Can we make the Digital World ethical?

A report exploring the dark side
of the Internet of Things and Big Data

by Peter Warren, Michael Streeter and Jane Whyatt

NETOPIA
Forum for the Digital Society

Foreword

The digital revolution connects people. Or so we're told. Our assumption about the Internet and digital technology is that it is about people communicating with people. Anyone with a laptop and broadband can make services for a potential global audience. Or so the story goes. The conflicts online involve real people: trolls on forums, privacy issues, grooming, hackers and man-made viruses. Or that's what we like to think.

But what if the digital revolution is about machines communicating with other machines? What if algorithms, software bots and smart devices make the most traffic? The stock market report in the newspaper you read with your morning coffee is not made by an editor; rather, it is, but that editor is a bot, compiling data from stock market servers. Most of the trades on that stock market are done by machines. The ads on that newspaper's web page (and all other web pages) are published by algorithms, and bought by bots on micro-second ad exchanges. When you tweet your thoughts on the latest market trends, that tweet is read, analysed, retweeted and stored by bots (often with human-looking account names). And you haven't even finished that morning coffee yet.

Within five years, a majority of online traffic will be machine-generated. Humans will be in the minority in terms of connectivity. That is the complete opposite of the way we think about the internet today and it raises many questions: can machines be accountable for mistakes? Small mistakes, sure, but what about medical treatment bots, self-driving trucks, or automated weapon systems? All of those technologies exist today. The trends of Cloud computing, big data and smart devices accelerate this development, and machines increasingly make decisions without human involvement.

Done right, this could be a blessing: scores of bots making your life easier. But the past six months' revelations about privacy abuse by both government and private organisations suggest that it's much more complicated. The legal consequences of this technology must be addressed, yesterday. Most importantly, what ethical considerations go into these systems?

ITGP has published this paper on behalf of Netopia to address these questions. It is intended as a starting point for a conversation, rather than a final answer. I hope you will find it interesting, intriguing and inspiring, as I have.

Brussels, February 13th 2014

Per Strömbäck
Editor Netopia

Contents

Introduction
Section One – What is the Internet of Things?
Section Two: How the IoT will work in practice
Section Three: The dark side of The Internet of Things
Section Four– Putting ethics (and better code) into the machine
Conclusion

Introduction

When the history of the 21st century is written it will be known as the Age of the Machines. Over the coming decades, machines will be intimately involved in every part of our lives – from the cradle to the grave – and playing a key role in the poorest of states as well as the wealthiest countries.

We and our machines will generate an ever-increasing flow of digital data, and the machines themselves will add a torrent of material that they generate independently.

According to the telecoms giant Cisco[1], by 2020 machine-generated data will exceed the traffic generated by people. These inanimate objects will fine-tune themselves and collect information from us and the billions of sensors built into what has become known as the Internet of Things (IoT). This is set to become a reality by around 2015.

According to those helping to bring it into existence, the Internet of Things will have the potential to connect virtually anything to the web or wireless networks, from packs of African hunting dogs in Botswana, to lamp shades on dining tables in Bristol and coffee-making machines in Dijon.

This Internet of Things will link computers, cars, mobile phones, clothes, fridges, food, fields, plants, planes and people. Nothing will escape measurement, because everything means something according to the theory of the IoT.

The same will be as true for us as it will be for our machines. In the same way that our health will be monitored so, too, our cars will be watched to ensure that they are performing at peak efficiency both for our wallets and for the environment. We will hear a lot more of the word 'smart': smart means finely-tuned, online and efficient, hence smart phones, smart grid, smart city, smart house, smart car, smart pants.

The Internet of Things will create a world where everything is measured and we, too, are measured in relation to the huge pool of 'big data' that the machines and sensors gather. Indeed, 'big data' will be the biggest product of the Internet of Things.

First, the good news. The IoT could be one of the greatest boons to humankind in history, enabling unparalleled understanding of our lives and our relationship to our planet. Using smart technology, big data and the IoT we should be able to improve our lives and reduce our impact on the world simply by making life more efficient.

The bad news is that there could be – some would say there is already is – a downside. As suggested by the recent revelations from the former NSA contractor Edward Snowden[2], the IoT has profound implications for us in terms of surveillance, privacy and consumer rights. As consumers we are at risk of becoming simply one component of the IoT: a component at the mercy of the sensors in the street and the analytical software engines and algorithms in the machine.

It will not stop there. As this mass of IoT-generated data becomes greater, computer systems will require more and more autonomy to allow them to reach conclusions about us. Some of these programs may even become self-programming, giving them even greater power. Machines will really start to govern our lives[3].

It would be a major mistake, though, to think of these machines as human-like 'androids' encased in an outer skin. On the contrary, many of the machines with which we will cohabit in the coming decades will simply be pieces of software moving freely but unseen through the internet as 'software entities'. Their influence will be subtle, unspectacular and often unseen, but no less profound for that.

The key question will arise: will we humans remain in control of the process, or will the process begin to control us? There is concern about our growing inability to keep up with the pace, scale and implications of technological change. In that sense, are we already beginning to lose control of the machines?

Arising from this key question are fundamental issues that should concern us all, no matter how what our attitude is to technology. In this mass of digital communications, machines and big data, where do human rights such as privacy stand? Where is the human dimension in a world dominated by machines? Can or should a system of ethics be imposed on computer software and the Internet of Things itself?

Are existing legal frameworks and approaches able to adapt to the coming machine age? Should machines themselves be accorded some form of 'rights' in order to better protect we humans? If so, what kind of body should be responsible for assigning those rights and inputting controls into the machine world? What rules will govern the makers of the machines and the 'Lords of the Clouds'?

One of the most important issues of all, we believe, is that we are entering a major age of technological change with little real public discussion of the implications that this will have on all our lives.

In compiling this report we have canvassed the opinions of a number of leading experts in the fields of computing, robotics, philosophy, the law, the Internet and cyber security. As we shall see, many experts feel it is indeed time that the public in general and politicians in particular woke up to the technological changes that we as a society are starting to encounter.

This is emphatically not a plea to slow down the pace of technological change, nor some modern-day Luddite anti-technological argument. As we say, the Internet of Things and all that goes with it could bring exceptional benefits to the human race.

Instead, this report is intended to help further the debate on just where we as citizens and consumers stand in relation to these changes. It is also important to remember that, in the midst of the complex and often exciting technical changes that are being developed to help the human condition, we should ensure that humanity itself is not left out of the equation.

Section One – What is the Internet of Things?

Put simply, the Internet of Things[4] (IoT) is a collection of sensors attached to objects – any kind of object – that will form an enormous data collection system. These sensors will connect to the outside world – usually the Internet – wirelessly using Radio-Frequency Identification (RFID) technology or a number of other radio technologies, or by SMS.

Already there are some 10 billion devices wirelessly connected, with research[5] suggesting that by 2020 that number will be 30 billion devices. The range of objects that will have sensors is limited only by our imagination; obvious examples include computers, cars, mobile phones, clothes, fridges, food, fields, plants, planes and people. Other less obvious ones include meat cooking in ovens, balls used in sport to prevent their being lost, pets, lampposts and keys.

The technology was pioneered on the International Space Station where every object is given an IP address which is known to small spherical robots that follow the astronauts around as they conduct their work. This allows the astronaut to know instantly where to find objects that they may need and cuts down on the time lost trying to find tools for a task. The lists of things that they need can even be sent up to the astronauts from Earth.

The idea of an RFID-connected world was set out at the Massachusetts Institute of Technology (MIT) in the late 90s, though the term 'Internet of Things' is usually credited to an English researcher, Kevin Ashton, who noted in June 2009 the limitations of a people-driven Internet. 'The problem is, people have limited time, attention and accuracy - all of which means they are not very good at capturing data about things in the real world,' he wrote.

'Ideas and information are important, but things matter much more. Yet today's information technology is so dependent on data originated by people that our computers know more about ideas than things. If we had computers that knew everything there was to know about things - using data they gathered without any help from us - we would be able to track and count everything, and greatly reduce waste, loss and cost.

‘We would know when things needed replacing, repairing or recalling, and whether they were fresh or past their best. The Internet of Things has the potential to change the world, just as the Internet did. Maybe even more so.’[6]

Indeed, the IoT can capture real-time data from anything from a jumbo jet – which, in a single journey across the Atlantic Ocean, can generate 640 terabytes of data from its four engines – to a wind turbine or a chair. In each case, the information collected will depend on what you want to measure and therefore which sensor you use and what you program the device to collect. Even something as simple as a chair, for example, can provide a whole range of data: on its location, when someone last sat on it, who is sitting on it now, what clothes they are wearing, the material, its weight, age, wear, size, owners, history, designer, manufacturer, place of purchase. The list is endless: in deciding what data is important, we create the IoT[7].

Each of these internet-enabled information-gathering objects could have a web page – even our clothes – and thus a virtual web identity that a huge range of other internet-enabled machines will be able to interrogate. We will need to consider whether we actually own our things and what rights to information ownership will confer. As Ashton rightly points out, however, they are all machines that owe their role and their *raison d’être* to humanity, even if humanity may not be the only entity that has an interest in them.

For the purposes of this report, we must engage with Ashton’s idea of a world of information and what that means in reality. Generating such a mass of data and its manipulation by machines is already presenting a significant challenge to humanity.

Big data

The key ‘product’ of the Internet of Things will be the data it generates. For the purposes of the current discussion, the most important aspect of this will be the gathering of what is called ‘big data[8]’. The technology research company Gartner has defined the term this way: ‘Big data is high volume, high velocity, and/or high variety information assets that require new forms of processing to enable enhanced decision-making, insight discovery and process optimization.’ As is so often seen in the world of technology and technology writers, this description is weighed down with jargon.

In simple terms, big data is the result of modern technology’s ability to gather and store lots of facts – data – quickly, efficiently and securely, and then analyse it in a way that makes the world more efficient. Or that, at least, is the theory. Big data is said to have three main attributes – volume, velocity and variety – or the three Vs. Some definitions now add a fourth V – ‘veracity’. Time could also be added as a final factor, making the term ‘three Vs and T’ or ‘four Vs and T’.

Examples of big data include the 39.5 million tracking requests from customers of United Parcel Service per day, the 172,800,000 card transactions processed by VISA every day and the 500 million tweets sent a day[9]. Such mountains of data are meaningless, however, without the ability and equipment to quickly extract meaning from them. Computer power, analytical software and – most importantly – our needs all have to be deployed so humans can create the algorithms necessary to extract meaning from the data. If used properly, this can yield some stunning results. For example, the US power company GE’s research into big data concluded that in aviation, a 1% reduction in fuel consumption could result in $30 billion in savings over 15 years. Meanwhile, a 1% improvement in the efficiency of gas-fired power plants could produce $66 billion in fuel savings globally.

Meet the machines behind the Internet of Things

SENSORS
CONTROL ROBOTS
PERSONAL COMPUTERS
MOBILE DEVICES
SERVERS
MAINFRAME COMPUTERS
ROBOTS
ALGORITHMS
SOFTWARE ROBOTS

As can be seen, today's machines are not the android-style robots of science fiction. It is perhaps precisely because of their highly technical, anonymous nature that there has been so little general discussion of their creeping impact on our lives.

Sensors

Sensors are the most basic form of computational organism. Like the tentacles of an undersea crustacean or sea-anemone, they are programmed to respond to the environment around them and in response to specific stimuli. Modern sensors have now become so sophisticated that they can survive conditions totally hostile to biological mechanisms, can travel to planets and be subjected to incredible pressures, temperatures and stresses. Some are designed for use in space craft, others are so incredibly small that they are imperceptible to the human eye and have been dubbed 'smart dust'. Sensors have now been developed that can exist at the nano-scale, others that – while made of metal – are thinner than cling-film, as flexible as skin and yet virtually indestructible.

More common sensors are available that can be injected into most mammals – a list that includes people, dogs and mice.

These sensors can work in an active or passive mode. This means they are either executing the task that they have been meant to perform – like the Mars rover – or simply sitting in 'sleeper' mode waiting to execute a task or transmit information upon receipt of instructions.

Some of the most classic examples of this are the RFID devices familiar in clothing shops or ticket-less mass transport systems, such as underground trains. These are turned on by directing a radio wave at them or by passing into a radio field. Larger devices can either contain their own power source or have technology built into them that allows them to generate their own power. Even if that fails, they can still be activated by a radio beam.

When coupled with the potential for light, heat, movement and chemical sensors, and anything that their programmers care to build into them, the Internet of Things carries a huge potential for information gathering.

Sensors will be placed out in the wild to monitor our environment and the condition of the systems that we depend upon. They will also be used to run our offices and homes. Via the smart-grid, we will be able to see for the first time exactly what resources our houses are consuming and be able to remotely control them from our mobile phones. The same is true of offices, factories and supply lines. Many industrial processes are already controlled remotely, often off-site by engineers working from home or in distant offices. In an increasing number of cases, that control will be relinquished to machines themselves.

The Internet of Things will in effect become a set of nerve endings for the web, which will be able to detect the presence of an object of interest by the movement of the web – and then home in on it.

Control robots

Dr David Levy, author of *Love and Sex with Robots: The Evolution of Human-robot Relationships*, views control robots as the lowest form of machine life. They are machines built to execute one purpose, utilising code that is designed to allow them to perform only that function.

Control robots include a family of important control systems known as programmable logic controllers (PLCs). These are the most basic form of computational control. They sit immediately on top of sensor systems and are set to respond quickly and 'authoritatively' to the information that is sent to them.

Limited in their computational power, PLCs are deliberately restricted fail-safe devices designed to operate in often extreme conditions to control things like nuclear reactors and carry out a restricted set of functions. They are often deployed in critical services such as the utilities, water, gas,

electricity, communications and other key areas of the infrastructure.

Personal computers

These are now the grandfather clocks of the computer world. The PC has become the computational heirloom in the modern house. Often unused – but often left on – the PC is still nonetheless a rich aggregated source of historical data on the modern family and shares with the domestic router the distinction of usually being the main route onto the web.

In the near future home computers will be still used to provide an access and control point. This function will be increasingly turned over to the Cloud as remote 'thin client' tablet-based computing replaces the desktop PC as the web interface of choice.

Mobile devices and avatars

It is in this area that the Internet of Things promises an explosion of connectivity. From the ubiquitous mobile phone, to the car, satellite navigation, the home and lifts; an incredible number of mobile devices will be deployed in our lives. We will become used to deploying our own sensors in mobile situations to monitor anything from potential rodent infestations to security and health applications, all of which will be connected to our mobile phones.

Some people will opt for household robots to help in those situations; others will prefer a more fluid internet- and sensor-based approach.

In the world of the Internet of Things sensors, as we have seen, will be everywhere, and our main way of accessing the data from these sensors will be mobile devices.

Indeed, as we are already seeing, a seemingly endless array of applications is turning the mobile phone into a digital Swiss Army knife. The addition of biometric technologies linked to body, location, behavioural pattern recognition and health monitoring systems will allow for robust security via the phone. On the basis of this research, the authors think that the mobile phone will in effect become the first mass robotic device and will become the focal point for consumer calls to legally protect machines.

A key development in this trend will seem like the beginnings of artificial intelligence on our phone – avatars. This development has been predicted for over a decade now, and like many of the other technologies mooted at the same time, it is just beginning to come to fruition. These avatars will act as our gateway into the Internet of Things and to a large extent become our guides through it.

These avatars could either be models of us, or some figure we have chosen that appears on our mobile phone. Initially they will look like a character in a computer game or in a video format, but eventually they will become 3D figures transmitted from the phone itself and will act as a reporting mechanism on the apps deployed on our mobiles.

The primary factor governing their adoption will be perceived usefulness. Like the mobile phone itself, adoption will be driven by the usefulness of an avatar that manages mundane tasks for us. Trustworthiness will be a secondary – though still important – factor. As we are now seeing with the fallout from the Edward Snowden affair, people are aware of their personal interests and are beginning to recognise that the privacy of their information is a right, and should be protected as such.

The key to these avatars will be in their effectiveness and their loyalty, a fact not missed by the eight technology companies which have stated that surveillance must end[10]. Their response is not altruistic but a response to commercial pressures: people and companies want their data to be private. Increasingly, they will insist that their interests are looked after by avatars and that those avatars stay loyal to them. Like their physical possessions, avatars will belong to a single, individual person. People will insist that the device and the avatar are extensions of their life. This will set a paradox between the different levels of ownership of a thing, but people will insist that if they have bought it, that they have rights of primacy over what it does. They will also insist that they have control over it. Once this occurs they will personalise the avatar.

In effect, these avatars will be acting as our personal assistants, telling us about our finances, giving us important information that we have to respond to and managing parts of our lives. It is possible that they will eventually be sold to us as a service by large

companies who will undertake to use our data and invest our funds.

More and more, given the complexity of the system that is developing, we will begin to see these programs as personal protective devices and as the most important piece of technology that we own.

According to Dr Jonathan Cave, Senior Research Fellow at RAND Europe, we are already seeing the emergence of companies offering 'identity as a service', and 'privacy as a service'. We believe it is inevitable that these will eventually be incorporated into services provided by the avatars created to look after our interests in the complex new world of the IoT. These avatars will examine the terms and conditions of the agreements relating to the software that we download. They will also alert us to the device settings that relate to use and ask for our instructions on how they should be changed to protect our security and privacy.

This is not science fiction, but a development that is now only a few years away. Though it may alarm some, Professor Murray Shanahan of London's Imperial College would welcome such a move. 'I think an avatar-based artificial intelligence is very much the way that we will start to see something that looks more like artificial intelligence of the sort we imagined when we were boys and girls. I think that is the way it's probably going to come,' he says.

'The unification of the apps is something that is very much lacking in the technology at the moment as the left hand of your smart phone doesn't know what the right hand is doing most of the time, and the moment that you start to integrate things and personalise them it's going to start to look a lot more like a little alien intelligence in your hand. I have to say that that is something that I find an exciting prospect. I don't find that a scary prospect.'

Servers

The central nervous systems of the new world order, data centre servers operate as the clearing house for all of this IoT-generated data as machines fire off messages at close to the speed of light.

The servers currently operate separately, according to the needs of the companies that have deployed them in their data centres. But in the world's largest data centre in Las Vegas, run by Switch, no such boundaries exist over the data according to Jason Mendenhall, the company's head of Cloud operations. Switch inherited a technology pioneered by the disgraced US energy company Enron that allows the companies using the data centre to empty their data into a communal pool where it can be 'interrogated' to make use of the data patterns. Once the data has been made anonymous – stripped of names and addresses, for example – the patterns derived from it can be applied, says Mendenhall, to socio-economic groups, businesses and other situations.

While this may appear an attractive proposition for some, many civil liberties groups are critical. They point out that it is a relatively trivial technological task to use other databases to identify individuals from within this pool of 'anonymised' data and that when DNA databases are added to the mix, anonymisation will be little more than a myth. Industry appears to share these doubts about how such data can remain truly 'anonymous'. The authors have spoken to sources at a company running a loyalty card system for a global supermarket chain. They confirm that it is possible to accurately mine through anonymised data and then work out how individuals in a street will vote based upon their purchases. This is data that is now offered for sale to political parties.

Mainframe computers

Like the PC, these are now 'antiquated', the steam engines of the modern computer world. They remain the repositories of massive amounts of information, yet their current role has not been defined. Rumours of their imminent demise have circulated for more than 25 years. Many think that their final role is as the super-computer systems used by intelligence agencies and government computer centres where the data can be protected. Plans exist for the consolidation of government data – in other words pooling it – along the same lines as those being developed in the Switch data centres as discussed above.

Robots

The popular machines of science-fiction, the idea of automated figures and self-operating machines dates back to the Ancient Chinese, Greeks and Egyptians. The modern notion of a human-like android that works for us owes its roots to the Czech playwright Carel Capek, who coined the word robot in his 1920 play *R.U.R – Rossum's Universal Robots*. At the start of the

21st century, the human-like robot is finally nearing reality in Japan, where demographic changes mean that there will soon be five elderly people to each young adult. The Japanese Government has thus embarked on a programme of robotics to generate the carers that will be required to look after this elderly population. Japan has also developed all of the other software that robot technology needs to support those more recognisable robots.

Algorithms

The younger of the software machines, algorithms are already used by businesses, governments and intelligence agencies to mine large amounts of data for information. While this may concern many, the use of such algorithms does not alarm Professor Murray Shanahan, an expert on cognitive robotics. 'A lot of artificial intelligence technology is very passive – it does not have plans or intentions, it's not sitting there doing anything sinister, it's not thinking about doing anything but simply passively extracting data,' he says. 'It's a use of AI technology that I would prefer to having hundreds of human operators sitting there listening to what was going on.' This does, however, implicitly raise concerns over who it is that defines the algorithms' parameters.

Software robots/'bots'

There is an ongoing dispute about the difference between a software robot and an algorithm; some experts argue that the software robot is simply a more sophisticated version of an algorithm.

If an algorithm is a software machine created to search for specific information, a software robot or 'bot' – of which the computer viruses Flame or Stuxnet[11] would be good examples – is created to carry out numerous tasks and to have a degree of autonomy. According to Professor Neil Barrett, author of 'The State of the Cybernation' and a former high level adviser to the UK Government, the UK Home Office and the EU, the next generation of software robots is expected to involve a degree of self-programming or decision making, based on the situations the robots will encounter via the Internet.

Section Two: How the IoT will work in practice

The arrival of the Internet of Things will, in time, have a major impact on our everyday lives. One important area that will be affected is health. Let us see how this new world might work in practice by considering how a typical heart patient – we'll call him Charles Rabbit – will be treated in the future.

Charles is in his mid-60s and has been under treatment for a heart condition for a few years. His father had a similar problem but Charles's treatment is altogether different. While Charles's father had to go into hospital or see his GP for regular monitoring, Charles's monitoring will be done at home. The sensors woven into Charles's clothes will be continually monitoring his condition and sending data back to his doctor via his avatar. This is a smart program based in the computing Cloud that Charles bought to look after his interests. It usually takes the form of a 3D talking figure – a visual avatar – that can appear on phones or on web pages to receive instructions. In Charles's case, he has opted for a complex machine endlessly performing calculations that he calls 'the Difference Engine'.

Using a combination of pre-programmed information on what Charles likes plus information the Difference Engine has stored on previous supermarket purchases, the avatar knows the sort of food that Charles favours, and it uses that information to check ingredients against a list of menus supplied by the doctor.

By cross-referencing through both lists and the contents of the Internet-enabled fridge, the avatar is able to produce a selection of evening meals for Charles to choose from. The fridge, using an RFID scanner, notes when ingredients are used and adds items to a list that will be bought from the supermarket.

The supermarket itself will also be monitoring Charles's consumption habits. It will offer incentives on certain items based on information it has asked the Difference Engine to supply, and on information the supermarket has culled from its loyalty card scheme.

After the meal, sensors in the house record Charles's activity. This data will include which rooms he has visited, whether he has taken his medicine, how much alcohol he has drunk and indeed whether he is moving, lying down or standing.

In winter, this detailed information enables the Difference Engine to achieve significant cost savings on Charles's fuel bills. Knowing which areas Charles habitually uses, it preheats certain rooms and reduces heating in others.

The Difference Engine combines this knowledge with information taken from weather forecasts that allow it to achieve the most comfortable blend of heating and clothing for Charles. Information of this kind is also cross-referenced with data about other patients with a similar demographic and condition.

The Difference Engine also sends information about Charles to his daughter, Ada, so that she knows her father is all right. Ada had, in fact, requested permission to 'virtually accompany' her father at all times but this was declined by the Difference Engine because Charles had not wanted it. He felt it would be too intrusive.

The Difference Engine had, however, sent information to Ada and to the doctor when Charles had not taken his medicine for two days in a row, and allowed them into the house when they called to find out why.

Meanwhile, the continuous monitoring of Charles's health data is being carried out by centralised computer systems. These systems map small changes in his condition and alert the system to changes in Charles's body that need to be corrected. This information is sent to the Difference Engine to implement; the only role for the doctor now is to intervene to ensure co-operation.

Charles is happy with the system that he has helped create. But inevitably there are downsides. One is privacy: although Charles stopped his daughter having full access to the Difference Engine, he is worried that others in the system may be 'snooping' on him.

Furthermore, as his monitoring involves a certain amount of expense, a large health company has underwritten the cost in return for access to Charles's data. The company has good reason to do this: Charles's condition is considered to be particularly interesting and as a result the data is potentially lucrative.

Charles is not well-off and so the offer of help to ensure his well-being is welcome. But at the same time he and

his daughter Ada are concerned about the company's use of Charles's data. Because the data is totally specific to him, they want to know whether he has copyright over it. Is the father's data worth more than the company is paying him? So far, their lawyer has simply told them that this is a grey area.

Another thing that troubles Charles is what happens if something goes wrong in his care and treatment. Who will be held responsible? The makers of the Difference Engine? The health company sponsoring him? The health service? The doctor? Himself? So far, neither Charles nor Ada has been able to get a clear answer from anyone on this.

Section Three: The dark side of the Internet of Things

At the moment, we are like individual bees. Bees become powerful when they are connected and form a beehive, becoming capable of incredible things. I think that is what is going to happen with pervasive computing. We will go from being individual bees to being part of a hive, forming a meta-intelligence.

The Internet of Things will create a situation in which, for the first time, we will live in a world mapped by machines, a world that we will increasingly inhabit outside our physical bodies. This world of data has huge implications for all of us. If not monitored properly, it will have a profound and potentially disastrous impact on our civil liberties and potentially rob us of our capacity for self-determination.

We call this the dark side of the Internet of Things. While there has been increasing media discussion about the very real benefits that this new era of technology will bring, we argue there has been worryingly little debate about the potentially harmful effects it could have. So, let's have a look at some of the areas of potential concern.

1. Privacy and integrity
2. Misleading data analysis – when the sums do not add up
3. Acting in the right interest

Privacy and integrity

The Internet of Things means that the opportunities for massive surveillance and intrusion into our lives by governments and large organisations will reach ever-greater proportions. It will potentially become an 'Age of Data Surveillance' or 'dataveillance'.

It is important to note that the data gathered about us may not come from information that we have chosen to give to governments or financial organisations. It may well come from information that we have made public elsewhere in the digital world. For example, at a recent meeting organised by the big data analysis company Splunk – which has built a database 'mining' tool that can run real-time analysis of Twitter at the same time as incorporating other databases – it was noted that attempts by the German government in 1980s to increase its census

Dr Adrian Cheok, Professor of Pervasive Computing, London City University

data had met significant popular resistance and had to be shelved. Yet the German government is now able to gather information about citizens that is significantly more detailed than that sought by the proposed census questions, simply by analysing social networks and other databases.

According to many of the experts we have spoken to, it is already possible to take information from such sources. Highly detailed pictures can be created of behaviour, location, health, finance, buying patterns, driving habits and internet searches. This ability to 'map' us will only increase as the Internet of Things grows in scale.

It is worrying how little consumers and citizens realise the importance of the data trail we leave in the digital world. As far as the algorithms and software machines are concerned, we are our data – warts, inaccuracies and all – and it will be virtually impossible to hide from our data.

Beware the spy in your fridge

Should we be worried by the risk of household appliances such as fridges being linked to the Internet?

Quite simply: yes.

In November 2013 Hacker News[12] revealed that devices secretly fitted with remote sensors had been found in Russia - the source for these was given as China.

While the aim of the exercise is unclear, one theory is that it could have been an attempt to introduce a radio spying system capable of logging onto internal communication systems. This would provide a way of gaining access to an important network containing vital information.

Such subterfuge is not new. USB sticks still in their shrink-wrapped packaging have been found to contain Trojan computer programs, and Microsoft computers made in China have been found to contain malware systems.

Fred H. Cate, Distinguished Professor at the Indiana University Maurer School of Law (photo: Ann Schertz)

On one device given to the authors - a USB stick bought in China by a top engine designer for a UK car company - programs stored on the stick had attempted to take engine blueprints from the company and hide them out of sight of the Windows operating system.

Given that there are already many documented cases of similar attempts to steal data it is inevitable that the Internet of Things will become a target. This will be due to the speed with which new sensors are introduced, the lack of thought given to their security, and their potential to be reverse-engineered as spying devices.

Smart meters are one such system[13]. They not only give a unique insight into who lives in a house through their power use, but also allow anyone else logging onto them a detailed picture of what is being done in a home by monitoring power usage against particular rooms.

The data processing industry is able to overlay large amounts of data from a variety of sources, which will then throw up what's known as anomalous data. This is a very effective way of finding people who are trying to hide, such as fraudsters. Unfortunately, it can also highlight people who, for perfectly legitimate personal reasons, are seeking to keep something private, such as their sexual orientation.

This level of sophisticated analysis will become even easier as the Internet of Things gets bigger. Mobile phone and locational information will reveal location and patterns of behaviour that can be cross-referenced against other people with known behaviours, such as sexual orientation – in other words, the data will 'out' people.

This is particularly worrying for people who may have a genuine need to hide. This can include former spouses, people evading organised criminals and many other individuals who have a legitimate reason to escape detection. Already many cases exist of people using social media networks to stalk people. The authors have been told by police in the UK of a worrying upsurge in cyber-stalking that involves the manipulation of social media data. We have also been told that during an investigation into a drug smuggling gang operating in the south-east of England that the gang had used mobile phone scanners to discover that it had been infiltrated. It then used hackers to try to break into the phone company's database to discover the identities of the informants using those phones.

Professor Viktor Mayer-Schönberger, Professor of Internet Governance and Regulation at Oxford University, notes how in the United States a technology company has bought the entire US offender list and published it online. Many would argue that those people who have already served their sentences should be given another chance. As the law stands, however, that data can be mined to show where they are living, having a detrimental effect on their employment opportunities and efforts to rehabilitate themselves.

These former offenders can pay a sum of money to be removed from the list, which, according to Mayer-Schönberger, amounts to blackmail on an already economically-challenged group. In any case, removing themselves from such a list may not give them much protection in the future. Using big data, analysts can easily discover economically-inactive individuals, people who do not appear on electoral rolls and so on.

These individuals can be mapped against the typical pattern of an incarcerated criminal, and the resulting data fed to credit companies and others offering financial services. In this way, even those who avoid leaving a digital trail can be identified by commercial interests. Our data-driven profiles will not just record what we do, but also what we do not do. *Not* tripping a sensor will be something that software robots will pick up on and use to make decisions about us.

Thus the combination of credit reference software robots and other AI robots working on behalf of other institutions will need to be scrutinised – and questions raised about how they use data. Professor Fred Cate says a key issue is how you can control and protect that data. Both Cate and Mayer-Schönberger argue there has to be a legal lifespan built into the data.

Misleading data analysis – when the big data sums do not add up

None of this might matter if we could be sure – really sure – that such data gathering was always in our individual interests. But, of course, we cannot be sure. Already, large corporations, governments and intelligence agencies are using algorithms to 'interrogate' big data for usable patterns, trends and information. Companies and governments look for large groups of people to target for a particular service they want to offer, while intelligence agencies sift the same information to find smaller groups of people of interest to their law and order goals.

Curiously, the intelligence agencies and financial institutions often work in similar ways to profile their 'targets'. Dr Bjoern Rupp, Chief Executive Officer of GSMK Cryptophone, an expert on encryption and phone interception who has worked as an advisor to the German Government, says that so-called 'data robots' work their way through government-intercepted data and a mix of other databases, some obtained from the commercial sector, to identify patterns.

The algorithms are fed the financial and communications records of a known terrorist or criminal organisation and then used to analyse the huge data pool that has been collected. 'The current data hides a lot of information inside so you can easily determine not just who called who, but who was travelling from where and when and how many people that they were in contact with,' says Rupp.

Adding further information from the internet makes the systems' potential profiling ability even more powerful. 'From the data you can employ advanced data mining technology and then find out, for instance, not only who the person is but you can then profile that person to find similar people to them in the database. You can effectively say to *[the algorithm]* I'm trying to find a certain person and the system will generate that person for me even though I don't know them yet,' says Rupp.

Such group profiling is made easier for the intelligence agencies because it is also information that financial institutions look for. Algorithms used by the financial services sector now seek to discover groups of friends and associations between particular groups in their records. These can then be overlaid with data from social media groups to fine-tune marketing activities. In the case of the banks, the algorithms seek out information showing groups of debit or credit cards used at the same time and place. This information indicates that the potential marketing targets are taking part in a group activity, such as swimming or football, or watching sports events.

According to Rupp, data robots working for intelligence agencies will also focus on such data clusters. In their case, however, they will try to determine whether, for example, participation in a five-a-side football team is not simply cover for a more sinister purpose.

A specific concern will be the way that insurance companies use big data, perhaps to minimise their exposure to risk. To put it another way, insurance companies could use big data to prevent having to pay out on claims by effectively disenfranchising entire groups.

An example concerns the insurance industry in Britain and its attitude towards people who live in houses built on flood plains. Due to their concerns about fears of human-induced climate change and the high incidence of flooding and damage to houses built on flood plains, the insurance industry is threatening to withdraw cover from such homes. This means that anyone wanting to buy such a home would be unable to get a mortgage. Those already owning them will not be able to sell them.

The insurance industry, using statistics culled from the emerging Internet of Things – patterns of climate change, previous flooding incidents, likely flooding models and sensors placed in all of the major rivers recording river flow and flood water patterns, and which are showing year-on-year increases – is demanding that the UK Government covers it for providing cover.

Thus, without government intervention, those people living in flood plains will become an unprotected group. The insurance industry has sought to strengthen its argument by pointing out that projected patterns suggest that areas which would not have been considered high risk areas for flooding could one day be under threat.

The same scenario could be developed for people suffering from obesity or other similar conditions. Big data, as Cheok says, can produce a wonderful world, but not when you are on the receiving end. 'We have to consider that there are already cases of ID theft and people running up a bill for you without your knowing.

But when you are extending more and more of your intellect onto the internet what we are going to see is that there is going to be a battle between the people in society who genuinely want to make good use of this data, as in the case of healthcare, and those who don't.

'It is genuinely useful if you can be connected to your doctor for 24 hours of the day and they can see all of the information about your body, but on the other hand health insurance companies, if they have that data, can eliminate from their policies anyone who the remotest chance of getting sick in the next ten years or so. So there is a bad side to having so much data on the internet,' says Cheok.

Minority Report myth

There is growing concern among experts about the current approach to analysing big data. The root of this concern, according to Professor Mayer-Schönberger, stems from a misunderstanding of big data. Mayer-Schönberger says it is being used wrongly to predict people's patterns of behaviour instead of being seen as a record of what has happened.

'The problem is that as human beings we want to see the world as a series of causes and effects and therefore we are tempted to abuse big data analysis – which can only tell us *what* is going on – to know *why* this is going on, so that we can then connect guilt and individual responsibility to individuals. This is precisely some of the abuse that we see coming out of the NSA and Prism debates,' says Mayer-Schönberger.

Indeed, many politicians have fallen prey to the temptation to see big data as a universal panacea, according to Mayer-Schönberger.

'As we have seen not only is [big data analysis] being used in the US to prevent terrorist attacks, but it is also being used to go after petty crime by the FBI and local police forces. Then you have a very powerful tool that cannot tell you anything about individual responsibility – it only tells you "what", not "why" – and it is being used for the purpose of assigning individual responsibility and causality,' says Mayer-Schönberger. While at present there is continuing uncertainty about exactly how one derives information from data, the point is that – given the speed with which the new world of data and the attendant IoT is developing – it makes sense to err on the side of caution before using data in this way.

Mayer-Schönberger says there is a risk in falling for the myth depicted in the 2002 Steven Spielberg film *Minority Report*, in which the police apprehend criminals before they commit a crime. 'Minority Report has a very strong rosy premise and that is to avoid having victims,' says Mayer-Schönberger. 'The problem with it is that we don't let fate play out and we don't know whether a person would have committed a crime; we make an assumption that they will because every prediction based on big data is probabilistic.

'With big data, there is a risk of predictive social control and a system of social control which slaughters human volition at the altar of collective fear,' he argues.

Another issue surrounding big data is the extent to which it is 'anonymised'; in other words, removing the information that would identify an individual. Though the data processing industry claims that by anonymising they are not using an individual's data against their wishes, this is challenged by a wide range of experts including Mayer-Schönberger. The critics suggest one can 'reverse out' individuals from the data and identify them. There are some who claim that 'anonymisation' is one of the data processing industry's greatest lies.

Acting in the right interest

Consumer rights

A major concern for our rights as consumers is the way that machines direct us according to their interests and not ours. Experts such as Dr Jonathan Cave warn about the growing influence software machines have over our lives. Cave says that software machines will make use of what they know about us to present information which may not be to our advantage. Because the search engines that we have used know a certain amount about us and our previous buying decisions, they are keen to exploit that by turning us into a buyer of something, by a process known as 'filter bubbles' – a feedback loop where recommendations only reinforce existing patterns.

As Dr Rupp states, 'if you are not paying then you are not the customer'. Thus, if you use a 'free' internet service, such as Google or Facebook, it is not acting in your interests; rather, it is acting in the interests of the customers who *are* paying to present information to you. 'As technology changes, our concept of what our rights should be may change: it could be that privacy empowers me to use my data, or it could be that it becomes a market opportunity for someone to collect my data,' Cave says.

To take a small but familiar example, when Google picks up that you are interested in – say, lawnmowers – it fields a number of adverts down the side of your search that relate to your search and then fine-tunes that list according to data that it holds on us. Research has shown that people do not often go beyond the first search page, so companies pay for search engine optimisation (SEO) – i.e. they buy in experts to make sure that they are in the top three results in a search and they keep on refining that to maintain that position. Thus, the search is not in our interests - it is in fact a series of adverts competing for our attention.

Moreover, other research has shown that if a web page does not load in eight to 11 seconds then we will go to another site. In this way, the system is already directing us and to this extent we are at the mercy of the machines. Other information about us is stored, based upon our profile and information that Google or other systems have culled about us, and later used to serve up offers it thinks may be relevant.

We know of one mid-level executive who was very embarrassed when, while doing a web search in front of colleagues, he saw adverts being served up about

Viktor Mayer-Schönberger, Professor of Internet Governance and Regulation at Oxford University

Björn Rupp, CEO of GSMK Cryptophone

Caribbean cruises for gay people. These had evidently resulted from previous searches that he had made. Of course, the machine did not know that at certain moments this information could be embarrassing and was trying to be helpful. The individual had not realised the implications of logging into his personal Google profile.

There is clearly a conflict of interest in using a search engine to search for information while that same search engine system is selling our personal interests to other companies seeking to establish a commercial advantage.

The situation will become even more worrying with the emergence of avatars, which as we have seen are predicted to become our 'personal representatives' in the world of the Internet of Things. Software developers will home in on our avatars' code and seek to develop their own code that will make decisions for the avatar. Companies and their software engineers will try to force data out of the avatar that we may have instructed our avatar to withhold. There will be a need to protect the avatar and the ways that it interacts with other software.

Yet another concern is the impact on behavioural science – the study of how we actually make decisions and behave. The Internet of Things will have a big impact on this, as it will yield huge amounts of data on

how we actually behave, where we are and what we are doing.

It is important to note that this information will not all come from the online world, but from our offline behavior, too. Supermarkets, for example, will be able tell which objects we stood in front of in a shopping aisle and which products made it into our trolley at which time. This can then be cross-referenced with other decisions we have made.

'Losing control of our technology'

Dr Michael Anderson, Associate Professor of Computer Science at the University of Hartford and an expert on machine ethics, says this debate raises an important question: who are the machines working for? By 'machines' in this context we do not mean 'dumb' robots, but instead software entities in the Internet that are engaging with a human search or action. It is these programs that cause the greatest concern to Professor Murray Shanahan, an artificial intelligence specialist, who has a particular interest in the risks associated with AI intelligence.

'I think before we should be worrying about humanoid robots taking over there is some concern about artificial intelligence that is not embodied in quite that way but which is in the devices that we carry around with us, and in the internet, and in the cloud and so on,' he says.

'We are going to see a lot more AI technology embedded in our surroundings and in the internet and in systems that are connected to the internet – and that's where I think that we need to worry. Not so much about being taken over in some sort of science fiction scenario – but of losing control of our technology.'

One reason we may lose control is simply that of complexity: the machines and the systems they run have become so complex that no-one understands the mechanism any longer, and it could already be developing its own momentum.

Eric Schmidt, the former Google CEO and advisor to US President Obama, has pointed out: 'The Internet is the first thing that humanity has built that humanity doesn't understand--the largest anarchy that we have ever had.'

Ethical machines?

This question of control is a key area that is often overlooked in discussions of the Internet of Things. Exactly who will these systems serve, what will or should drive their decisions, and how will humans ultimately retain full control of what is doing on? Will the driving force be large corporations, governments – or citizens and consumers?

These are questions echoed by Dr Anderson, who is particularly concerned about the ethical dimension of the machine age. 'Should the robots be trying to tell you something, for example, should we have whistleblowing robots?' he asks. 'Should we have ethical machines in the stock exchange systems, that are making the decisions based on the buyer in pursuit of profit or decisions that are in the interests of the employees of a particular company which could be put out of business due to a buying decision?

'If there is a sales robot, should it be trying to sell you something because it wants to make as much money as possible for the person who has paid for the development of the robot – or for the buyer who wants to make the best possible choice for themselves?'

This is a particularly problematic area due to the limited understanding legislators have about the way that the Internet works and the relatively poor representation of consumers in this new global marketplace. While stock market regulators have quickly evolved a whole host of mechanisms to ensure – in theory, at least – probity in the markets, including sophisticated software analysis of market patterns, similar systems of control have not yet been developed for the Internet at large. This has given companies the freedom to evolve highly sophisticated systems for 'market rigging'.

Among the most obvious of these is the use of 'search engine optimisation' to promote a company, as touched on above. For some years, and despite calls from European countries to prevent this, companies have spent considerable sums of money to influence web searches. This can involve buying a search for a particular word or writing a web page in such a way that it improves the search engine ranking. Pages are deliberately written – arguably 'programmed' – to appeal to the search robots themselves and achieve a high hit rate for SEO, while programmers abuse the scoring system used by the search engines to also achieve the same end.

Many leading Internet thinkers, such as Jaron Lanier, are now arguing that there is a need to reverse humanity out of the machine that has developed and start building a new system that meets its needs. Since the start of industrialisation, humanity has built machines that it has imposed upon itself, and computers are a very good example of this. Changes in our behaviour can be seen by the way that people in the United States and elsewhere now mould themselves to 'fit' forms and profiles in order to create credit histories. In effect, they have begun to make themselves more like machines in a bid to succeed in a world where machines are increasingly making autonomous decisions.

Cyber crime and the Internet of Things

Each of the experts interviewed for this report expressed profound concern at the way this world was developing, the speed at which it has been occurring, and above all the lack of debate surrounding it. This is especially concerning given the abuses of personal information that have occurred due to the algorithms used by the NSA and GCHQ. Legislators, in particular, have arguably focused too much on issues such as encouraging big data and controlling 3D printing of firearms, and not enough on the collection of personal information and its impact on individual rights.

According to Melissa Hathaway, US President Barack Obama's first advisor on cyber security strategy, the debate, particularly in the US, has been virtually paralysed by political concerns.

'The next few years are going to be crucial for the internet and the US is not in the best possible place to respond to that – at the moment it is facing a number of financial issues and the government itself is not particularly cohesive.

'This is coming at the same time that a number of important things are happening; in October 2014 there is a significant meeting to discuss the future of the internet and that is the UN World Summit on Information Society and it will be important for the US Government to think about what is the positive narrative for the internet and how will the US work with allies and other countries to promote the economic health and well-being of the internet,' she says.

This is an important point because much of the key Internet infrastructure is still in the US, and thus the US has an important part to play in that debate. After the revelations about Prism, many countries may not object to acknowledging the important role the US plays in the Internet, but is still something that has to be acknowledged.

It is also important to note that political attention on the development of the Internet and Internet of Things has been diverted by other issues. These issues may initially appear more pressing, such as the war on terror, concerns over global warming, and the economic repercussions of globalisation.

All of this meant that technology issues have not received the full consideration that they need, particularly in the area of law and regulation. These issues include cyber crime, the changes wrought by social networks, the ramifications of the rapid and wholesale penetration of information technology into our lives, and so on.

This is a point Hathaway underlines. Every household is now equipped with Internet-capable devices – not just mobile phones but laptops, tablets, smart televisions, eBook readers and PCs – and to these we will rapidly add fridges, smart meters and cars. All of these devices will be connected to the Internet. We will be able to connect to them from inside the house via Bluetooth or Wifi, accessed remotely from almost anywhere, we will be able to grant access to our family and friends; all this to manage our homes and devices.

All of this is great and makes life more efficient, but it also makes us more vulnerable to attack from the unscrupulous.

Self-programming software?

Any potential problems with the Internet of Things and how it will increasingly dominate our lives will only grow when even more sophisticated software enters the scene. The next generation of software robots may involve a form of self-programming or decision making, based on the situations they encounter via the Internet. This is a risky step for any technology, given the possibility of computer viruses running amok. For example, 25 years ago the first internet worm – the Morris Worm – jammed the fledgling Internet after developing in a manner unforeseen by its creator; Stuxnet, which was designed

to be a stealth virus and was targeted only at Iran's nuclear industry, still managed to attack power plants in Asia and caused damage to the US oil company Chevron. It has already been reported that a programming virus has attacked the IoT[14].

According to Professor Neil Barrett, any move to self-programming systems would be extremely worrying. Barrett suggests that the initial moves towards artificial intelligence will involve locking down parts of the code. 'The way that such software works is through "adaptive" programs. These have a fixed core of functionality, augmented by a set of varied additional functions, pre-programmed by the author but switched in or out by the program as required. Something very like this is used to make mutating computer viruses, for example.'

This creates a self-programming, autonomous program, and is a step that an organisation should take only if it is completely confident that it is in control of what that system can do – a difficult guarantee for anyone involved in technology to give. This is why many of the experts we interviewed consider software machines to be the greatest potential threat to humans from the new world of machines.

Old and new data

A more prosaic concern for policy-makers is what to do with the large amounts of 'old' data still stored on 'old' computer systems run by governments and some large companies. This data is currently separate from the 'new' world of big data.

The chief issue is how to migrate this old data from the legacy systems where it is stored, and whether its value can ever be fully realised. There are also concerns about data protection laws that prevent much of the old stored data from being mixed or 'consolidated' with other data because. Mayer-Schönberger explains, when the data was collected it was for explicit pursposes, not necessarily those that people may wish to use it for in the future.

This has led to various initiatives aimed at finding Cloud solutions to allow data to move from government data pools onto the web. This has been dubbed the G-Cloud, or Government Cloud. For government and big businesses such as financial institutions, this represents the final move from the secure computing environments of the phone age, onto the Internet, where data is in computers with restricted access in secure locations and in the Cloud.

The move to the Cloud is a step fraught with risk, however. Currently much of the data is held on legacy systems and in differing coding architectures. For both governments and big business this represents a huge problem as it is being stored against the risk of loss while offering no benefits. The efficiencies promised by the Internet of Things will be not be fully realised until this old data is harnessed in conjunction with the data generated from these other sources.

Banks and financial institutions are at a disadvantage against less restricted, Cloud-based competitors. For example, mobile phone companies are able to develop the single customer profile that the banks have not been able to deliver. Governments are also in a difficult situation, but for different reasons: if they try to pool together all their data they are at risk of being accused of developing a 'Big Brother' computer system to monitor their citizens' behaviour. Banks, on the other hand, have a better track record than governments in terms of their ability to protect data. This is arguably due to commercial pressures arising from the possibility of reputational loss, litigation and the banks' view of data as an asset.

On the other hand, a failure to develop a government variant of the single customer profile much beloved by business marketeers will lead to accusations of government technological backwardness, incompetence and inefficiency. In the light of the expected explosion in health data, this will be politically difficult for a government.

It will be even more difficult for governments to justify failures of centralised computer systems, particularly in countries like the UK at a time when budget cuts demand efficiencies that can only be delivered using 'smart' technologies such as the smart grid and smart cities.

Section Four– Putting ethics (and better code) into the machine

We have seen how the Internet of Things and big data will throw up major problems for consumers and citizens – problems that have barely been grasped by most policy-makers. With this growing complexity, the risk of machines making unanticipated actions increases the potential for unintended consequences. There are key issues, too, about our reliance on data at a time of massive data generation, data storing and data preservation, which have the potential to both obscure results and generate injustices.

Perhaps the greatest issue that we now face is caused by our blind faith in machines. We have invested them with certainty and we trust them. Part of the reason for this is an odd confusion that has conflated the machines of the industrial age with the machines of the information age: we trust that machines will do what they are meant to do.

We assume that our cars will start, that our washing machines will wash and that our electric drills will bore holes. When their mechanical controls are replaced by software controls we still assume the same thing because most of us are unaware that this has happened. As we have seen with the Snowden affair, however, the extent that software systems have penetrated our world is not widely known by the population at large, and the ramifications are only now being appreciated.

While it could be said that this has generated risks to privacy and freedom, those risks are merely hanging off the dominant ethical concerns.. This is not simply the ethics of covert surveillance of populations; it is the ethical impact of creating a world so complex it is incomprehensible to humanity and beyond its control.

This world is becoming so complex that those involved in its creation warn that it has the capacity to do things that we are unaware of. Furthermore, it could start a chain of events that we would be the victims of – either as a series of decisions, or as the result of machine error leading to a catastrophe.

So how can we start to improve the system? We can begin by ensuring that only the safest and best code is used in this complex system. Until now, we have had a poor understanding of this issue: computer games consoles are currently more secure than medical computers that control patients' lives. In the Higgs Boson experiments on the large hadron collider at CERN the scientists had to employ a number of code specialists to weed through the programs to ensure the veracity of the code and thus the veracity of the experiment's results.

As Professor Shanahan makes clear, perhaps the greatest potential risk is the lack of human restraint in the system and the potential for the system to make decisions that have an impact on us without our knowing.

'I think if we get it right this symbiotic relationship is beneficial and that largely this technology is pretty good stuff, but we can get things wrong and because of that, that can mean bigger implications for us today than it did in the past.

'A minor programming error in the past might just be confined to your desktop whereas now something can be released into the 'wild' and cause all kinds of problems, and I think the potential impact of small engineering mistakes – let alone malicious mistakes – is going to increase as time goes on,' he says.

As Professor Shanahan and Dr Cave point out above, embedding artificial intelligence technology in systems increases the risk of losing control of technology. The potential process of machine 'evolution' increases this risk, possibly to the detriment of humanity.

'We all know about computer viruses and computer viruses that can become increasingly intelligent, that can be made to be increasingly intelligent - they can also be made so that they can improve themselves so that they can "evolve" in which case they can change in unpredictable ways,' says Professor Shanahan.

'So there you would have little packages of intelligence that were moving around, as it were, but you will also have AI that is in static systems that is doing all kinds of things like deciding whether we should be given a mortgage or insurance or surveillance systems

making decisions about us,' he says, admitting that given the potential for catastrophe, there is now a need to implement a much more rigorous system to check computer code before it is released.

'Certainly one thing that *[we can do is]* try to build our code so that we are better able to verify in a formal mathematical way whether it is working properly – and whether its security has been violated in some way.'

The perils of untested software

We still allow the computer industry to road-test unfinished software in beta (trial) form. Gary McGraw of the computer security software company Cigital, says: 'In some cases the beta software is doing things such as controlling nuclear power stations.' McGraw notes that many politicians are unaware of technology issues and suggests that in the field of computer security Europe is 18 months behind Washington – which is itself off the technological pace.

'Washington lags very much behind the cutting edge of technology and computer security is very much at the cutting edge of technology. It's a little like when buildings were going up faster than the legislation in places such as San Francisco and Chicago [in the 19th and early 20th centuries] and there were no fire codes and it took burning a couple of cities down to the ground for us to say: "Maybe there's a better way to do this".'

As Professor Shanahan points out, culpability for problems caused by software in the system currently lies with the computer manufacturers. This means there is massive potential exposure to a disaster, one that the computer industry[15] would rather not consider.

'Putting ethics into it is a difficult thing to do, because it is very much like passing the buck by the engineers to the computer and saying that "the computer says no" and "the computer says kill" and that's a very back-to-front story – because the responsibility is down to the programmer to make sure that the thing works correctly. We are not envisaging yet some kind of future where the AI is genuinely autonomous like we are, and having consciousness,' says Shanahan.

Professor Susan Anderson and her husband Dr Michael Anderson are adamant that computer systems should not be deployed in situations where the consequences are unclear. 'We've always said that if the ethics isn't clear for a machine functioning in a particular domain we are opposed to putting machines in the domain and we say that repeatedly,' says Professor Anderson.

Many in the technology industry would reject this as idealistic and unworkable. After all, much of the modern world is already run on software and machines, and restricting their use for security or ethical reasons could have economic consequences. Howard Schmidt, President Obama's former cyber security czar, admits that balancing the interests of commerce and security is not an easy task.

'We would be in a situation on the cyber security committee where we would say "no, that's it - we are going to pull the plug and stop this right now",' says Schmidt. 'And then I would go into the economics committee and they would say "no, you just can't do that".'

Already there are calls for a radical overhaul of the base code of the Internet and computing on which we rely to make it more secure, and to build security in from the beginning. Before Bill Gates ceded control of Microsoft, he committed the company to adhering to the Trustworthy Computing Initiative[16] to improve the company's software.

According to many observers such moves, while welcome, are not enough. The amount of poor code already developed at high speed to meet commercial pressures has left us dependent on an Internet system that is as unsafe as the car industry was in the 1930s. It is onto this unsafe – some would say rickety – infrastructure that we are now planning to launch the Internet of Things. This is a process for which no one person or organisation has overall responsibility, while people releasing software do so with no concern for any over-arching architecture or infrastructure.

In other words, there is no guidance to state whether or not you have released a safe or unsafe vehicle. As a result, the Internet and computing have to a large extent become an 'ethics-free zone'. As we have seen from the actions of companies and organisations seeking to harvest our data, there is little concern for the rights

of individuals because the computer code and the Internet of Things turn them into data and strip them of their humanity. The same is true of computer software as we have seen from the row over the NSA's use of data culled from the mobile phone app 'Angry Birds' – a game mainly played by children. The NSA's 'exfiltration' of data is exactly the same as the actions of tens of thousands of companies that have built apps for exactly the same purpose[17].

It should come as no surprise, therefore, that if there is such a wilful disregard for individual privacy rights, then the same holds true for the development of other software systems.

Indeed, there is widespread ignorance of the fragility of the system, and of our dependence on it. As a number of experts have pointed out, the power grids in both the US and Europe are particularly vulnerable because of this uncontrolled evolution.

While the results may represent a practical risk to humanity, there are also ethical considerations about how this situation has been allowed to develop. This problem is also predicted to accelerate due to the emergence of the IoT.

It would be better, say some observers, to introduce the equivalent of a Federal Drug Administration to prevent the roll-out of untested systems, and to ensure that safeguards are built in – and to build in a system of control that allows human beings to effectively assert their rights.

A European software certification agency would, inevitably, be criticised early on for being unwieldy or for slowing the pace of commercial competition and hampering the development of software in Europe. But demands for light touch legislation will only be tolerated until it is deemed that legislation is essential because light-touch administration has failed. Post-event legislation frequently follows rapid technological change, as has been mentioned with the automotive legislation of the 1930s, and the close control exerted on the avionics industry by bodies such as the European Aviation Safety Agency following concerns over the safety of air travel.

Industry-specific legislation is also drawn up following particular crises. This was the case with large companies such as Enron and WorldCom that led to the Sarbanes-Oxley and Basel II legislation. the regulations that have been ushered in following the credit crisis in the US and Europe, and the reform of the US hotel industry following the rape of singer Connie Francis and her subsequent $2.6m lawsuit against the Howard Johnson Motel group. The **Sarbanes-Oxley Act** of 2002, is a US federal law that set new accounting standards for US public company boards, management and public accounting firms. The bill was enacted following a number of high profile accounting scandals surrounding the collapse of large companies such as Enron, Tyco and WorldCom. The act contains eleven titles or sections ranging from additional corporate board responsibilities to criminal penalties and requires the Securities and Exchange Commission to implement rulings to comply with the law.

Basel III is a set of banking regulations set up by the Basel Committee on Banking Supervision, a committee of major banking supervisory authorities set up to improve the quality of banking supervision worldwide. Basel III extended the remit of Basel II to take into account the financial crisis of the late 2000s and now supersedes the Basel II regulations.

These are examples of a post-event legislative culture that can be avoided, according to Professor Susan Anderson and her husband Michael, by the introduction of a new form of computer code that places ethics at the heart of the new communication systems.

Ethical dialogue with the machines

'It's not a matter of there being a set of ethics for machines and another for human beings; we argue that there is just one thing called ethics. We want to make sure that machines have this ethics built into them,' says Professor Susan Anderson, who asserts that this needs to be an exhaustive process.

'In order to try to capture the ethical principles needed we need to have a dialogue with the machine that is centred just around whatever the domain is that the machine will be functioning in, and try to discover the ethically-relevant features that the machine will have to encounter or deal with, the *prima facie* duties that the machine should be aware of and the decision principles that in the last analysis should govern its behaviour[18].'

Professor Anderson says that in the course of a dialogue between the machine and one or more ethicists, the machine would be able to 'tease out' ethical elements that are relevant to its domain. 'Like, could someone be harmed? That is something that ethicists feel is ethically relevant and should be taken into account,' she says. 'Also, in the area of biomedicine, respect for the

autonomy of the patient is another example of an ethically-relevant feature and then from that *prima facie* duties are discovered by figuring out what the ethicist says that the correct action is, and whether that involves maximising or minimising the features in question.'

Supporters of the idea of ethical code argue that it could be rolled out on a country-by-country basis. Using this model, individual states or areas will be perceived as politically mature and democratic because of their willingness to deploy ethical code.

This process would involve underlining the key ethical requirements for the machine. 'So harm is something that you would want to minimise, respect for autonomy is something that you would want to maximise, causing benefit is something that you would want to maximise,' says Professor Anderson.

The problem comes, she says, when these 'prima facie' duties come into conflict with one another, as we saw in the example of Charles and the Difference Engine described earlier.

'So for example you might have a situation where a machine is trying to remind a patient that they have to take their medication and the patient says that they don't want to take it now,' she says. 'You have a conflict between whatever the purpose was of taking that medication to prevent harm, or cause a benefit, with respect for the autonomy of the patient. It will then depend on input from the doctor to help the machine to figure out what should be dominant.'

Professor Anderson adds: 'This will allow the machine to be able to work out at what point it will hit the time when the patient will be harmed, and the medication reminder system needs to inform the doctor and say "you'd better intervene, there's a real problem here".'

Data protection and privacy

Issues surrounding data protection and privacy will become ever more important with the advent of the IoT. Almost all of the experts we spoke to agree that there is a need for protection for particular machines, certain data and the programs that manipulate that data, and to ensure that this is ethical.

Above all, there is the question of how to approach the data that is being generated. This is especially important because it will include our personal data.

In Professor Adrian Cheok's view, the pervasive nature of the new Internet of Things will mean that privacy becomes impossible, and that the only option left open to us will be to be as transparent as possible. 'I think that what is going to happen is that the majority of us will, by default, just become totally public because of the amount of data that is online about us, because for the average person it is just a lot easier. People use credit cards now because it is more convenient, data use will be the same. We will use our data to make a transaction and to say who we are. Most of us will go transparent,' he says.

This solution only works, however, if we are guaranteed that the IoT and related artificial intelligence systems are also utterly transparent. This transparency will allow one to see what is being done and how it relates to individuals.

This approach will also require adding a new concept into the Internet world, says, Professor Mayer-Schönberger, that of 'relevance'.

He argues that much of the data that is stored about us is no longer relevant and reflects an inexact picture of what we are now. Moreover, that inaccuracy will be imported into the big data collected by the IoT and distort its usefulness. Past data, he argues, may simply no longer reflect who we are.

'I may have once had a girlfriend who was keen on gardening so we did this as a mutual activity but we have now split up and I am no longer interested in gardening but Amazon and Google still try to direct me to gardening books. In so doing they might be upsetting me rather than pleasing me,' says Mayer-Schönberger.

'Digital tools prioritise the preservation of data over deletion, we have built that by default into the system but it does not reflect us. We start to forget things almost immediately and that has an impact on our decision-making and our ability to abstract. Too much information gets in our way,' he argues.

Professors Cate and Cheok, meanwhile, agree over the need for the transparency of data. Professor Cate proposes that our data should have binding conditions attached to it, governing how it is used. The point is reinforced by Professor Mayer-Schönberger. 'The biggest issue relating to data is, how data will be re-used,' he says. 'How it is collected will be of less importance – how it is *used* is the important issue.'

In other words, the Andersons' concern with the ethics of the machines themselves should also be extended to data, its use and its deletion.

Data integrity and quality

Given that the essence of the Internet of Things is the generation of data, and that crucial policy, commercial, military and consumer decisions will increasingly be made on the basis of that data, the data's integrity has to be viewed as sacrosanct.

In the future, the backbone servers of the Internet will have to be zealously guarded against attacks by hackers because of the potential impact upon humanity. According to Melissa Hathaway, organisations have to acknowledge that they owe a duty to the people whose data they collect.

'I think that governments or the private sector have to realise that information is their greatest asset,' she says. 'Putting more and more data into data centres and not really thinking of putting in place the appropriate safeguards for those assets is unacceptable. We are seeing more and more breaches and people are beginning to realise that their data is vulnerable.'

Limitations of the law

Meanwhile, lawyers across Europe admit they face profound challenges keeping up with the pace of technological change. According to Michael Drury, former Director of Legal Affairs for GCHQ, developments in areas such as social media alone have quickly made legislation obsolete.

As a result, Drury says we are currently dependent on technology companies imposing ethical constraints upon what they are doing with data because legislation does not exist to guide their actions. For example, according to Drury, when the UK's Regulation of Investigatory Powers Act (RIPA) was drafted it did not envisage the development of social networks or Cloud computing.

Another good example was the EU e-commerce law that did not make allowance for the rapid uptake in ADSL broadband connectivity even though the technology was known when the legislation was drafted. According to Sir Bryan Carsberg, the first director of the UK's telecoms regulator Oftel, the organisation only ever expected mobile phone use to be at around 500,000; there are now three phones for every person, a figure more or less replicated across Europe.

'How do you define and safeguard for the future? It is a very difficult thing to do, given that no one knows what developments will occur next and no-one really knows what the future development of social media sites will be, to take one example,' says Drury.

'I think that there is a case that due to technological change that we may be on the edge of what can be legislated for under the law. Any statute may be potentially unwieldy and there may be a case to look at a set of principles, defined by a code and regulated by a standing committee.'

Larry Lessig, the noted American legal academic and technological thinker, has argued that the law should give way to computer code. He says that if we want to control what is possible, code is much more efficient than law. This conclusion backs the views of Professor and Dr Anderson on the necessity of introducing ethics into the computer code itself.

The development of an unprecedented system for the collection of data from humanity has coincided with great weaknesses in the protection of the interests of those whose information is collected – namely, us. This is because of the pace of technological change, a lack of understanding of technology among legislators, a regrettable lack of political attention and, most importantly of all, a lack of understanding of a system that humanity has become frighteningly dependent upon.

The case for machine rights – to protect humans

As we have already seen, one question that has been raised is whether there should be some form of 'rights' for the machines that will be helping to run our digital world. It should again be made clear that what we are concerned with here is not a 'robot charter' for super-intelligent androids, an issue beloved of science fiction writers. Rather, it means asking whether to confer some rights on these machines in order to better protect the people they work for.

Warwick University's Dr Cave argues, for example, that there is a case for the creation of an ethical framework for the protection of the smart phone avatars discussed earlier, and which are currently under development.

'I am not saying that machines should have rights in and of themselves, but I do think that two things are true,' he says. 'Firstly, that if they do not have something that looks like a right – the power to take decisions and act on them, for example, or to learn from experience and behave in ways that they were not originally programmed to do, to act as autonomous systems – I don't think my interests, our interests, would be served by our networked interactions.

‘Moreover, the internet as it exists would not exist because it depends on these autonomous systems operating. The question as to whether they should have human rights, though, depends on whether, in acting on the internet, we are acting as human beings.

‘Because if I am being nudged around by all of this information so that I am responding to it but it is impossible for me to know or verify that information and I simply react to it, then I have acted – but I cannot be said to have “decided” or to have made a choice.’

Humans becoming more like machines…

There is an irony here. Traditionally, the Turing Test[19] is used to determine whether a machine is acting intelligently, akin to a human being. But Dr Cave wonders whether, confronted by the vast mass of data around us in the digital world, it is we humans who are at risk of behaving more like machines.

‘So it could be that the Turing Test gets failed in the other way,’ says Dr Cave. ‘It’s not so much that machines can masquerade as human beings, but that human beings, in a sufficiently immersive and interactive world, begin to behave like machines because they know that the decisions that they are making are too hard for them to understand, or they don’t have enough time to make them properly, or the consequences are so awful that if they thought about them they would not actually choose at all.’

Consumer rights and machine insurance

The issue of machine rights may seem theoretical and remote from the consumer, but this is not so. Given the increasing role machines play in our lives – and their semi-autonomous nature – the question will arise when something goes wrong: who do I sue, the machine or humans?

In their 2011 book *A legal theory for autonomous artificial agents* the philosopher Samir Chopra and the lawyer Laurence F. White[20] make a powerful legal and philosophical case for giving ‘autonomous artificial agents’ a form of legal status. This status would be analogous to a legal ‘agency’ status: ‘people’ with the legal authority to act on our behalf.

Chopra and White further argue that such artificial autonomous agents should be given legal ‘personhood’, taking their place alongside humans and corporations as legal entities that can, in theory, be sued. ‘There is no reason in principle that artificial agents could not attain such a status, given their current capacities and the arc of their continued development in the direction of increased sophistication,’ they write[21]. In terms of punishment, the authors of the book say that artificial agents that control money ‘would be susceptible to financial sanctions, for they would be able to pay damages...and civil penalties or fines[22]‘. Chopra and White also note that such agents could also be restrained in other ways, including by being ‘disabled’ – in other words, turned off.

One risk of making software machines liable is that it opens the way for yet more time-consuming and expensive litigation. This is why Chopra and White, and others, have floated the idea of insuring machines against damages they cause. ‘One move ... would be the establishment of a registry that would stand behind registered autonomous artificial agents and insure them when things go wrong, so as to provide some financial backing to the idea of artificial agent liability[23]‘.

In a conference paper written in 2012, Dr David Levy went even further. Admittedly, he was talking specifically about robots for household use or entertainment purposes, but the principle holds for any ‘intelligent’ software-based entity. He suggested a compulsory no-faults strict liability insurance scheme that would pay out when something goes wrong, whoever is to blame.

One reason Levy is so keen to see a no-faults insurance system – a level playing field for all – is that he fears the impact that widespread litigation would have on software and robot development. ‘One of the negative effects of all this litigation is that the growth of robotics as a research field and as a branch of commerce will be stunted because commercial robot development, manufacture and marketing will become such risky businesses,’ he suggests.

The same problem could affect the developers of Internet-based software programs that form the Internet of Things. It is a problem Chopra and White also address in their book, noting that while software providers up to now have largely been given legal protection that would be thought unacceptable for dangerous tangible goods, that situation looks set to

change as more and more software is embedded in machines and objects. ‘Suppliers of defective artificial agents may face increasing liability under professional liability theory, particularly if the judiciary comes to recognise software engineering as a profession with applicable codes and standards[24],’ they write.

Fears that insurance might reduce accountability – as developers would fall back on the fact that they were insured – may be outweighed by the fact that litigation would lead to increased premiums for those involved. This factor already exerts considerable pressure on professionals such as architects who have to carry liability insurance for the buildings that they create.

In his 2012 paper, Levy highlights the obstacles to progress that the threat of litigation can cause. He cites the example of a 1970s computer program called MYCIN developed at Stanford University in the US to identify bacteria that caused severe infections such as meningitis, and to recommend suitable antibiotics treatment. A comparison between the program and five human experts at Stanford Medical School showed MYCIN’s ‘acceptability’ performance was 65%, significantly better than the human experts whose ratings were between 42.5% and 62.5%.

Despite this superiority, says Levy, the MYCIN software was never used in clinical practice. ‘One reason was the legal objections raised against the use of computers in medicine, asking who should be held responsible if the program were to proffer a wrong diagnosis or to recommend the wrong combination or dosage of drugs,’ he wrote[25].

Conclusion

The Internet of Things and the era of big data will bring great benefits. Many of those benefits risk being overshadowed, however, if the real problems posed by this new technological revolution are not addressed.

As we have seen, we are now seeing an unprecedented and unregulated explosion of data, data gathering and data analysis, which our leading lawyers say the law is unable to keep up with.

Also, when there are regulators in virtually every other field – from medicine to transport and communications to energy – the only area we do not regulate is computer software. Yet it is this very computer software that will control the Internet of Things and, with it, the fabric of the world we live in.

Too much focus, we believe, has been placed on the technological advantages that the Internet of Things and big data can bring. Not enough attention has been given, on the other hand, to the impact on humans of living in a world in which we increasingly hand over control of everyday functions to machines and to the new gold standard of the modern world: big data. The benefits of the IoT have been stressed while the dark side of the changes has been largely ignored.

There is now, therefore, an urgent need for policy-makers to consider key practical questions about how to ensure the IoT works for the people, and not independently of the people.

There is also an urgent need to ensure that the system itself is safe and protected. At the moment it is worryingly vulnerable.

Blueprint for action

To address some of the issues raised in this report we recommend the following:

1. Consideration should be given about how to bring ethics into computer programs/software to ensure that human consumer rights and privacy are protected. Citizens' privacy needs to be much better protected from the world of big data, whether through protecting access to that data in the first place or, as many of the experts we have spoken to suggest, placing controls over how that data is used once it has been gathered. The rights of citizens and consumers in relation to the Internet of Things and Internet software need to be codified in a short and simple form. This could include giving machines some form of legal status to ensure that we humans are given extra protection.

2. We call for an end to the current practice of road testing software on the population at large. New software destined to be used in the public arena must be properly regulated and checked for safety and compatibility before it is released. This requires the setting up of a new European technology regulation body; it would essentially be a software-focused equivalent of the Federal Drug Administration in the United States.

While we are aware that the IT industry does not just mean large organisations such as Microsoft and Google and that it is a vibrant and developing industry, the costs of funding this new software regulation body and proving software should not be shouldered entirely by the smaller companies, and they should be helped to 'prove' their work. The patent system is currently unwieldy due to costs and is a significant disincentive to companies to try to work within it. This has led to many companies trying to find ways around the issue. At the same time, we believe it would be unfair for the taxpayer to fund software regulation. To protect the interests of both consumers and small-scale developers, we suggest the IT industry provides a sliding fund for the proving of technology, based upon company size.

Another key function of this new Europe-based technology regulation organisation should be to inform governments and politicians of the significance of technologies. Much good work has already been done by the EU in bringing companies such as Microsoft to account. This has meant that the EU is now seen as taking a lead in this area. This new organisation would set the benchmark for the rest of the world and ensure that Europe is seen as a centre of probity.

The new technology body would also have the key role of informing the public. There is an urgent need to increase the awareness of the population at large about the significance of the Internet of Things and what it means for them. This is something that the IT industry is not currently doing. It has a vested

interest in promoting the benefits of technology and not its demerits.

A final role for this new technology regulator should be that of an infrastructure planning agency to understand exactly how much of the internet system is European and what we control. Its remit would be to draw up contingency plans to bring back limited parts of that infrastructure under European control in the event of a widespread attack upon it.

3. We call for the development of technology that can make data anonymous and at the same time produce valuable data that is of benefit to society as a whole. We contend that the only way that this may be possible is by the development of an ethical computer system that stipulates how the Internet of Things can use information.

4. We suggest there is a need to reinforce what we call 'device sanctity'. As smartphones, devices and the software they use become increasingly personalised, it is important that these devices are loyal to the individual who owns them. Devices considered to have 'a human interest' need to be properly protected against incursions from both the state and cyber criminals; a protection enshrined in law.

5. Primacy of interest. It is now possible for a number of different groups to have an interest in a device such as a smartphone – the person who bought it, the telecommunications company that runs it on our behalf, companies such as Facebook, Google or LinkedIn to whom we have granted an interest in our whereabouts, the government and the police. It is essential that the order of primacy in this interest is made clear and asserted.

Individuals should have to actively opt into the Internet of Things if the use of their device is being solicited by another party, and the implications of signing in should be made clear.

In return for services offered by the IoT there should be a 'cooling off period' before those wishing to use a service can participate. Data must not be used without an explicit 'buy-in' from the person concerned.

We suggest that consideration should be given to imposing compulsory insurance for computers and devices. and for those who are producing software for those devices, for the Internet and for the IoT.

We believe the issues are major ones; nothing less than the future safety of the internet and the acceptance by citizens of this new technological world are at stake.

While most consumers seem to have embarked on a deep love affair with their smartphones, devices which, as we have seen, will be most people's main contact with the Internet of Things, this technological love-in cannot be taken for granted.

If, over the coming years, more and more people feel alienated, lost and no longer in control of the world they live in, there could be a significant backlash against the machines, software and all things technological.

Up to this point in history, humans have been able to touch, see and intuitively understand how the world around them works. This reassuring handle on the world will start to disappear with the advent of the Internet of Things, which is increasingly likely to be seen as vast, complex, hidden and mysterious.

We have seen in the recent – and ongoing – financial crisis how the complex world of finance lost the trust and confidence of many people when they were confronted with the real world impact of vast transactions and operations that they did not understand and were seen as being damaging to society's interests.

How much greater will the risk of alienation become if people feel they are suffering as a result of the complexity of the everyday world itself, one that is perceived to be run by machines and not always in the interest of us, the consumers?

This is why the technological optimism of the new digital world must be accompanied by pragmatic policies, rules and workable legislation to reassure people that they are still the masters of the world in which they live.

Visible, concrete, practical and robust measures need to be adopted to show citizens that the technological world is both safe and here to serve people – and not the other way around.

That way, the new age of machines can do what it was surely always intended to do: make life a little easier and more efficient for we humans.

About the authors

Michael Streeter
Michael Streeter is an author and former Fleet Street executive who worked for The Independent, the Daily Express, the Mirror and the Daily Mail. He was also editor of the Scottish Daily Express and launch editor of the Daily Express website.

Peter Warren
Peter Warren is an award-winning newspaper and TV journalist acknowledged as an expert on technology and computer and internet crime. He wrote the first articles highlighting the potential for the Internet to be abused by paedophiles in 1989 and, as a result, was asked to brief the first UK police force to respond to the danger, and the Greater Manchester Police Obscene Publications Squad on the issues the technology has produced. He has also set up and currently chairs, the Cyber Security Research Institute, an organisation pulling together the UK's top academic and business experts in the field of computer security with leading journalists in a bid to raise awareness of cybercrime.

Jane Whyatt
Jane Whyatt is an experienced journalist and radio producer currently specialising in new technology and its impact on our lives. She produces PassWord with Peter Warren, the UK's only live news talk show about new technology on independent radio. Jane has made hundreds of BBC radio programmes as an indie producer working with Robin Lustig and also worked as a regional radio news editor. Her company, Angel Media Productions CIC, is a registered provider to all BBC networks and has just produced a documentary about immigration for BBC Radio 1, transmitted in March 2014.

Sources

http://share.cisco.com/internet-of-things

1. The science fiction writer and futurist thinker Isaac Asimov predicted a world where robots were dominant entities and, in an attempt to deal with it, suggested his now famous three laws of robotics:

1. *A robot may not injure a human being or, through inaction, allow a human being to come to harm.*
2. *A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.*
3. *A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.*

However, Asimov's idea that robots would start to resemble humans seems as far away as ever. Some think that Asimov's sci-fi rival Arthur C Clarke was closer to the mark with his creation of the robot Hal in the novel and film 2001: A Space Odyssey.

Edward Snowden (born June 21, 1983) is an American computer expert who leaked huge amounts of classified US intelligence documents to the press.
He began his career working as a system administrator for the Central Intelligence Agency and as a counter intelligence trainer at the Defense Intelligence Agency.
While working for the CIA in Switzerland Snowden began to question the morality of the work he was doing. Despite this he went on to work for the computer company Dell, a contractor with the National Security Agency, a posting that took him to Japan.
While at Dell Snowden began obtaining the 200,000 documents that he passed to journalists working for the London Guardian on a number of top secret surveillance projects aimed at collecting communications data across the world..
In March 2013, Snowden joined Booz Allen Hamilton a computer consulting firm working for the NSA in Hawaii. In June 2013 Snowden became an international *cause celebre* when it emerged he had leaked thousands of secret documents on the illicit surveillance techniques that had been developed by the NSA and the UK Government Communications Head Quarters.
The documents revealed the existence of a number of massive surveillance programs run by the US, UK, Australia, Canada and New Zealand.
Among the systems revealed by Snowden were:
Boundless Informant, a big data analysis and data analysis visualisation tool that provides the NSA with summaries of its worldwide data collection activities, including data on US citizens in contravention of NSA assurances to the US Congress.
Prism a system for accessing stored communications data held by telecommunications companies. **Xkeyscore**, a system used to track everything that an individual does on a computer anywhere in the world.
The **Fascia Database** which holds trillions of device location records.
Tempora, a system used to store internet communications taken from fibre-optic cables.
Windstop, this is the collective name for a number of systems that collect information from Google and Yahoo networks.
Optic Nerve, a system that collects webcam images globally.
Dishfire, which collects text messages from around the world.
Squeaky Dolphin, a program that collects and analyses information from social media networks and the **Joint Threat Research Intelligence Group**, a cyber attack system used to discredit individuals, plant misinformation and disrupt computer communications.
The release of the material was called the most significant leak in US history by Pentagon Papers leaker Daniel Ellsberg.
Snowden has variously been described as a hero, dissident, traitor, patriot and whistleblower.

3.

4. http://en.wikipedia.org/wiki/Internet of Things
5. https://www.abiresearch.com/press/more-than-30-billion-devices-will-wirelessly-conne
6. http://www.rfidjournal.com/articles/view?4986
7. http://www.bigdatanews.com/profiles/blogs/ge-s-clarion-call-on-cloud-and-analytics
8. http://en.wikipedia.org/wiki/Big_data#cite_note-22
9. http://www.sas.com/big-data/
10. http://www.theguardian.com/commentisfree/2013/dec/09/tech-giant-companies-open-letter-white-house
11. http://www.bbc.co.uk/news/technology-18393985
12. http://thehackernews.com/2013/11/russia-finds-spying-microchips-planted 1.html
13. Google's acquisition of Nest is one such example, a purchase that unites a smart home with a data agglomeration giant http://www.theguardian.com/technology/2014/jan/13/google-nest-labs-3bn-bid-smart-home-devices-market
14. http://allthingsd.com/20131130/a-new-worm-proves-that-the-internet-of-things-is-vulnerable-to-attack/
15. As many of our experts pointed out, the 'computer industry' is in reality a disparate mix of groups that includes start-ups, self-taught developers, university projects, amateurs and a host of other interested groups. This is analogous to that of the car industry in the 1930s, which spawned a whole host of dangerous technologies until it had safety constraints imposed upon it. Though the challenges inherent in implementing changes in computer software should not be underestimated, arguments put forward in this report that an insurance policy is needed to write code in the new world could be one solution to this.
16. http://en.wikipedia.org/wiki/Trustworthy computing
17. According to research carried out by the computer security company Bit9 of 1.4m apps on Android mobile devices 180,000 of those removed data and of those 120,000 were malicious
18. A difficult process, but there is the beginning of a consensus that the world needs to start talking about the fundamentally important issues that computer code raises, and there is growing awareness of the need for an international organisation to oversee computer code. World leaders have already begun to acknowledge that there is a need to have mutual agreements about the development of weaponised software, evidenced by the talks held by British Prime Minister David Cameron in

China on the issue in December 2013.

http://en.wikipedia.org/wiki/Turing test

Laurence F. White and Samir Chopra, *A legal theory for autonomous artificial agents*, University of Michigan Press, 2011.

20. *Legal theory*, p 189
21. Ibid, p 167
22. Ibid p 149

It has been suggested by Ben Hammersley, UK Editor of Wired Magazine that many jobs will soon be 'algorithmised' such as insurance and law, a process that has already begun. Many insurance quotes are now already decided upon by software and increasingly the law is being logged onto computer databases that algorithms can search according to plaintiff circumstances and precedent. A process set to continue and broaden. Tractors and combines working in fields are already roboticised and can be satellite controlled to achieve specific results such as depth of ploughing, and fertiliser injection according to information produced from the satellites of soil deficiencies.

The IoT will allow this process to be further fine-tuned, it will also allow the roboticisation of delivery replacing lorry drivers with robot operators. In December 2013, Amazon announced that it was experimenting with the development of robotic delivery drones.

23. *Legal theory*, p 149
24. Dr David Levy, *When Robots Do Wrong*, Conference on Computing and Entertainment, 3-5 November 2012.

NETOPIA

Forum for the Digital Society