Artificial Intelligence in the World – Brief Course Outline

- Lecture 1
  - Overview
  - Key questions
  - History
  - Who is doing AI research & development, and why?
- Lecture 2
  - Software
  - Technology components of AI-based systems
  - Key application – Self-driving cars
  - People
    - John McCarthy
  - History
    - Alan Turing
    - From 1945 to 1965
- Lecture 3
  - Key application – Robots
  - People
    - Andrew Ng
  - History
    - From 1965 to 1980
- Lecture 4
  - Key application – Game-playing (chess, Go, Jeopardy!)
  - People
    - Kevin Kelly ("The Inevitable")
      - *Dialogue with Kevin Kelly*
      - *"What makes these things inevitable?"*
  - History
    - 1980 - 1989
- Lecture 5
  - Key application – Speech input and output
  - People
    - Ken Olsen & Gordon Bell (minicomputers)
  - History
    - Bob Metcalfe (Ethernet) & Alan Kay (Alto, "Dynabook")
  - Movies
- Lecture 6
  - Key applications – Virtual Reality & Augmented Reality
  - People
    - Vint Cerf
  - History
    - Doug Engelbart, Bill Gates, Steve Jobs (IA, DOS, Mac)

Lecture 7

Other key applications

a. Business optimization (factory, warehouse, delivery, …)
b. Recommendation engines (Amazon, Netflix, …)
c. Military (drones, aircraft, ships, …)
d. Biology (drug research, medical diagnosis, …)
e. Financial (market prediction, insurance risk, …)

People

Tim Berners-Lee; Larry Page & Sergei Brin

History

From Frank Rosenblatt to ANNs

Lecture 8

Other applications

Language translation

Chatbots

Drones & other UAVs, Unmanned ships

Ethics & dangers of AI

Ethical questions

General Artificial Intelligence (GAI)

Future AI

Ada's Algorithm: How Lord Byron's daughter Ada Lovelace launched the digital age, by James Essinger, Melville House (paper) 2015

With amazingly detailed research Essinger presents the 19th century life of Augusta Ada Lovelace and her father (Lord Byron) and mother (Lady Lovelace). He supports Ada in claiming her rightful place as the first person to comprehend the potential of computers and to write an algorithm for Babbage's unbuilt Analytical Engine. Full of details of the doings and prejudices of the Victorian upper class in which Ada found herself.

The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies, by Eric Brynjolfsson and Andrew McAfee, W.W. Norton, 2014

These two MIT scholars present their understanding of our current computer age as being analogous to the Industrial Revolution (and first machine age) that started with the invention of the efficient steam engine. In particular, they believe that we have only started into this "second machine age" and that the pace of change will only accelerate. Artificial Intelligence is addressed in a full chapter. They propose some changes in education and government policy to deal with the displacement of workers and the growing inequality in our society.

Homo Deus: A Brief History of Tomorrow, by Yuval Noah Harari, HarperCollins 2017

What shall humans do after we have conquered famine, plague and war? Oxford-educated Harari, a lecturer in history at Hebrew University in Jerusalem suggests that we will seek immortality and expanded mental powers by using knowledge of biology and algorithms. This opens up questions of what narrative will guide our future, since the old ones seem to have been made obsolete. With directness, candor and humor that will probably offend most current academics and religious leaders, Harari proposes that we revisit our assumptions about the widely accepted scientific viewpoint and our assumptions about what it means to be human.

Tubes: A Journey to the Center of the Internet, by Andrew Blum, HarperCollins 2012

What's behind the modem or internet connection in your home? Blum, a journalist, goes in search of the *physical* structures that form the internet and finds anonymous buildings in Palo Alto, The Dalles (Oregon), Frankfurt, Amsterdam, London and New York containing fantastic arrays of light-carrying "tubes" where the internet makes its cross-connections. He also explores server farms, undersea cables, and the substructure of New York City to learn where our bits actually go when we surf the web.

The Economist

Autonomous vehicles

Who’s self-driving your car?

The battle for driverless cars revs up

Sep 24th 2016 | From the print edition

But can it fly?

WITH its successful test of robo-taxis on the streets of Pittsburgh last week, Uber has dominated recent headlines on autonomous vehicles. But behind the scenes three groups—technology giants such as Uber, carmakers and a whole fleet of autoparts suppliers—are in a tight race. Each is vying to develop the hardware and software that make up the complex guts of a self-driving vehicle.

A couple of years ago tech firms appeared well ahead in this battle. But, Uber aside, they have dabbed the brakes of late. The recent departure from Google of Chris Urmson, the company’s figurehead for autonomous vehicles and the man who once promised it would put self-driving cars on the road by 2017, is a significant reversal. The recent slimming of the team at Apple that is devoted to building an autonomous electric car, also shows that tech firms are not having it all their own way (though Apple’s possible tie-up with McLaren, a British maker of sports cars and Formula 1 racing team, would be one way to put its carmaking ambitions back on track).

Carmakers, meanwhile, are making more of the running after a slow start. Despite recent safety concerns, Tesla, an electric-car maker, is making progress with its Autopilot system. In 2017 Volvo, which is also working with Uber to get cars to drive themselves, will test self-driving cars by handing them for the first time to a select group of ordinary motorists. And in August, Ford said it would launch a fully-autonomous car, without steering wheel or pedals, for car-sharing schemes by 2021.

All parties recognise that the biggest profits from autonomy will come from producing an “operating system”—something that integrates the software and algorithms that process and interpret information from sensors and maps and the mechanical parts of the car. Tech firms probably have the edge here.

But carmakers and suppliers are not giving up easily. So they are involved in a bout of frenzied activity to keep control of the innards of self-driving cars. In July, for example, BMW, Mobileye, an Israeli supplier that specialises in driverless tech, and Intel, the world's biggest chipmaker, said they were joining forces.

Another strategy for carmakers is to develop autonomous driving in-house. They are hoovering up smaller firms that have useful self-driving technology, notes Andrew Bergbaum of AlixPartners, a consulting firm. Ford has put money into a lidar company (lidar is a type of remote-sensing technology), and into another that sells mapping services. It has also acquired two other firms that specialise in machine-learning and other artificial-intelligence technology.

The losers in this race look likely to be the big parts-makers, whose relationship with their main customers could become strained. Over time carmakers have largely ceded to them the job of developing new technology. If they turn back the clock and reintegrate vertically that may leave less business for the suppliers.

The tech giants still have huge advantages. As well as their financial resources, they are in the best spot to claim the big profits from the operating system. Apple's plans to build a car may be swiftly revived if it buys McLaren. And Google is ahead in machine-learning, the vital element in developing algorithms that will eventually replace drivers. But carmakers are coming up surprisingly fast on the inside lane.

From the print edition: Business

The Economist

Aviation and robots

Flight fantastic

Instead of rewiring planes to fly themselves, why not give them android pilots?

Aug 20th 2016 | From the print edition

THE idea of a drone—an aircraft designed from scratch to be pilotless—is now familiar. But what if you want to make pilotless a plane you already possess? Air forces, particularly America's, sometimes do this with obsolete craft that they wish to fly for target practice. By using servomotors to work the joystick and the control surfaces, and adding new instruments and communications so the whole thing can be flown remotely, a good enough lash-up can be achieved to keep the target airborne until it meets its fiery fate. The desire for pilotlessness, though, now goes way beyond the ability to take pot shots at redundant F-16s. America's air force wants, as far as possible, to robotise cargo, refuelling and reconnaissance missions, leaving the manned stuff mostly to its top-gun fighter pilots. This could be done eventually with new, purpose-built aircraft. But things would happen much faster if existing machines could instantly and efficiently be retrofitted to make their pilots redundant.

Shim Hyunchul and his colleagues at KAIST (formerly the Korea Advanced Institute of Science and Technology) think they can manage just that. They plan to do so by, quite literally, putting a robot in the pilot's seat. As the photograph shows, this robot—called PIBOT (short for pilot robot)—has a human body plan, with a head, torso, arms and legs. The head is packed with cameras, which are thus in the same place as a human being's eyes, and the arms and legs can operate an aircraft's controls, just as a human being would.

Call me George

To design PIBOT, Dr Shim and his colleagues broke the task of piloting down into three areas—recognition, decision and action. They then developed the machine intelligence and sensory software needed for a robot to carry all three out well enough to fly a plane.

The recognition part was fairly easy. Trainee pilots have to learn to ignore irrelevant stimuli and concentrate on the instruments, which is trivial for a robot. And most recognition tasks during flight involve reading simple text displays and markings, tasks for which modern optical-recognition software is more than adequate. For looking out of the cockpit, meanwhile, PIBOT has edge-detection software that recognises features like the horizon and runway markings.

Decision-making is similarly simple to program in. Here, PIBOT works like a standard autopilot, following the rules set down in the handbook of whichever aviation authority has to approve it. Programming in the actions consequent on these decisions, though, was trickier. Every such action—for example, flicking a particular switch or moving the joystick a prescribed amount—has to be expressed as a combination of arm- or leg-joint movements that have to be calculated precisely and then added to the robot's memory.

The first PIBOT, a scaled-down version based on a commercially available 'bot called BioLoid Premium, was demonstrated in 2014. Though just 40cm tall, this had the same articulation as a full-sized device. When strapped into a cockpit simulator with miniature controls, it was able to go through a complete flight sequence, from turning on the engine and releasing the brakes to taxiing, taking off, flying a predetermined route and landing safely at the destination. Crucially, it was then able to do the same in a real, albeit miniature aircraft—though it needed some human assistance with the tricky procedure of landing.

Now, Dr Shim has unveiled PIBOT2, a full-sized version of his invention. This flies a simulator as well as its predecessor did, though it has yet to be let loose in a real cockpit. If it can outperform that predecessor in the landing department, then it will fulfil the United States Air Force's requirement for a "drop-in robotic system" that can be installed quickly without modifying an aircraft—and will do so at a unit cost of $100,000, which is $900,000 less than the cost of converting an F-16 for a trip to the great shooting gallery in the sky.

From an air force's point of view there is a lot to like. PIBOT's autonomy removes the risks of jamming or loss of a communication link that goes with remote control. The robot is immune to g-forces, fatigue and fear, requires neither oxygen nor sleep, needs only a software download—rather than millions of dollars of flight training—to work out how to pilot an aircraft, and can constantly be upgraded with new skills in the same way.

Moreover, Dr Shim sees the military use of PIBOT as just the beginning. It could also provide an economical replacement for a human co-pilot on commercial flights. It could revolutionise ground transport, too—providing, as an alternative to purpose-built driverless cars, the possibility of a robo-chauffeur. Dr Shim says he is already working on a PIBOT able to drive a car, a task which is, he says, "easier in some parts and more difficult in others" than piloting a plane. If successful, this approach could turn millions of existing vehicles into driverless ones quickly and easily. And the owner could still put the robot in the back seat (or even the boot) whenever he wanted to experience the old-fashioned thrill of taking the wheel himself.

Aviation safety

Flight response

An artificially intelligent autopilot that learns by example

Sep 17th 2016 | From the print edition


Just relax and enjoy the view, Captain

ON JUNE 1st 2009, an Air France airliner travelling from Rio de Janeiro to Paris flew into a mid-Atlantic storm. Ice began forming in the sensors used by the aircraft to measure its airspeed, depriving the autopilot of that vital data. So, by design, the machine switched itself off and ceded control to the pilots. Without knowing their speed, and with no horizon visible in a storm in the dead of night, the crew struggled to cope. Against all their training, they kept the plane's nose pointed upward, forcing it to lose speed and lift. Shortly afterwards the aeroplane plummeted into the ocean, killing all 228 people on board.

French air-accident investigators concluded that a lack of pilot training played a big part in the tragedy. As cockpits become ever more computerised, pilots need to keep their flying skills up to date. But pilots are also in short supply. In July Airbus predicted that 500,000 more will be needed by 2035 to keep pace with aviation's expected growth. That means there is pressure to keep aircrew in their cockpits, earning money, rather than in the simulators, taking expensive refresher courses.

Help may be at hand, though, from artificial-intelligence (AI) experts at University College London (UCL). Inspired by the Air France tragedy, Haitham Baomar and his colleague Peter Bentley are developing a special kind of autopilot: one that uses a "machine learning" system to cope when the going gets tough, rather than ceding control to the crew.

Today's autopilots cannot be trained, says Mr Baomar, because they are "hard coded" programs in which a limited number of situations activate well-defined, pre-written coping strategies—to maintain a certain speed or altitude, say. A list of bullet points (which is what such programs amount to) does not handle novelty well: throw a situation at the computer that its programmers have not foreseen, and it has no option but to defer to the humans.

Mr Baomar suspected that a machine-learning algorithm could learn from how human pilots cope with serious emergencies like sudden turbulence, engine failures, or even—as happened to the Air France jet—the loss of critical flight data. That way, he says, the autopilot might not have to cede control as often, and that, in turn, might save lives.

AI takes off

Machine learning is a hot topic in AI research. It is already used for tasks as diverse as decoding human speech, image recognition or deciding which adverts to show web users. The programs work by using artificial neural networks (ANNs), which are loosely inspired by biological brains, to crunch huge quantities of data, looking for patterns and extracting rules that make them more efficient at whatever task they have been set. That allows the computers to teach themselves rules of thumb that human programmers would otherwise have to try to write explicitly in computer code, a notoriously difficult task.

UCL has lots of experience in this area. It was the institution that spawned DeepMind, the company (now owned by Google) whose AlphaGo system this year beat a human grandmaster at Go, a fiendishly complicated board game. The UCL team has written what it calls an Intelligent Autopilot System that uses ten separate ANNs. Each is tasked with learning the best settings for different controls (the throttle, ailerons, elevators and so on) in a variety of different conditions. Hundreds of ANNs would probably be needed to cope with a real aircraft, says Dr Bentley. But ten is enough to check whether the idea is fundamentally a sound one.

To train the autopilot, its ten ANNs observe humans using a flight simulator. As the plane is flown—taking off, cruising, landing and coping with severe weather and aircraft faults that can strike at any point—the networks teach themselves how each specific element of powered flight relates to all the others. When the system is given a simulated aircraft of its own, it will thus know how to alter the plane's controls to keep it flying as straight and level as possible, come what may.

In a demonstration at a UCL lab, the system recovered with aplomb from all sorts of in-flight mishaps, from losing engine power to extreme turbulence or blinding hail. If it were to lose speed data as the Air France flight did, says Mr Baomar, the machine would keep the nose low enough to prevent a stall. The newest version will seek speed data from other sources, like the global positioning system (GPS).

To the team's surprise, the system could also fly aircraft it had not been trained on. Despite learning on a (simulated) Cirrus light aircraft, the machine proved adept with the airliners and fighter jets also available in the database. That is a good example of a machine-learning phenomenon called "generalisation", in which neural networks can handle scenarios that are conceptually similar, but different in the specifics, to the ones they are trained on.

UCL is not the only institution interested in better autopilots. Andrew Anderson of Airbus, a big European maker of jets, says his firm is investigating neural networks, too. But such systems are unlikely to be flying passenger jets just yet. One of the downsides of having a computer train itself is that the result is a black box. Neural networks learn by modifying the strength of the connections between their simulated neurons. The exact strengths they end up with are not programmed by engineers, and it may not be clear to outside observers what function a specific neuron is serving. That means that ANNs cannot yet be validated by aviation authorities, says Peter Ladkin, a safety expert at Bielefeld University in Germany.

Instead, the new autopilot will probably find its first uses in drones. The system's versatility has already impressed delegates at the 2016 International Conference on Unmanned Aircraft Systems in Virginia, where Mr Baomar presented a paper. The system's ability to keep control in challenging weather might see it used in scientific investigations of things like hurricanes and tornadoes, says Dr Ladkin—some of the most challenging flying there is.

From the print edition: Science and technology