Part of the conference series
Breakthrough science and technologies
Transforming our future

Machine learning

Conference report

THE
ROYAL
SOCIETY

Introduction

On 22 May 2015, the Royal Society hosted a unique, high level conference on the subject of machine learning. The conference brought together scientists, technologists and experts from across academia, industry and government to learn about the state-of-the-art in machine learning from world and UK experts. The presentations covered four main areas of machine learning: cutting-edge developments in software, the underpinning hardware, current applications and future social and economic impacts.

This conference is the first in a new series organised by the Royal Society, entitled *Breakthrough Science and Technologies: Transforming our Future*, which will address the major scientific and technical challenges of the next decade. Each conference will focus on one technology and cover key issues including the current state of the UK industry sector, future direction of research and the wider social and economic implications. The conference series is being organised through the Royal Society's Science and Industry programme, which demonstrates our commitment to reintegrate science and industry at the Society and to promote science and its value, build relationships and foster translation.

This report is not a verbatim record, but summarises the discussions that took place during the day and the key points raised. Comments and recommendations reflect the views and opinions of the speakers and not necessarily that of the Royal Society. Full versions of the presentations can be found on our website at: **royalsociety.org/events/2015/05/breakthrough-science-technologies-machine-learning**

Foreword

The advent of machine learning has begun to transform a broad range of economic and social processes, from internet search via voice interfaces, to large-scale data mining, and the finance sector. The combination of large datasets for the training of learning algorithms, combined with cheap powerful computing hardware, has triggered a diverse set of applications for machine learning technologies.

While there is currently a heated public debate on the moral and ethical challenges future artificial intelligence may pose to humanity, it is important to also focus on the real-world value and economic benefits that have arisen from automated machine reasoning. The greatest imminent benefit will, likely be in improved medical diagnosis, disease analysis and pharmaceutical development.

For many attendees at the event the keynote on the medical potential of machine learning had the greatest resonance and clear potential benefit to wider society. In an age of economic challenge, combined with increasing demands on health care provision, there is an urgent need to extract value from the vast datasets that now exist in the life sciences.

A final point arising from this event was the common theme that the UK needs to invest in the advanced skill base required to realise these technologies, if we are to deliver economic impact for the UK. Specifically, the underlying demand for STEM graduates to make use of, and develop machine learning technologies, was a clear priority, in terms of future public policy on this subject.

It is hoped that this form of combined industry, academic, and public-sector conference, as hosted by the Royal Society, can help address the many complex challenges and unlock the potential of machine learning for the benefit of society.

Dr Hermann Hauser KBE FREng FRS

Dr Robert Ghanea-Hercock

Current and future developments in machine learning software

The conference presented state-of-the art developments in machine learning software, including both deep learning and probabilistic models. Although there are fundamental differences between the two techniques, both will complement the other and already probabilistic models are being used to reduce the costs of training complex deep neural networks. As Professor Christopher Bishop FREng of Microsoft Research Cambridge noted, what is ultimately wanted is a system that can achieve great performance on any application.

Neural networks and deep learning

In imitation of their biological counterparts, neural networks are systems of interconnected 'neurons', computational nodes that pass messages back and forth and adapt to changing inputs. In the recent advancements in deep learning, as described by their pioneer Professor Geoff Hinton FRS (University of Toronto and Google), these nodes are placed in a series of hidden layers which sit between the input and output layers.

FIGURE 1

In a neural network, inputs are fed into the input layer at the bottom. Each hidden layer detects combinations of active neurons in the layer below, filtered by the 'weights' of the connections (shown as coloured dots), transforming the input to give a decision at the top. The system is trained by back propagation where the decision is compared to the correct output and the weights adjusted to ensure a more accurate decision in the future.

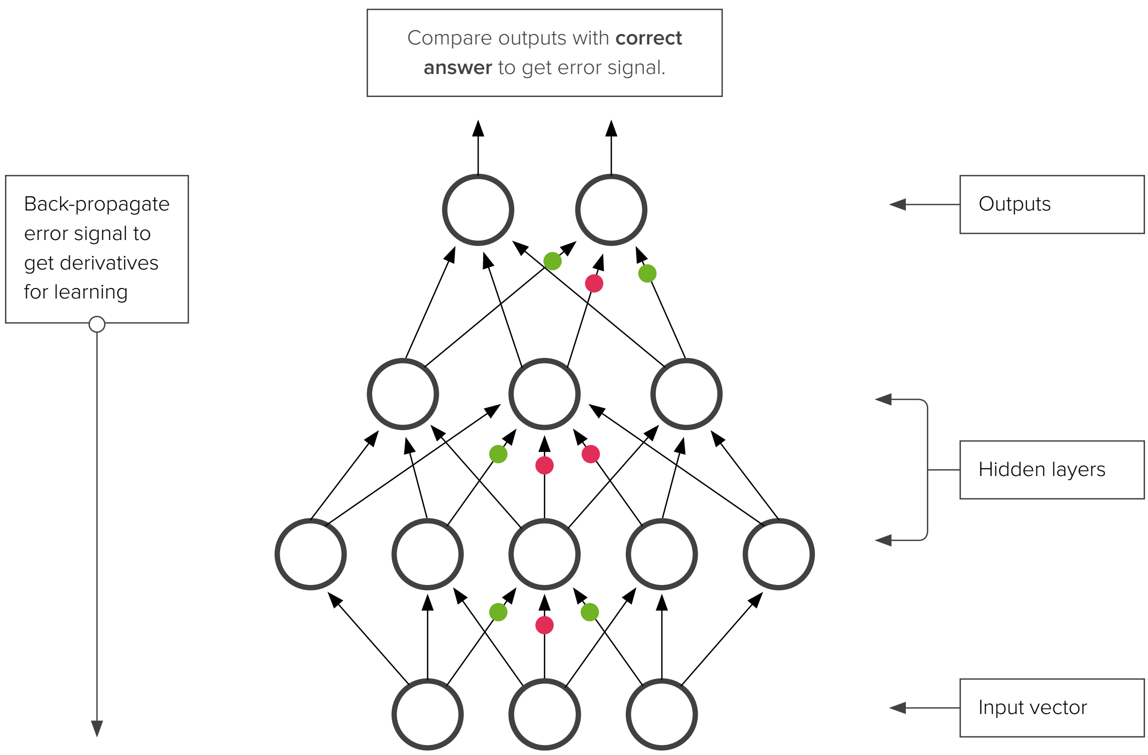

Learning is achieved incrementally through the layered architecture of the network using a process called backpropagation. For example, with object recognition, humans recognise objects through many different features. A deep learning system starts with just a batch of pixels and uses the intervening hidden layers of nodes to group the pixels into elements of the image and then into features it can use for classification. Backpropagation compares the difference between an expected, 'true' output and the actual output of the neural network and uses the difference to work backwards recalculating the statistical model, or 'weights', that the network uses to improve its accuracy.

> "For deep learning, it's very clear that government funding for translational research if anything slowed it down. What was important was curiosity driven research"
>
> **Professor Geoff Hinton FRS, University of Toronto and Google.**

Increased computational power, driven by usage of graphics processor units (GPUs), and increasingly large labelled datasets which can help train the networks, have brought forward advancements in deep learning technology. In just the past few years, Hinton's deep learning techniques have become state-of-the-art in speech, text and image recognition and translation software (see section 'Applications of Machine Learning' for more details).

Probabilistic programming

Many aspects of machine learning and intelligence depend crucially on the representation of uncertainty, the problem of inferring cause from a known effect, which was solved by the 18th century clergyman and mathematician Thomas Bayes. Bayes' rule is used to update the probability for a particular hypothesis in response to new data, and it underlies an increasing number of everyday systems such as search engines (which predict the right web pages to answer a given query).

Probabilistic models use known outcomes to estimate the probability of an event occurring again or of it being the result of a particular cause (the "prior" in Bayes' theorem). Professor Zoubin Ghahramani FRS, the University of Cambridge, believes that probabilistic modelling "offers a framework for building systems which reason about uncertainty and learn from data, going beyond traditional pattern recognition problems".

Although neural networks are not trained in an overtly probabilistic manner and are excelling in pattern recognition and translation, there are still many aspects of learning and intelligence where handling uncertainty is essential, include forecasting, decision making, working with limited or noisy data and automated scientific modelling. However, deriving the probabilistic models and algorithms needed is time-consuming, error-prone and resource inefficient.

Both Ghahramani and Bishop are developing more universal model-based approaches which will make the process more efficient and help systems choose the right algorithms to use. Ghahramani is developing probabilistic programming to make it "easy to write down models and modify them and see which ones are consistent with the data and which ones are not." To reduce expense and resource usage, these models must then evaluate data in as few experiments as possible through a sequential decision-making process he described as "Bayesian optimisation."

> "Many problems in machine learning and AI require the evaluation of a large number of alternative models on potentially large datasets. A rational agent needs to consider the trade-off between statistical and computational efficiency."
>
> **Professor Zoubin Ghahramani FRS, The University of Cambridge.**

In the era of big data, it is not possible to analyse large datasets with hundreds of available models as this is “computationally very expensive”, stated Ghahramani. To solve this problem, he believes machines need to act as “rational agents” that consider the “trade-off between statistical and computational efficiency”. The system he is developing runs each model on a dataset for only a few seconds, then evaluates their performance to forecast how well they will do over a longer timeframe. The machine will then allocate more resources to the most promising model.

Bishop outlined a single development framework for creating a wide range of bespoke models that can automatically generate an algorithm to solve real-world problems. Bishop’s lab is partway to the goal of overcoming an issue common among the many new entrants into the field – how to identify the right algorithm for the problem they’re trying to solve. If, instead, they can generate an algorithm automatically after defining the model that represents the problem, they need not know whether the result corresponds to an algorithm that’s already known and published. If model-based machine learning is a philosophy, the practical application of it is probabilistic programming. A system can be generated automatically in just a few hours from a probabilistic program, which simply describes a model that encodes the assumptions about how the data is generated: simple to describe, transparent, easy to share with others. The hard part, writing the inference code, is done automatically by this probabilistic program compiler.

Both Ghahramani and Bishop have created open platforms for others to use their model-based techniques. Ghahramani’s Automatic Statistician is an automated data analysis and data science tool, which generates new models by combining others like Lego Bricks. Meanwhile, Bishop and his colleagues at Microsoft have created the downloadable tool infer.net, which “provides state-of-the-art algorithms for probabilistic inference from data.”

Artificial general intelligence

Dr Demis Hassabis, of Google DeepMind, is focused on the problem of creating “artificial general intelligence” (AGI), in which machines can learn automatically from raw inputs and operate across a wide range of tasks. If solved, the results could then solve “everything else”.

> “Our mission at DeepMind… [is] a two-step process. Step one – solve intelligence, then step two, use it to solve everything else.”
>
> **Dr Demis Hassabis, Google DeepMind.**

Creating AGI involves two principles – emulating biological reinforcement learning and embodied cognition. In reinforcement learning, a machine finds itself in an environment and interacts with it by using its noisy, incomplete observations from sensory apparatus, to model its surroundings as best it can and respond with the best action to reach its goals. The results are fed back into the cycle. Embodied cognition holds that a truly thinking machine must be embedded in sensory motor data streams.

DeepMind are testing their AGIs in computer games, which provide rich, high-dimensional data streams that require the machines to make decisions and reach goals on an independently built platform with no testing bias. DeepMind’s “end-to-end agents” start with raw pixels and end by deciding what action to take. Assigned to 1980s Atari games, DeepMind’s agents have already done better than professional human players in many, such as Space Invaders and the TORCS racing car simulator. Next up will be other 3D games, strategic games such as Go, then simulators and even real applications for robots.

The next step to AGI will be adding transfer learning, which presents three challenges. The machine has to identify the current environment’s features, working out which can be ignored; re-representing those as abstract concepts; and selecting among them and appropriately applying stored prior knowledge. Humans do this effortlessly but uniquely, and how to do it well in machine learning is unknown. Hassabis thinks abstract concept formation is the key.

Future hardware solutions

The era of intelligent machines, argued Simon Knowles, Chief Technology Officer at UK semiconductor company XMOS, may represent the most significant transition in computer science since the field was invented. In terms of hardware, the move from 'dumb' algorithms to probabilistic data models made up of billions of graphs will require new approaches to how data is addressed and operated upon. Consequently, not only does machine learning require huge amounts of computing power, it will also need the development of new computer architecture to support it and open the door to parallel computing.

Next generation silicon processor for machine intelligence

At the most basic level, all current computing is arithmetic performed on data held in dynamic memory. The computing speed of contemporary silicon processors are limited by two problems which must be solved to achieve more efficient machines: how much data can be kept near the processor logic gates and the power density.

> "Machine learning is surely a massive opportunity and will become a massive market for computing machines. It will justify… the invention of new processors"
>
> **Simon Knowles, XMOS.**

Knowles described how machine learning is beginning to take conventional silicon processors beyond their capabilities. Already, graphic processor units (GPUs) are preferred to conventional CPUs, as they are able to perform much greater numbers of calculations needed for machine learning. However, GPUs were created for graphics and high-performance computing applications where communication between memory and processor units is of less importance than in parallel computing. This is not the case in applications where parallel computing may be needed, such as machine learning, and even the most powerful GPUs only operate at 1 – 2% efficiency. Cramming more computing power onto the chip at the expense of memory is not a solution, as it runs into the problem of having to reduce speed to maintain the power density.

FIGURE 2

An 'Intelligent Processing Unit' will place more memory on the chip, thereby increasing communication speed between memory and the floating points used for calculation ('flops'), giving it higher performance and more power efficiency than a standard GPU.

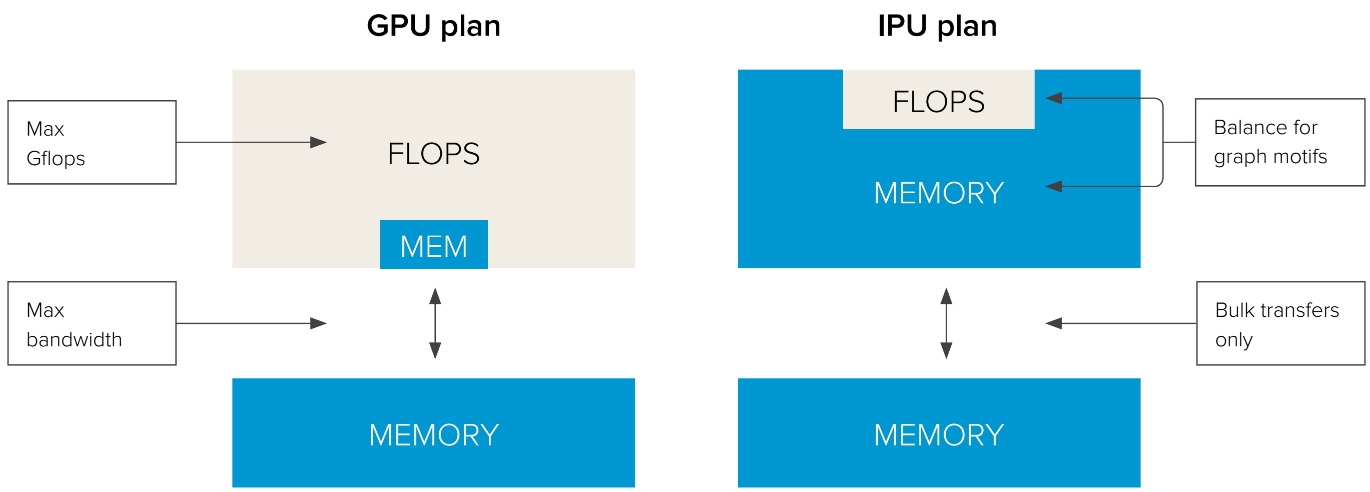

As modern processors are optimised for computing and memory access at a given power density, this low level of efficiency means they are massively underused in machine learning, wasting both power and silicon. If more memory was placed on the chip, it would require less power than needed to go 'off-chip' to fetch it. Furthermore, processing probabilistic data models requires access to data scattered across the memory, which is not how modern computers work and machine learning will require different memory access motifs. Currently, there is "too much emphasis on arithmetic, not enough on communication", said Knowles.

Rather than continue to use modern GPUs which weren't developed for machine learning in mind, a new generation of 'intelligent processor units' (IPUs) are needed with massive available parallelism. The huge market opportunity of machine learning makes it feasible to invent and develop IPUs. The IPU will use most of the chip for data storage, use external bandwidth only for transferring data in bulk, and limit arithmetic units to a much smaller area. Fewer arithmetic units means less cooling is required; this design therefore allows faster clock speeds, uses less power fetching data from off-chip, and delivers much faster overall performance.

Neuromorphic computer architectures

Neuromorphic approaches, mimicking neuro-biological architectures, to machine learning are not new but offer an opportunity to accelerate our understanding of the brain and to reciprocate that knowledge to build better computers.

FIGURE 3

SpiNNaker's design allows systems to be designed to scale as needed: 1 billion neurons at 1,000 neurons per core will mean 18 cores per chip, 48 chips per board, 24 boards per card frame, 5 card frames per cabinet, and 10 cabinets.

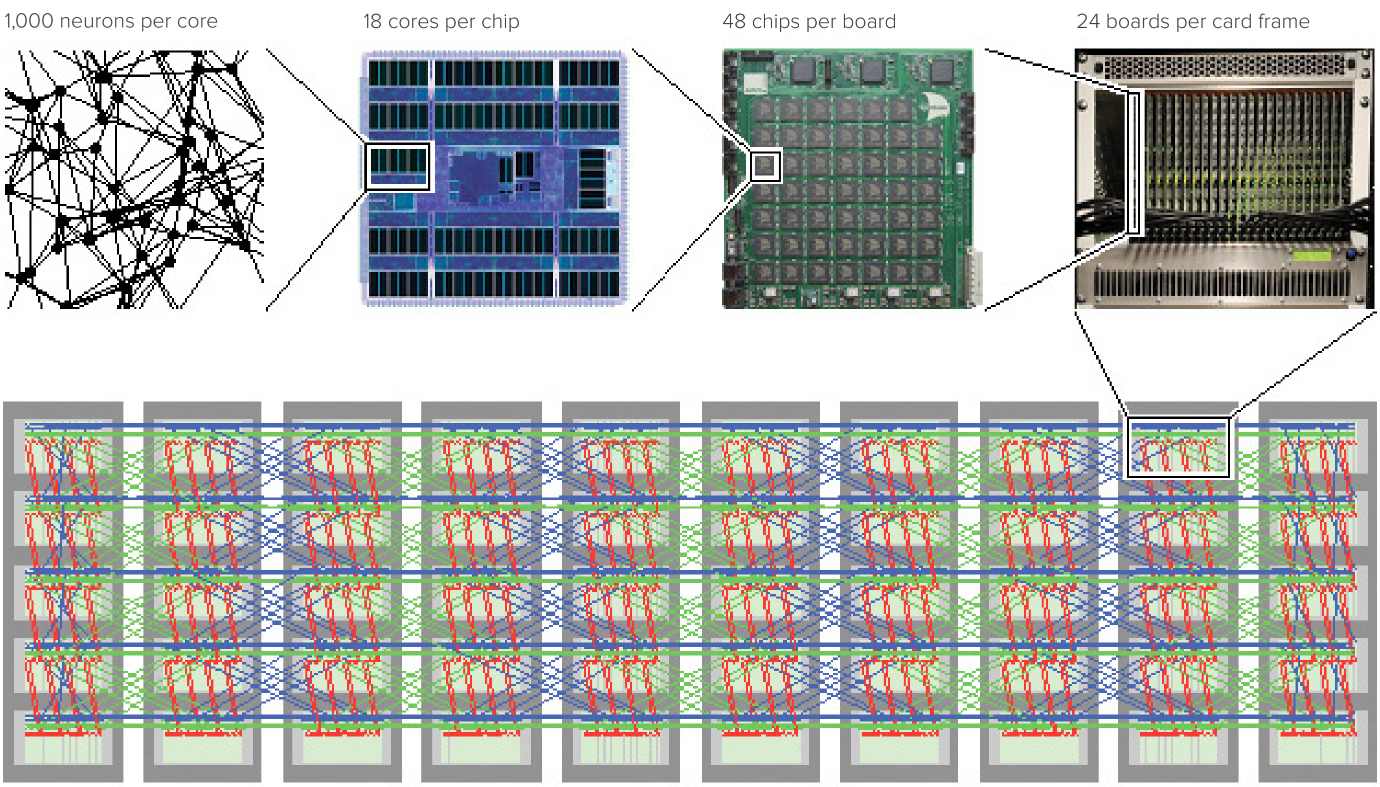

5 card frames per cabinet, and 10 cabinets

The SpiNNaker project, being built as part of the Human Brain project at the University of Manchester by Professor Steve Furber CBE FREng FRS and colleagues, is a neuromorphic architecture inspired by the brain. Its aim, by placing a million ARM mobile phone processors into a computer, is to create a scalable approach to overcoming the parallel computing problems of brain modelling.

SpiNNaker will be a universal machine that can model the many network topologies found in the brain, overcome synchronisation issues found in parallel computing and be energy efficient. Although once complete, it will only represent 1% of computing power of the human brain, its scalability means it should be able to overcome the constraints on computing resources typically faced by cognitive scientists.

Neuromorphic computers have yet to find any applications, but Furber said that recent forecasts were more positive, making it an emerging technology to watch. Other major neuromorphic projects are being developed by IBM (TrueNorth) and at the Universities of Stanford and Heidelberg, all with different optimisations and approaches.

Although SpiNNaker has been developed to better understand the brain, it has been used as a puzzle solver and to analyse big data. Furber believes that once better understanding of how animals store and process data is found, progress in neuromorphic computing, and machine intelligence in general, will start to accelerate further.

Machine learning as the killer app for quantum technology

In December 2014, the UK launched the five-year National Quantum Technologies Programme with £270 million, including £50m for innovation, in public funding to exploit its potential benefit to the UK.

Several applications are suitable for quantum computing: code-breaking, algorithm search and simulations, and it has already been applied to some calculations where classical techniques don’t work, but is machine learning quantum technology’s ‘killer app’? This is the question that was posed by Professor Simon Benjamin from the University of Oxford.

Today’s binary computing is based on switches that are on or off, 0 or 1 (bits). In a quantum device the system can be in two states simultaneously, a ‘superposition’ of quantum bits, or ‘qubits’, and the amount of information that can be stored rises exponentially.

The theory has existed for some time, and now teams such as Benjamin’s are building the hardware needed for quantum machine learning. For many years, it was predicted that cascading errors would doom quantum computing but now it’s being shown that it’s possible to make the machine operate without calculation errors even though the components are constantly misbehaving. Theorists have suggested that if 99.9% of all the qubits operate without errors, then the device as a whole can be made to behave. Ten years ago, 90% was tough to reach in a laboratory but the team at the University of Oxford have now surpassed the 99.9% threshold.

Measuring for errors in a quantum device is challenging. As those familiar with Schrodinger’s Cat might suspect, measuring individual qubits mid-calculation collapses the superposition. How can errors be detected if the process can’t be measured? Teams at IBM, Delft, and UCSB (recently acquired by Google) are experimenting with the use of ‘parity checkers’, interspersed between the qubits. The parity checkers test adjacent qubits collectively and can detect errors as they occur – hinting that while bringing quantum machines under full control remains expensive, it is not impossible.

It’s now reasonable to say that quantum machines will happen, although the time to a finished device is still unknown. Quantum technology is attracting growing corporate investment from companies such as IBM, Microsoft and Google and examples of quantum machine learning have already demonstrated promising results. The challenge laid down to software developers by Benjamin is to design robust algorithms that by their nature can tolerate small amounts of error during execution.

Applications of Machine Learning

Already well established in the financial sector and search-engine queries, machine learning algorithms are now underpinning voice, text and speech recognition software, network security and robotics. These sections summarise some of the applications described during the conference.

Deep learning for image and text recognition

The recent advancements in deep neural networks have already led to step-change advances in text, speech and image recognition, thanks to more powerful computer processors and large enough datasets to train the networks. For example, speech recognition leapt forward when in 2009 Hinton's deep neural networks, pre-trained using backpropagation and powered by GPUs, broke the then record for accuracy. By 2012, Google was using Hinton's systems in the Android voice search and neural networks are now found in all the major mobile phone voice recognition devices.

Hinton's deep convolutional neural network won the 2012 ILSVRC ImageNet competition, which evaluates algorithms for object detection and image classification at large scale, scoring an error rate of 16%, 10% better than any other competitor. By 2015, Google, Microsoft and Baidu reduced the error rate to below 5%, the same as humans.

FIGURE 4

Hinton's neural network gives this picture the caption "a close-up of a child holding a stuffed animal". But does the network 'see' the child holding the animal or does its language model assume 'holding' as a best fit for the sentence?

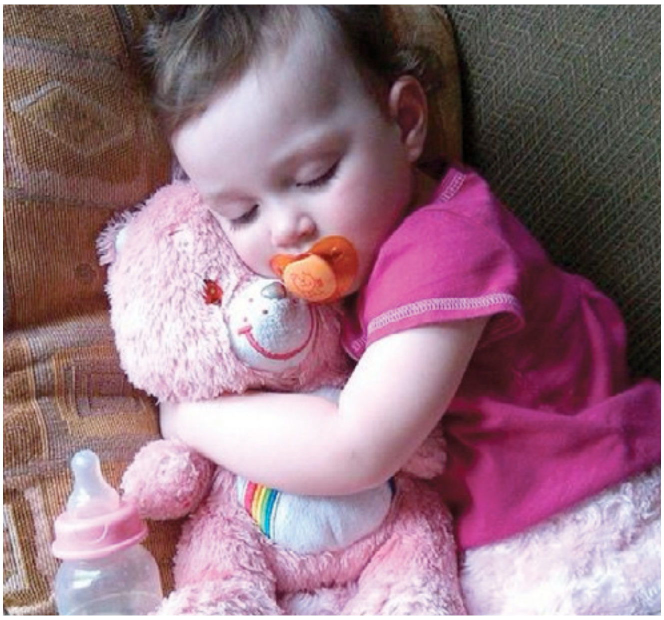

© Oriol Vinyals, Google

Working with Google, Hinton has devised a new method of machine translation. The method in use until recently used a huge database of original and humanly translated documents and statistical analysis to determine the most likely translation between pairs of languages. The new effort, working between French and English, feeds an English sentence to the hidden units one word at a time. The hidden units develop a 'state vector' to represent the thought the sentence expresses, and a 'word vector' for each word. Using that thought vector as a limiting condition, the network then uses probability distributions to create a set of all the possible French sentences. Beginning with random weights, the two networks – encoder (English) and decoder (French) – are trained together using backpropagation to maximise the likelihood of producing a correct translation. On a medium-sized dataset this works about as well as phrase-based translation and is improving much faster. Using more data – such as the 25 – language database of translated European Parliament documents – to jointly train encoders and decoders for many languages is expected to fuel rapid progress.

Hinton is now combining vision and language to produce accurate captions for images, although as the network is trained to recognise objects but not relationships, it's not clear how much it understands the latter.

Hinton believes that converting language into a sequence of thought vectors can be used to model why one set of vectors leads to another, in other words 'natural reasoning'. "This will rapidly change the level at which we can understand documents", said Hinton, although he added that in terms of resources, we are still "a few orders of magnitude" off reaching human-level understanding.

Improving online media services

Machine learning algorithms have been used by online media services, such as Xbox Live and Netflix, for a number of years. Professor Chris Bishop gave two examples from Microsoft's Xbox Live – a matchmaking system that matches video gamers of similar skill levels in multiplayer online gaming, and a recommender which makes personal recommendations for movies, games and music.

The matchmaking system, TrueSkill, works by estimating the outcome of games using skill levels which are modelled as a Gaussian distribution. After each game, an inference algorithm then works backwards to update the system's assumptions about each player's skills and modify the skill distribution. Originally coded by hand over many months, the inference algorithms can now be written in a matter of hours by the probabilistic program compiler Bishop and his team have been developing at Microsoft Research.

The recommender uses features about the user, such as age and gender, and incorporates it with known likes or dislikes to help make personal recommendations for the user. Preferences are viewed probabilistically, and each choice provides the recommender with more data which reduces uncertainty and allows it to be more confident about the user's choices.

Given the millions of data points these systems will encounter, Bishop is "particularly interested in… using approximate inference algorithms that are cheap to operate at scale"

Providing enterprise network security

Combatting enterprise network security threats is a big data challenge. "You have a lot of data from events that are happening in your network and you want to identify which of them are bad," said Dr Miranda Mowbray, HP Labs. The vast number of events mean that security teams can become overwhelmed if their security software raises too many alarms, either true or false, and end up turning off their "very expensive computer network security software".

> "There have been cases where enterprises have bought very expensive computer network software, and then turned it off."
>
> **Dr Miranda Mowbray, HP Labs.**

Traditional anti-virus protection works by identifying suspicious signatures (strings of characters) in the virus' code. As modern viruses are polymorphic, which means the signature looks different each time, machine learning can offer a more robust alternative by identifying anomalous or rare classes of these signatures.

Mowbray gave the example of malwares, which contain an encoded domain name used to communicate instructions between an infected machine and the malware controller. Known domains are blacklisted but the malware authors can bypass this by using domain generation algorithms (DGAs) which generate new batches of domains each day.

One machine learning approach is to use a classifier which is trained using logistic regression on known good and bad domain names. It can then recognise features in the domain name, such as positions of characters or suffixes, which are anomalous. Instead of raising an alarm for each detected domain, the classifier identifies any machines in the network trying to access 'dodgy' domains, so they can be disinfected, increasing robustness and reducing the number of false alarms.

HP are also developing an anomaly detection algorithm for when the malware author uses unknown DGAs. This technique looks for rare 2-length distributions in the domain name, the section between the machine name and top-level domain or country code, e.g. mallard.example.com or eider.example.ac.uk. Targeting this part of the domain name makes it more difficult for malware creators to match their codes to rare 2-length distributions in each network, offering greater security. In one analysis of a single enterprise network, HP were able to uncover 19 DGAs, nine of them unknown, using this algorithm.

Looking to the future, Mowbray said that machine learning will be critical in detecting new threats attacking the Internet of Things. She ended with a cautionary note, saying that although she was unaware of malware creators using machine learning she didn't expect this to remain the case, "you should always ask what the response of the malware authors will be."

Robotic vision

Current robots are able to do a lot of tasks already, but what's missing, said Professor Andrew Davison of Imperial College London, is the perception and control that links robots' manipulators with a vision system.

Davison, whose work with Dyson has led to the recently launched 360 Eye robot vacuum cleaner, described how a robot arriving in a new space must simultaneously build a map of the space and localise itself within that map, a process known as simultaneous localisation and mapping (SLAM). Over the past two decades, developments in robotic vision have been driven by brute-force computing methods using the processing power of GPUs. More recently, better 3D depth imaging and mapping have been brought about by another advancement in gaming technology, Microsoft's Kinect Fusion camera.

The aim now is to build "maps of the world that don't just have raw geometry", said Davison, but have "meaning". The most obvious approach to do this would be to build a dense, bottom-up 3D reconstruction of a space and then apply a 'learned object detector' to label the objects in the room. Instead, Davison's SLAM++ approach attempts to detect objects early and then, once detected, bringing in a high-quality model of that object from a database into the visual reconstruction generated. In other words, the resulting system is more efficient as it updates at the level of objects instead of pixels.

FIGURE 5

The Dyson 360 Eye™ vacuum cleaner, whose visual recognition system was developed in collaboration with Professor Andrew Davison of Imperial College London.

FIGURE 6

The machine learning software gives an accurate prediction to the whole image of a road junction (left), describing it as “a road side that is naming two roads which is at a cross road”, but is unable to identify a subset of the same image (right), which it describes as either a “nematode”, “sea snake”, “sand bar” or “hockey puck”.

In the future, the role of machine learning in robotic vision will be to augment and replace hand-engineered elements of the system, making it more robust and accurate overall. Already the data being obtained from SLAM systems will be able to provide data to train and develop the next generation of deep learning algorithms.

Low power hardware in combination with parallel, heterogeneous algorithms and neuromorphic processor networks will eventually lead to mass-market devices.

Machine learning in telecommunications

It is not only in newer industries, such as online media providers or mobile phone speech recognition, where machine learning is making headway. Even in a large, established sector like telecommunications, usage of machine learning is widespread. Dr Simon Thompson, BT Research, described how his company uses machine learning throughout many parts of its operations, from content recommendation and fraud detection to customer relationships and automated route planning for their field engineers.

However, Thompson described the difficulties of implementing machine learning in older industries such as telecommunications and banking, where there is a problem with legacy systems with thousands of databases connected into a “federated data estate.” “Getting data out of a legacy system is an expensive job”, he observed. Even when acquired, there are a myriad of steps needed to prepare the data before it can be analysed using machine learning, including dealing with noise, missing values, and orphaned records.

The current limitations in machine learning mean we need to view it currently as a tool used by people to manage, process, and deploy solutions to problems. For example, today’s image recognition software are good at recognising whole images but perform less well when looking at sub-components of them. This failure is an artefact of the training regime and objective function of the deep network used and represents a gap between the conception of the system and the reality of it as implemented, i.e. between the idea of what the machine has ‘learned’ and what it can actually do in practice. Overcoming this gap is one challenge for turning state of the art algorithms into systems that deliver real world business value.

Future applications and impacts

The need for better human-machine collaboration

The integration of computers into human life, including in the analysis of big data and the Internet of Things will need a change in the role of computer systems and how they are perceived. Professor Nick Jennings FREng, Chief Scientific Advisor to the UK Government on National Security, declared that computers "need to do an awful lot more than they've done in the past" to help us in our activities. "Agent-based computing" will see machines moving from passive to proactive and become agents that suggest ideas and hypotheses to us.

Machines will need to become better at complex problem-solving and in working in collaboration with humans and will require a "fundamentally different" way of thinking about computers. The wrong-type of interactions, described Jennings, are where humans abdicate their intelligence to a "dumb machine", (e.g. following a sat-nav blindly), and will need to be avoided.

> "Humans will not suffer machines that just say 'do X or do Y'; humans want to hear their doctor robot tell them why they should take the medicine"
>
> **Dr Simon Thompson, BT.**

The smart and effective means of collaboration will be where computers combine with human problem-solving and optimise decision making. This involves using smart algorithms that sift data with experienced humans who use their expertise to make a first decision. The machine then rearranges the data allocation accordingly, and so on.

Jennings gave one example from his career where this approach has been used successfully in disaster reduction and response, with the charity Rescue Global. An unmanned aerial vehicle flies over an area and notifies humans if there is something important that needs attention, who can then initiate a targeted and appropriate response.

> "Medicine is so complex, the challenges are so great… we need everything that we can bring to make our diagnostics more precise, more accurate and our therapeutics more focused on that patient"
>
> **Sir Malcolm Grant, NHS England.**

Healthcare – the low-hanging fruit?

Perhaps the field where machine learning has the greatest potential is in healthcare, where it offers the possibility of combining a rise in quality and efficacy of care and treatment, with improvements in efficiency. Furthermore, the huge datasets held by a healthcare system such as the NHS are optimal for training machine learning algorithms. Quoting US businessman Vinod Khosla, the Chair of NHS England, Sir Malcolm Grant suggested that the "next two decades will see a far greater contribution to human health from data science than they will from bioscience".

Grant described how diagnostics is "the low-hanging fruit" for machine learning in medicine, where automated testing is already replacing humans. Image pattern recognition techniques could play an important role in analysing the forty million imaging investigations run by the NHS each year. Grant gave the example of how, in 2011, researchers at Stanford University developed a machine learning technique for breast cancer diagnosis. Analysing tissue samples, the machine not only replicated human levels of accuracy but identified new structures in the cancerous tissues not seen by pathologists.

Patients and clinicians will need to adapt to the use of computers in diagnosis. The rapid growth in medical knowledge makes it impossible for doctors to keep up to date and more computers such as IBM Watson will be needed to help them. Although one day we may be able to do away with physicians and talk straight to the machine, Grant believes that this eventuality remains a long way off. For the foreseeable future, diagnosis will be undertaken by clinicians and machine working together, with the expertise of the former improving the performance through learning and reiteration of the latter.

Another large dataset being created by NHS England is that of up to 100,000 human genomes sequences by Genomics England. As sequencing becomes cheaper, it is no longer data collection that is the problem, but its interpretation. Again there is a role for machine learning to support the developing fields of personalised medicine and theranostics, where treatment is targeted to the specific needs of the patient.

The unique combination of the NHS' 55 million patient records and our expertise in life science gives the UK a huge opportunity to become the world leader in this field, believed Grant.

Control, understanding, regulation and acceptance
As machine learning becomes established in more areas, questions about quality control and certification arise. Depending on the dataset and task, machine learning can already provide very reliable answers, e.g. in object recognition. However, as these techniques are applied more broadly, certifiable levels of accuracy may be needed, e.g. in clinical diagnosis or chemical analysis.

As the examples of image recognition software given by Hinton and Thompson suggest, greater understanding of how methods such as deep neural networks work and come to decisions is one area of investigation.

Developments in understanding how and why machines come to decisions will not only lead to technical advancements but help humans tolerate and accept them when working in collaboration.

Jennings believed that standards organisations may arise to accredit accuracy with stamps of approval but he felt that these would function at the system level and not at the level of individual machine learning algorithms themselves.

Public resistance to machine learning, especially in an application such as healthcare, may arise from a misunderstanding about "black-box" decision making, suggested Bishop, even if the machine is more accurate than a human. Fears about big data, privacy and ethics will also be present as training machine learning algorithms will require using data from which personal and sensitive information can be derived. You "can't just do anything" with the data, observed Thompson, as it needs to be treated "with great respect." However, although sharing personal data with technology companies will raise "suspicions", Grant believed that successful introduction of machine learning in one field, e.g. driverless cars, will lead to greater public acceptance in others, e.g. healthcare.

The longer-term future
The recent advancements in software, hardware and the access to large datasets upon which algorithms can be trained is rapidly accelerating progress in machine learning and intelligence. A recent survey of leading experts in artificial intelligence, carried out by Professor Nick Bostrom and co-workers at the University of Oxford's Future of Humanity Institute, found that half believed human level machine intelligence could be with us by 2050.

> "The transition to the machine intelligence era is of momentous significance… it's the last invention we'd ever need to make"
>
> **Professor Nick Bostrom, The University of Oxford.**

Machines who are intelligent and curious could identify interesting areas of enquiry to aid human analysis, and may even be better at science than humans themselves, Bostrom explained. Solutions to complex and intractable scientific challenges which require analysis of big data, such as climate change or particle physics, may be beyond the capability of a human lifetime, and artificial intelligence, Hassabis suggested, will be employed as a "meta-solution" to help us.

It may be that artificial intelligence ends up being able to do science better than humans, making it the last great technological advancement we ever make. Given the implications of such "superintelligent" machines, Bostrom believed a field of enquiry around the "control problem" of artificial intelligence must be established as part of long term scenario planning. In part, this may involve ethics boards overseeing responsible use, such as that created by Google DeepMind, but will also need collaboration between the leading minds in mathematics and computer science, both in industry and academia.

Summary

The boom in machine learning has been built upon improvements in hardware design and the rapid growth in large datasets over the past decade.

This has led to a renaissance for techniques such as deep neural networks and backpropagation which require more powerful computers and larger training datasets than were originally available to them, in order to operate successfully. Graphical processor units, not traditional CPUs, are offering the advancements in power needed for machine learning workloads and quantum processors may soon offer even greater advantages.

Already established in the fields of finance and the World Wide Web, machine learning is now pushing advances in speech, text and image recognition, cybersecurity, scientific research and robotic vision. The need to analyse large datasets in the healthcare system makes this application a “low-hanging fruit” for machine learning that the UK can exploit, thanks to the combination of its world leading life science sector and the patient data available through the NHS.

Looking to the future as machine learning pervades more areas, from analysing clinical trial data to the Internet of Things, important questions are raised. How can the technology be used to interrogate huge datasets of personal information in an ethical manner? What effects will machine learning have on future workforces and how do we ensure a ready supply of skilled users and researchers? How will machine agents work collaboratively with humans and, in the longer term, how do we avoid the existential threats posed by an artificial “superintelligence” motivated towards non-human goals?

In order for machine learning to be successfully applied requires intelligent collaboration between machine and user and the narrowing of the performance gap between algorithms operating on training datasets and on real-world data estates. Users will need to understand how to prepare the data appropriately before analysis and how to work with machines in the decision-making process. Examples of how machine learning has succeeded in one area may help lower barriers to acceptance in others.

With a strong research base, adoption of machine learning by industry and a growing and innovative start-up sector, the UK is poised to be at the forefront of the machine learning era.

Acknowledgements

Organisers

Dr Hermann Hauser KBE FREng FRS
Co-founder, Amadeus Capital Partners Ltd
Dr Robert Ghanea-Hercock
Chief Researcher, BT

Chair

Professor Andy Hopper CBE FREng FRS
University of Cambridge

Speakers

Professor Geoff Hinton FRS
University of Toronto and Google
Professor Zoubin Ghahramani FRS
University of Cambridge
Professor Christopher Bishop FREng
Microsoft Research Cambridge
Professor Steve Furber CBE FREng FRS
The University of Manchester
Simon Knowles
Chief Technology Officer, XMOS
Professor Simon Benjamin
University of Oxford
Dr Demis Hassabis
Engineering Director, Google DeepMind
Dr Miranda Mowbray
Research Engineer, HP Labs
Professor Andrew Davison
Imperial College London
Dr Simon Thompson
Head of Practice, Big Data & Customer Experience
Sir Malcolm Grant CBE
Chairman, NHS England
Professor Nick Bostrom
Future of Humanity Institute, University of Oxford
Professor Nick Jennings FREng
Chief Scientific Advisor to UK Government
on National Security

Full biographies and videos of the presentations can be found on the event website: **royalsociety.org/events/2015/05/breakthrough-science-technologies-machine-learning**

The Royal Society is a self-governing Fellowship of many of the world's most distinguished scientists drawn from all areas of science, engineering, and medicine. The Society's fundamental purpose, as it has been since its foundation in 1660, is to recognise, promote, and support excellence in science and to encourage the development and use of science for the benefit of humanity.

The Society's strategic priorities emphasise its commitment to the highest quality science, to curiosity-driven research, and to the development and use of science for the benefit of society. These priorities are:

- Promoting science and its benefits
- Recognising excellence in science
- Supporting outstanding science
- Providing scientific advice for policy
- Fostering international and global cooperation
- Education and public engagement

For further information
The Royal Society
6 – 9 Carlton House Terrace
London SW1Y 5AG

T +44 20 7451 2500
W royalsociety.org

Registered Charity No 207043
Issued: December 2015 DES3814