Deep in Thought: How 'Brainy' Computers Are Changing Our Lives

The Fourth Industrial Revolution is here – and together with technologies such as robotics and quantum computing, Artificial Intelligence (AI) is at the heart of this awesome new age of progress and possibilities.

Deep learning is a great example. Harnessing computer algorithms inspired by how the brain works, this fast-growing field of AI makes it possible to analyse vast amounts of data and detect patterns and features there – **generating insights that can improve and protect lives in areas as diverse as medicine, transport and anti-terrorism**.

UCL is just one of the SES members currently working at deep learning's cutting edge. In co-operation with other SES institutions, it's also putting in place the state-of-the-art supercomputing infrastructure essential to unlocking even more of the massive potential.
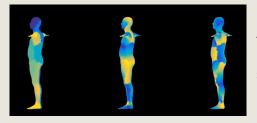
The Rise of Machines

Anyone familiar with 2001: A Space Odyssey and HAL 9000 – the 'sentient computer' in that sci-fi classic – may feel that the future has finally arrived. AI, machine learning, artificial neural networks: terms like these are now common currency as amazing advances (in driverless cars, for example) secure media attention and seize our imagination.

> *"Deep learning computing architectures have an excellent track record in analysing medical and natural images."*
>
> *Dr Delmiro Fernandez-Reyes, UCL*

So how does deep learning fit into the picture?
- **AI** is the overarching field of computer science concerned with developing machines that 'think' in a human-like way.
- **Machine learning** is a branch of AI that replaces 'programming by instructions' with 'programing by examples' – allowing a computer to learn for itself from observed data.
- **Deep learning** is a branch of machine learning, using artificial neural networks that mimic the multi-layered architecture of the human brain.

Click to watch

Deep learning is already firmly embedded in daily life. Google, Amazon, Facebook and Netflix are just a few of the household names already harnessing the technology – e.g. for image and speech analysis. But the potential extends much further and researchers at UCL's Department of Computer Science are fully focused on unlocking it.

A Match for Malaria

Every year sees around 200 million cases of malaria worldwide and around half a million people dying from the disease. As so often in healthcare, early and accurate diagnosis is critical to successful treatment. But testing the sheer number of people at risk from malaria has traditionally posed an overwhelming challenge.

A pioneering international collaboration called FASt-Mal (Fast, Accurate and Scalable Malaria Diagnosis System) aims to deliver a vital step forward. Funded by EPSRC and focused on sub-Saharan West Africa – the world's malaria hotspot – **FASt-Mal will use deep learning to help develop fully automated technology that can detect the presence of malaria in blood samples**. Incorporating robotics, machine learning and computer vision (the use of computers to analyse digital images), the system will scan blood samples for tell-tale signs of the disease much more rapidly and reliably than can be achieved by other means.
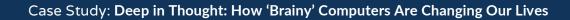

"We already know that several processes, such as counting the total number of red blood cells in a sample, can be carried out using traditional computer vision algorithms," Dr Delmiro Fernandez-Reyes of UCL explains. "**The deep learning part of the FASt-Mal system will target the most difficult task: counting the number of red cells that contain a parasite**. Deep learning computing architectures should be well-suited to this. They have an excellent track record in analysing medical images as well as natural images, while ensuring the privacy and security of patient data." The College of Medicine at Nigeria's University of Ibadan is a key partner in the initiative.

Vision of the Future

As well as developing specific applications, substantial effort is being directed at extending deep learning's capabilities and maximising its impact. UCL's Dr Iasonas Kokkinos is exploring how deep learning can optimise computer vision. "Traditionally, computer vision had been used in a select set of domains, such as security and surveillance, optical character recognition, inspection of industrial production lines, and image search and retrieval," he says. "But over the past five years the advent of deep learning has dramatically improved the accuracy of computer vision systems – enabling applications that were originally treated as sci-fi."

For example:

- In **driverless cars**, deep learning makes it easier for vision systems to track moving objects, to identify pedestrians, road signs and other vehicles, and to interpret the interactions between different elements of a car's environment.

- In **medical imaging**, deep learning makes it easier for X-ray, MRI or ultrasound images to be assessed automatically, quickly and effectively, improving the speed and reducing the cost of disease diagnosis.

Iasonas is currently focusing on combining the different aspects of the challenge involved (3-D surface estimation, object detection etc.) into a single 'universal' vision system. "I'm exploring how all these sub-tasks could be unified, to save time and cut costs," he explains. "That would encourage even greater use of computer vision, enhanced by deep learning, across an even wider range of applications."

The Computing Question

The significance and future potential of deep learning can scarcely be exaggerated. Nor can the key role played by recent advances in computing in terms of delivering the sheer computational power that deep learning depends on. **Most crucial of all has been the emergence of GPU (graphics processing unit) technology over the last few years**.

"A single GPU card can do the same work as ten, twenty or thirty conventional core processing unit (CPU) computers, cutting the time taken by processing tasks from weeks to hours," Iasonas Kokkinos says. "For me, as a researcher working in the field of deep learning, a key question is always 'how many GPU cards can I access and when?' Delays inevitably hold up the whole research process."

Boosting GPU capacity available to UK researchers is therefore vital to enabling continued progress in deep learning, and in many other disciplines. That's why initiatives such as **JADE** (Joint Academic Data science Endeavour) are so significant. Led by the University of Oxford and also involving King's College London, Queen Mary University of London, the University of Southampton and UCL, plus three other UK universities, the consortium will establish a new high-performance computing facility based on leading-edge GPU technology. With a focus on meeting the needs of machine learning, this major development will only deepen the pool of possibilities and broaden the pipeline of breakthroughs that deep learning delivers in the years ahead.

Contacts

Dr Delmiro Fernandez-Reyes
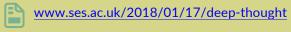Reader in Digital Health and Intelligent Systems

Department of Computer Science
University College London
E: delmiro.fernandez-reyes@cs.ucl.ac.uk

Dr Iasonas Kokkinos
Senior Lecturer

Department of Computer Science
University College London
E: i.kokkinos@cs.ucl.ac.uk

Discover more online

www.ses.ac.uk/2018/01/17/deep-thought

http://www.ses.ac.uk

@SES_Consortium