formulas, and no natural law forbids homosexuality or eating pork. Muslims and Christians go through life trying to gain points in their favourite virtual reality game. If you pray every day, you get points. If you forget to pray, you lose points. If by the end of your life you gain enough points, then after you die you go to the next level of the game (aka heaven).

As religions show us, the virtual reality need not be encased inside an isolated box. Rather, it can be superimposed on the physical reality. In the past this was done with the human imagination and with sacred books, and in the 21st century it can be done with smartphones.



YUVAL NOAH HARARI received his PhD from the University of Oxford in 2002, and is a professor at the Department of History at the Hebrew University of Jerusalem. His current research focuses on macro-historical questions such as: what is the relationship between history and biology? Published in 2014, Yuval’s book “Sapiens: A Brief History of Humankind” has become an international hit and is published in nearly 40 languages worldwide. “Homo Deus: A Brief History of Tomorrow”, a critically acclaimed book that examines the big future projects facing humanity in the 21st century, hit the shelves in 2016.

Some time ago I went with my six-year-old nephew Matan to hunt for Pokémon. As we walked down the street, Matan kept looking at his smartphone, which enabled him to spot Pokémon all around us. I didn’t see any Pokémon at all, because I didn’t carry a smartphone. Then we saw two other kids on the street who were hunting the same Pokémon, and we almost got into a fight with them. It struck me how similar the situation was to the conflict between Jews and Muslims about the holy city of Jerusalem. When you look at the objective reality of Jerusalem, all you see are stones and buildings. There is no holiness anywhere. But when you look through the medium of smart books (such as the Bible and the Qur’an), you see holy places and angels everywhere.

The idea of finding meaning in life by playing virtual reality games is of course common not just to religions, but also to secular ideologies and lifestyles. Consumerism too is a virtual reality game. You gain points by acquiring new cars, buying expensive brands and taking vacations abroad, and if you have more points than everybody else, you tell yourself you won the game.

You might object that people really enjoy their cars and vacations. That’s certainly true. But the religious really enjoy praying and performing ceremonies, and my nephew really enjoys hunting Pokémon. In the end, the real action always takes place inside the human brain. Does it matter whether the neurons are stimulated by observing pixels on a computer screen, by looking outside the windows of a Caribbean resort or by seeing heaven in our mind’s eyes? In all cases, the meaning we ascribe to what we see is generated by our own minds. It is not really “out there”. To the best of our scientific knowledge, human life has no meaning. The meaning of life is always a fictional story created by us humans.

In his ground-breaking essay, Deep Play: Notes on the Balinese Cockfight (1973), the anthropologist Clifford Geertz describes how on the island of Bali, people spent much time and money betting on cockfights. The betting and the fights involved elaborate rituals, and the outcomes had substantial impact on the social, economic and political standing of both players and spectators.

“Hence virtual realities are likely to be key to providing meaning to the useless class of the post-work world. Maybe these virtual realities will be generated inside computers”

The cockfights were so important to the Balinese that when the Indonesian government declared the practice illegal, people ignored the law and risked arrest and hefty fines. For the Balinese, cockfights were “deep play” – a made-up game that is invested with so much meaning that it becomes reality. A Balinese anthropologist could arguably have written similar essays on football in Argentina or Judaism in Israel.

Indeed, one particularly interesting section of Israeli society provides a unique laboratory for how to live a contented life in a post-work world. In Israel, a significant percentage of ultra-orthodox Jewish men never work. They spend their entire lives studying holy scriptures and performing religion rituals. They and their families don’t starve to death partly because the wives often work, and partly because the government provides them with generous subsidies. Though they usually live in poverty, government support means that they never lack for the basic necessities of life.

That’s universal basic income in action. Though they are poor and never work, in survey after survey these ultra-orthodox Jewish men report higher levels of life-satisfaction than any other section of Israeli society. In global

surveys of life satisfaction, Israel is almost always at the very top, thanks in part to the contribution of these unemployed deep players.

You don’t need to go all the way to Israel to see the world of post-work. If you have at home a teenage son who likes computer games, you can conduct your own experiment. Provide him with a minimum subsidy of Coke and pizza, and then remove all demands for work and all parental supervision. The likely outcome is that he will remain in his room for days, glued to the screen. He won’t do any homework or housework, will skip school, skip meals and even skip showers and sleep. Yet he is unlikely to suffer from boredom or a sense of purposelessness. At least not in the short term.

Hence virtual realities are likely to be key to providing meaning to the useless class of the post-work world. Maybe these virtual realities will be generated inside computers. Maybe they will be generated outside computers in the shape of new religions and ideologies. Maybe it will be a combination of the two. The possibilities are endless, and nobody knows for sure what kind of deep plays will engage us in 2050.

In any case, the end of work will not necessarily mean the end of meaning, because meaning is generated by imagining rather than by working. Work is essential for meaning only according to some ideologies and lifestyles. Eighteenth-century English country squires, present-day ultra-orthodox Jews and children in all cultures and eras have found a lot of interest and meaning in life even without working. People in 2050 will probably be able to play deeper games and to construct more complex virtual worlds than in any previous time in history.

But what about truth? What about reality? Do we really want to live in a world in which billions of people are immersed in fantasies, pursuing make-believe goals and obeying imaginary laws? Well, like it or not, that’s the world we have been living in for thousands of years already.

The article was previously published in The Guardian.

Epilogue

by Alexander Görlach

Looking into the answers that experts of various fields gave us for this publication, artificial intelligence and its implications for society are all but a Terminator-like scenario. On the contrary, speaking in science fiction terms may disguise the real challenges and therefore delay asking the right questions.

Undoubtedly, a major shift is at hand: it is an undeniable fact that machines are capable of processing and storing information on an infinitely larger scale than humans can. For the experts here, this is anything but a problem. As a matter of fact, they perceive embracing this fact as the first step in coming to terms with this new reality. Were machines not always designed to do their task better than their human makers? Why should it be different this time?

When it comes to the scenario of machines not only executing what they were designed for but also developing a life of their own, all experts say the same: this scenario is not likely ever to happen. If we are to go down the road of smarter machines, it will at least be a long way to go.

Experts make it abundantly clear that further pursuing artificial intelligence would not mean that machines will develop emotions and mistakes comparable to those of men. To the contrary: their infinite reservoir of data may make them see problems and solutions faster and therefore differently from the way their human makers can see and comprehend them. The ideology of singularity describes nothing different: artificial intelligence will understand what we humans will need before we even know this. Already today Google knows when a flu outbreak is under way, simply because people start looking for medication or the address of the nearest pharmacy.

This is not simple data aggregation alone. The algorithms themselves get smarter in reading and understanding the data available. They not only process and store, but also interpret. The measures they apply so far are only quantitative. Qualitative understanding, however, is on its way – a semantic comprehension. It is not quite certain that machines will prevail, for algorithms are also open to biases in the search they conduct. Biases, as we know from humans, are the source of errors. One crucial question, therefore, will be how to “insert” critical ethical comprehension into the processing of algorithms.

Take self-driving cars for example: when will an automated car brake, and when will it not? A human driver has to make a decision in a millisecond, based only on what he or she sees. Will the algorithm ultimately be blind to gender, age and race, and brake at all costs? But would this not, at times, endanger the passengers in the car or the bus? Ethical weighing is a crucial thing to be deployed. Without it, there is no larger application of artificial intelligence outside the confinement of a smart phone or a laptop.

Self-driving cars very soon will work in a surrounding they know; for instance, the streets of Palo Alto or Cupertino. Driving to San Francisco, just 40 miles away, may be impossible for a longer period. The lack of data that the car would need to navigate unknown territory is, at this

point, immense. As we just pointed out: it is not only data regarding navigation of streets, it is about data navigating the ethical implications of moving a vehicle in human surroundings.

The same goes for machine learning in medicine or law-making. At this point, machines will be better at creating an overview of all cases accessible online. Their suggestions in the end, however, should be informed not only by quantitative but also by qualitative criteria. No expert would deny the possibility of this happening. Yet the timeframe is rather blurry.

This leads to a stout geo-political implication of artificial intelligence: what ethical framework should apply? A Western one? And if so, from which tradition? Or should China run the show? Or even ethics deriving from Islam? What about proposing a profound atheism in artificial intelligence? Would artificial intelligence take into account human and civil rights and in consequence always be “good”? There is no reason whatsoever to believe that artificial intelligence will support us in outsourcing these crucial questions from our own consideration. An artificial intelligence God, as some fantasise, in the sense and the frame of science fiction movies, is totally out of the question. Since humans have an inclination toward the inexplicable, the utopian, frame it however you wish, they may find a certain inclination in themselves to subordinate themselves to what they perceive as an unquestionable “God-artificial intelligence-machine”.

This is not the only political implication of artificial intelligence: all experts express their concern to some degree or another when it comes to the impact of artificial intelligence on jobs and distribution of wealth in society. They have a belief that either new jobs will be created in the process, or that some sorts of jobs will never be replaced by robots or other sorts of intelligent machines. Since we attach our self-esteem and narrative about ourselves strongly to the work we do, the implications of artificial

“All experts express their concern to some degree or another when it comes to the impact of Artificial Intelligence on jobs and distribution of wealth in society”

intelligence will not only concern the question of how we make money to live by, but also how we gain and maintain or self-esteem in a world with less and less work.

When machines and robots create more and more wealth in societies, governments will have to ensure that the companies owning these technologies are subject to redistribution. We already see a disparity both globally and within the Western world in the accumulation and distribution of wealth. All experts agree that the artificial intelligence revolution has all the ingredients to amplify this development.

What is next?

by Alice Deißner



ALICE DEIßNER neé Steinbrück is Director of Strategies and Programmes at the Vodafone Institute. She is responsible for the strategic development of the institute's programmes as well as heading up research and policy work in our various areas of work.

Just imagine robots that can perform flawless backflips, recognise people and address them personally, or artificial intelligence that analyses the human genome and suggests targeted treatment. These and many other impressive examples of recent technological advances make it abundantly clear that what until recently would have been considered figments of a science fiction author's imagination is now reality. The future has arrived. But are we ready for it? Ready for technology that makes the unimaginable possible? Looking back at the mistakes and economic upheaval during the last industrial revolution in the 19th century, as well as the economic and social changes that came with it, it is clear we would be well advised to plan ahead for the digital future.

The Vodafone Institute wants to explore the potential associated with future technologies for an equitable and responsible society. Since it is clear that artificial intelligence can help improve the quality of life for many people on this planet, we asked distinguished experts and thought leaders in science, philosophy and industry for their views on the societal and ethical implications of this technology. As an action-oriented think tank, we now intend to define where action is required and by whom.

Although their opinions and perspectives on artificial intelligence differed greatly, all of our interviewees agreed on one crucial point: putting the right ethical and legal framework in place is a highly political issue and a tremendous challenge for governments and regulators around the world. A set of rules and standards must be defined to control the impacts of new

technologies on our societies. Political and technological determinism would certainly not be an appropriate mental orientation given the risks and opportunities ahead. Rather than stumbling blindly into the future, we have to establish the rules of the game. In order to guarantee that all members of society benefit from our digital future, the deployment of artificial intelligence must be regulated, transparent and fair, while leaving sufficient freedom for innovation and growth.

After having thoroughly analysed the interviews, the Vodafone Institute has decided that its future research will focus in more depth on the following aspects:

Sophisticated algorithms are becoming increasingly smart, and they are making more precise and accurate predictions than humans could ever make. Although they are supposed to help us eliminate the human subjectivity factor from the decision-making process, algorithms often end up perpetuating preexisting biases. If algorithms function as "black boxes" these biases may remain unchecked, quietly altering how the world operates. Engineers generally follow the instructions of their employers or customers, so it is hard to hold them responsible for the consequences of the tools they create. However, when a company decides to engineer something that the general public would not agree with the issue of ethics comes into play. In order to move away from a reaction-oriented prevention strategy, several measures have already been proposed for private-sector enterprises, such as the employment of ethicists working alongside development teams, or the provision of ethics training to computer science

students. This also raises the question of how governmental regulations can be implemented to guarantee fair and transparent algorithms, which has become a matter of considerable debate. In several countries – also at EU level – the debate surrounding algorithm watch initiatives has begun to attract the attention of political decision makers and regulators. The Vodafone Institute intends to consider these issues in more depth and facilitate dialogue on specific steps to be taken in politics and the private sector.

A second aspect that the Vodafone Institute will specifically focus on is the dystopian fear that artificial intelligence induces in the general public. It extends from robots making them redundant, to intelligent machines turning against humans, surveillance systems running amok and robots ultimately wiping out humanity. The fact that artificial intelligence can actually improve society and serve humanity is often ignored. This fear and distrust of artificial intelligence – and future technologies more broadly – could in fact become the new model of digital disruption. Issues such as security and protection will hence play a pivotal role in gaining people's trust and protecting consumers in the digital economy. Here, again, gaining trust will depend on the introduction of political regulations, on the one hand, and credible private sector initiatives, on the other.

These are, however, only two out of many aspects that require more research and thorough discussion. The Vodafone Institute will continue to explore challenging issues, provide action recommendations and, most importantly, encourage debate between academics, government and the private sector.

Imprint

Published by
Vodafone Institute for Society and Communications
Behrenstraße 18
10117 Berlin
www.vodafone-institut.de

Executive Committee Members
Inger Paus, Thomas Holtmanns

Responsible for Content/Editorial Management
Friedrich Pohl (Head of Communications), Alice Deißner (Director Strategy and Programmes)

Publishing House
Tempus Corporate GmbH –
Ein Unternehmen des ZEIT Verlags
Alt-Moabit 94
10559 Berlin
www.tempuscorporate.zeitverlag.de

General Management
Jan Hawerkamp

Project Management
Dr Joachim Schüring

Art Direction
Barbara Krimm, Christopher Delaney

Picture Editor
Beatrice Jansen

Proofreading, Corrections
Tom Harris, Lisa Wicklund

Production
Dirk Woschei

Print
Mediadruckwerk, Hamburg

Date of Publication
February 2018

Pictures:
Cover: Getty Images/Science Photo Library/PASIEKA
p. 3: Peter Rigaud; p. 4: Getty Image/Hulton Archive;
p. 5: Getty Images/Bloomberg Finance LP; p. 8: David Elmes;
p. 12-13: Boston Dynamics; p. 14: Soeren Stache dpa/lbn;
p. 15: Luciano Rezzolla; p. 16: akg-images/Africa Media Online/Ahmed Jallanzo; p. 19: akg-images/Stocktrek Images; p. 20: Getty Images/Fabrice Coffrini/AFP; p.23: Pierre Le Leannec; p. 24: IBM; p. 26: Getty Images/Forrest J. Ackerman Collection/Corbis; p. 28: Ian Scott; p. 30: Heinz Troll/EPO; p. 32: Nuria Oliver/Arduino Vannucchi; p. 34: akg-images/Universal Images Group; p. 37: Huw Price; p. 38: NASA; p. 40: akg-images/Walter Ballhause; p. 41: Andrew McAfee; p. 44: Getty Images/BSIP/UIG; p.47: Getty Images/David S. Holloway; p. 50: Ramy Zinger; p. 54: Amin Akhtar



FSC
www.fsc.org
MIX
Papier aus verantwortungsvollen Quellen
FSC® C107139



ClimatePartner

klimaneutral

Druck | ID 53533-1801-1012



Vodafone Institute
for Society and
Communications