DIGITAL LEADERSHIP

An interview with

Michael A Osborne

Associate professor
at the University of Oxford

The New Innovation Wave: Machine Learning and AI

Michael A Osborne

Associate professor
at the University of Oxford

Your latest research with Professor Carl Benedikt Frey and Citi on innovation shows that throughout history powerful interests have hindered innovation. Is history repeating itself?

History might be repeating itself. Technological progress has always created significant disruptions that powerful interests tried to combat to maintain the status quo. In Ancient Rome, the great author, Pliny the Elder, records the story of an inventor who discovered a way to manufacture glass that was unbreakable. Hoping to receive a reward, the inventor presented his invention to Roman Emperor Tiberius. Unfortunately for the inventor, Tiberius had him sentenced to death, fearing the loss of jobs and creative destruction due to the new technology.

The various innovation waves in history have led to dramatic upheavals in society but we have never seen large-scale technological unemployment. For example, at the beginning of the 20th century in the U.S., about 40% of our employment was engaged in agriculture,

whereas at the end of the 20th century, it was less than 2%. There was quite a seismic change in the makeup of employment due to technology. However, none of these changes actually affected unemployment much; the unemployment rate was about 5% at the beginning of the 20th century and was still 5% at the end of the 20th century.

> *Roman Emperor Tiberius, had [an inventor] sentenced to death, fearing the loss of jobs and creative destruction.*

Is the current digital innovation wave different from the previous innovation waves?

Yes, in the past technology has largely just substituted the human muscle, whereas now technologies can increasingly substitute human brain labor or human cognitive work. If machines can substitute for cognitive labor, as they have done in the past for physical labor, it's not very clear to what extent there might be any demand remaining for human labor. Another thing that is different this time is the accelerating rate of technological change. It is undeniable that we really are making quite dramatic leaps forwards in designing intelligent algorithms that might be able to substitute for human workers.

Can you give us your definition of machine learning and artificial intelligence?

Machine learning is a subset of artificial intelligence (AI) and can be defined as the study of algorithms that can learn and act. It's about the reproduction of some of the most quintessential human characteristics. The boundaries between all these different fields are not very clear. However, you might associate computer vision, language processing and speech recognition as other subsets of artificial intelligence.

Could you give us some examples of current applications of artificial intelligence and machine learning in business or non-business environments?

A great example of this is the retail business, where 'recommendation engines' are being used extensively. Take Amazon as example. It recommends products to millions of customers on the basis of historic data. Algorithms can process data, characterize the spending patterns of millions of customers, identify latent trends and patterns, and identify clusters and communities for recommendation.

> *In the past technology has largely just substituted the human muscle, whereas now technologies can increasingly substitute human cognitive work.*

Self-driving cars are another application where machine learning is really playing quite a pivotal role. Cars today have sensors, onboard computing and safety systems, and use intelligent algorithms to render themselves autonomous. These algorithms observe how you drive using the onboard sensors and quietly learn to travel from the current to desired positions on the basis of this data.

Speech recognition is one key application of machine learning. Voice recognition software, such as Apple's Siri, has been driven by advances in machine learning. Deep learning is becoming a mainstream technology for speech recognition and has successfully replaced traditional algorithms for speech recognition at an increasingly large scale. Historically, people were trying to induce speech recognition by recognizing speech features like words and phrases. But recently, algorithm designers have turned to massive data that has become available for characterizing human speech. If this trend continues, we are likely to continue to move from the 95% accuracy rate that we're able to achieve on our mobile devices today to nearly 99%. This would probably make these speech recognition techniques completely widespread.

> *Machine learning is a subset of artificial intelligence (AI) and can be defined as the study of algorithms that can learn and act.*

You mentioned 'deep learning'. Can you shed some light on this concept?

"Deep learning" is the new big trend in machine learning. It promises general, powerful, and fast machine learning, moving us one step closer to artificial intelligence. Deep learning is a set of techniques, inspired by biological neural networks, for teaching machines to find patterns and classify massive amounts of data. It tackles crisp and well-defined problems, such as classifying images as to whether they contain a car or a book, for example. It's worth noting that deep learning is still not well-understood, or provably robust, and that it is likely unsuitable for safety-critical applications.

What are the most promising opportunities for both AI and machine learning?

In our recent paper, "The Future of Employment", Frey and I have identified occupations that we thought were the most susceptible to automation. These are the kind of jobs that would be most easily automated within the near future using machine learning techniques. The model predicts that most workers in transportation and logistics occupations, together with the bulk of office and administrative support workers, and labor in production occupations, are at risk. I think that in all these kinds of areas, there will be an ever-increasing degree of automation using machine learning algorithms.

> *More or less anything that does not require one of the three bottlenecks – i.e. creativity, social intelligence and the requirement to manipulate complex objects in an unstructured environment – will be potentially automatable.*

More or less anything that does not require one of the three bottlenecks — i.e. creativity, social intelligence and the requirement to manipulate complex objects in an unstructured environment — will be potentially automatable in the near future. When self-driving cars become a reality, there is going to be an enormous economic impact for taxi and truck drivers. In the near term, self-driving technologies are going to have an impact in the automation of professions, such as forklift drivers, agricultural vehicle drivers, or mining vehicle drivers.

Are there sectors where you see a lot of potential for machine learning?

The retail sector, for example, offers tremendous opportunities. We have seen restaurants deploying tablets on tables. These tablets can be used to take orders from customers. They can also recommend customers a set of dessert options tailored to what customers might want based on what they have ordered previously. A customer might even have a profile that's built up over multiple desserts in the same restaurant or even a family of restaurants using the same software. So, you might get to the point at which this tablet can instantly know, as soon as you walk through a door, what your usual order is, and expedite the ordering of it.

Even in the traditional mall, the customer movement around the store can be closely monitored, indicating customer interest in products. The change here is that these intelligent algorithms can become more personalized and gather a lot more data about us as individuals. Therefore, they can make recommendations to us that are more directly targeted at what we want and what the store might actually want to promote.

Can individual companies invest in artificial intelligence?

The wonderful thing is that you don't necessarily need to be developing novel artificial intelligence techniques and machine learning algorithms to benefit from the state-of-the-art. It's relatively low cost, an individual corporation can just hire a data scientist or a machine learning PhD. There is this wealth of
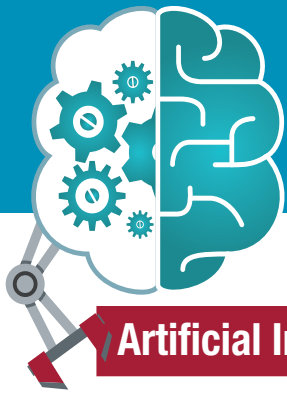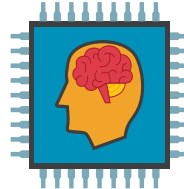
Making Machines Learn

Artificial Intelligence and Machine Learning

What is machine learning?
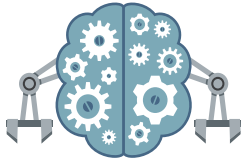
- A subset of artificial intelligence (AI).
- Study of algorithms that can learn and act.
- It aims to reproduce some of the most quintessential human characteristics.
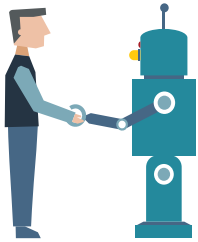
How big is machine learning?

Research on AI constitutes approximately **10%** of all computer science research today[1].

How hot is machine learning?

Venture capitalists have invested over **$300 million** in AI startups in 2014, a **twentyfold** increase from 2010[2].
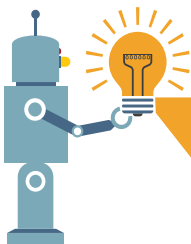
Working with machines

"We need to equip the next generation of workers with the necessary skills that they will need for an economy that is increasingly automated".

- Michael A Osborne

At the same time, organizations need to **protect** themselves **from deskilling** – loss of human skills due to intrusion of automation in occupations.

No heavy investments needed – a wealth of open-source, state-of-the-art algorithms already available, it costs very little to reproduce an amazing algorithm and it benefits everyone.

[1] Machine Intelligence Research Institute, "How Big is the Field of Artificial Intelligence?"
[2] Bloomberg QuickTake, "Artificial Intelligence", July 2015

open-source state-of-the-art algorithms already out that could be immediately deployed to applications. So, the bottleneck is really not investment, it's talent. The challenge is just attracting people with the right skills.

> "
> *There is this wealth of open-source state-of-the-art algorithms already out that could be immediately deployed to applications.*
> "

What are the risks for a company investing in AI or machine learning?

If you rush too fast and hard into automation, there are skills that will be irrevocably lost to organizations. In his book, "The Glass Cage: Automation and Us", Nicholas Carr explains how automation can separate us from reality and deskill us. One example that is particularly evocative is that of an airline pilot. Carr plausibly makes

the case that pilots' increasing reliance on automated flying reduces their own ability to actually fly a plane in any direct way. The case of airplane navigators is also very interesting – those who have not lost their jobs have been deskilled; more and more of their job is being handed over to an algorithm. As a result, when something unforeseen happens, when there is a kind of an unexpected event, the algorithm becomes useless, and humans take over, they won't necessarily have the same degree of understanding and skills that they might have had before the automation. We can argue that there have been catastrophes because of a human skill being lost owing to the intrusion of automation in that occupation. I think this is a real risk for many businesses; automation of labor might result in the loss of softer skills that weren't necessarily quantifiable but are yet essential.

> "
> *Automation of labor might result in the loss of softer skills that weren't necessarily quantifiable but are yet essential.*
> "

What would you recommend to CEOs of large organizations regarding machine learning and artificial intelligence?

We really need to develop in-house expertise in these techniques because, fundamentally, the increasing availability of data across a wide range of different industries is going to be transformational. If your business is not in the forefront of using this data, there certainly will be another business. The potential value that could be realized is going to develop into a real competitive advantage. I think companies need to be doing a lot more to recruit and hold on to talent such as data scientists and machine learning experts.

What are some of the ethics issues that we will face with machine learning or AI?

The big problem here — and one that the machine learning community is spending a lot of time on — is privacy.

Our objective is to have systems that could be verifiably shown to not intrude upon someone's private space. Ideally, you would have an algorithm that would process, for example, medical health records without ever returning to the users of that algorithm any personal details. These algorithms look at this dataset and return aggregate statistics
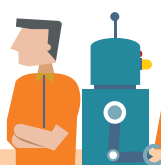
Occupations vs. Automation

Carl Benedikt Frey and Michael A. Osborne studied 702 detailed occupations to estimate the probability of computerization

About **47%** of total US employment was found to be at risk of computerization

21% of US employment and **24% of UK** employment is **highly creative,** involved in areas such as development of novel ideas

Creativity is a hindrance to automation – creative skills cannot be readily replaced by machines

As a result, **86% of US workers** in the highly creative category are **at low or no risk of automation**; the equivalent number for the **UK is 87%**

Source: Nesta, "Creativity vs. Robots: The Creative Economy and the Future of Employment", April 2015; Carl Benedikt Frey and Michael A. Osborne, "The Future of Employment: How Susceptible are Jobs to Computerisation?", September 2013

useful to researchers without ever returning to those people, the details of any individual patient. This relies upon the robustness of the algorithm.

> " *Value is delivered to us as customers or as consumers. Let's make sure that all this wealth that is created is distributed in a fair way throughout society at large.* "

How can governments and regulators try to mitigate the adverse effects of machine learning?

I feel the most imminent concern is the effect of machine learning technologies upon employment. People have rightly recommended that these technological events would necessitate the change of education. I think we need to equip the next generation of workers with the necessary skills that they will need for an economy that is increasingly automated. Governments do have a role in investing in the technological research. There are profound challenges here in getting governments to select the right type of technological development.

Is there an argument to slow down the progress of machine learning?

I don't think there is. An important thing to note is that all these technologies will be delivering enormous value to society. We have seen in the last 10 years the delivery of amazing technologies. Even the smartphone in your pocket, for example, has a range of apps that would have been unthinkable even 10 years ago. My point is that value is delivered to us as customers or as consumers. Let's make sure that all this wealth that is created is distributed in a fair way throughout society at large. Moreover, the wonderful thing about the technology is that it has very low marginal cost of reproduction. So, at the point that you develop an amazing new speech recognition algorithm, it costs very little to deploy that algorithm on everybody's mobile phone. Everybody benefits from the introduction of those techniques as customers and as consumers. The problem is that some people necessarily will be put out of work by these technologies. I think there is a challenge to make sure that there are jobs created to replace those that are eliminated by new technologies.

Michael A Osborne

Associate professor
at the University of Oxford

Michael A Osborne is an associate professor at the University of Oxford. He co-leads the Machine Learning Research Group, a sub-group of the Robotics Research Group in the Department of Engineering Science. Professor Osborne designs intelligent algorithms capable of making sense of complex data through the usage of techniques such as Machine Learning and Computational Statistics. His work in data analytics has been successfully applied in fields as diverse as astrostatistics, labor economics, and sensor networks. Professor Osborne is also the co-author of the widely publicized research "The Future of Employment: How susceptible are jobs to computerisation?" with fellow Oxford academic Carl Benedikt Frey. They concluded in the paper that about 47% of total US employment was at risk of being automatable. Capgemini Consulting spoke with Professor Osborne to understand Machine Learning and its impact on the business world.

Capgemini Consulting

Capgemini Consulting is the global strategy and transformation consulting organization of the Capgemini Group, specializing in advising and supporting enterprises in significant transformation, from innovative strategy to execution and with an unstinting focus on results. With the new digital economy creating significant disruptions and opportunities, our global team of over 3,600 talented individuals work with leading companies and governments to master Digital Transformation, drawing on our understanding of the digital economy and our leadership in business transformation and organizational change.

Find out more at:
http://www.capgemini-consulting.com/

Rightshore® is a trademark belonging to Capgemini

About Capgemini

Now with 180,000 people in over 40 countries, Capgemini is one of the world's foremost providers of consulting, technology and outsourcing services. The Group reported 2014 global revenues of EUR 10.573 billion. Together with its clients, Capgemini creates and delivers business, technology and digital solutions that fit their needs, enabling them to achieve innovation and competitiveness. A deeply multicultural organization, Capgemini has developed its own way of working, the Collaborative Business ExperienceTM, and draws on Rightshore®, its worldwide delivery model.

Learn more about us
at www.capgemini.com.

Contacts: **Didier Bonnet**, didier.bonnet@capgemini.com, **Jerome Buvat**, jerome.buvat@capgemini.com