Week 2 APRIL 24
Do you want to hear the most epic story ever told?

"A long time ago the atoms in your body were scattered across trillions of miles of otherwise empty space. Billions of years in the past there was no hint that they would eventually come to be configured as your eyes, your skin, your hair, your bones or the 100 billion neurons in your brain. Many of these atoms came from deep inside a star – perhaps several stars, themselves separated by many more trillions of kilometers. As these stars explode, they hurled parts of themselves outward in a flood of scorching gas that filled a small part of one galaxy out of hundreds of billions of other galaxies, arrayed throughout a gaping span of space and time, almost a trillion, trillion kilometers across.

Some of those atoms might have been in the shell of a trilobite, perhaps thousands of trilobites. Since then, they've been in tentacles, roots, feet, blood, and trillions, quadrillions of bacteria in between. Some may have floated in the eyes of creatures that once looked out across the landscape of 100 million years ago. Yet others may have nestled in the yolks of dinosaur eggs or hung in the exhaled breath of a panting creature in the depths of an ice age. For others, this is the first time they've settled into a living organism after drifting through eons of time in oceans and clouds, and part of trillions and trillions of raindrops or snowflakes. Now, at this instant they are all here, making you."

Caleb Scharf and Ron Miller (2017) ***The Zoomable Universe: An Epic Tour Through Cosmic Scale, from Almost Everything to Nearly Nothing***, Atlantic Books (UK)

Caleb Scharf and Ron Miller (2017) ***The Zoomable Universe***, Scientific American (2017), November, Pages 70-74

Noam L. Libeskind and R. Brent Tully (2016) ***Our Place in the Cosmos***, Scientific American, July

History of Artificial Intelligence

The whole idea that forces of nature and the earliest gods might possess an intelligence the equal or greatly exceeding ours, is an ancient idea going back at least to the earliest civilization with a written language – Mesopotamia – five to six thousand years ago and probably much earlier in the case of the treasure trove of cave art in Europe and well before that Africa, seventy to one hundred thousand years ago. So the cover of the New Yorker Magazine's October 23 cover picturing an assortment of mostly human looking robots going about their business – presumably in New York City - carrying lunch bags and boxes, walking robotic dogs, fiddling with a smart phones and carrying a cup of coffee, while one casually drops change in a cup held out by the one human, a street bum, shouldn't come as any surprise. It was the New Yorker Magazine's way of highlighting the potentially hostile invasion of artificial intelligence (AI) driven robots into our life – the subject of one of the articles in this edition by Sheelah Kalhatkar, '*Dark Factory: Can humans adapt to a robot warehouse*?" For those of us who grew up with the many iterations of the Star Wars and Star Trek mythologies, a world populated by robots and strange life forms – or at least as strange as their writers and creators could imagine – the whole idea that humans had some unique and exclusive hold on intelligence and consciousness, was simply at odds with the imagined galaxies they inhabited.

Since the introduction of the first analogue and later digital computers, engineers have dreamed of creating computers and robots possessing intelligence rivaling, or these days, exceeding our own. But it wasn't until silicon-chip based computers became powerful enough to crunch huge amounts of data were reconfigured to work more like carbon-based computers (brains) that artificial intelligence really took off. It was the latter step, the creation of silicon-based equivalents of the widely distributed and overlapping networks so characteristic of the brain that were the key to developing self-learning computers. Even so, the best AI systems, remarkable as they are at beating the most talented humans in Go or chess and reading CT scans or skin lesions, and recognizing faces, these systems have a long way to go before they challenge the broad range of cognitive and social intelligence of which humans are capable. Even so, AI is capable now of running warehouses and possibly soon, whole production facilities of for example, cars and sensibly translating between hundreds of language with a nuance missing from earlier versions. And given those achievements and time, it won't be long before broader cognitive and possibly empathic skills are mastered. Where does that leave us? In the world of employment, what might be left for humans?

Artificial Intelligence (AI) October 20, 2017

Artificial intelligence has been with us since the mid twentieth century and as long as science fiction has been around – and that's a very long time. In the 1980's *Gene Roddenberry* famously created the television series Star Trek from which emerged two sensate robots in the series, *The Next Generation*. The two were created with superhuman strength and number-crunching talents far beyond their human peers, although only one, Lore, possessed anything like human emotion, and he turned out to be the 'bad' guy. The computers imagined for the Starship Enterprise, in their various iterations, included handheld devices, which were mind boggling in their capabilities, including a hand-held medical device for making diagnoses and even repairing wounds with a quick scan. These days such a device doesn't seem so far fetched what with small micro-needle patches designed for diabetics to continuously monitor their blood glucose levels and programmed to subcutaneously deliver the requisite amounts of insulin as needed. In the three decades since the earliest *Star Trek* series, the rapid evolution of microprocessors, much larger memories and sophisticated software made computers much more powerful. And it wasn't long before they proved more than capable of trouncing humans at their most challenging games, beginning with *Chess*, and more recently the far more demanding game – *Go*. These highly publicized triumphs and others stunts such as defeating former champions in the popular game *Jeopardy*, remind us that computers have come a long way.

Now we're in the midst of a new revolution – computers capable of learning on their own without need of software code to cover every eventuality. The architecture and software algorithms of these computers were designed to mimic the widely interconnected and layered networks so characteristic of the brain, and configured to look for patterns and patterns within patterns, and relationships in mounds of data, large enough and complex enough to severely challenge – and exceed - the brightest humans. They also proved to be much less prone to making careless errors.

Humans solve problems in a logical fashion or so we've been led to believe. In medical school and later residency training I was taught to solve diagnostic problems in the traditional textbook sequence beginning with a thorough history, followed by the physical examination, ending with a provisional list of diagnostic possibilities based on the best fit for the history, symptoms and findings, what's most common and what diseases pose the greatest risks for the patient. This initial assessment was followed by selected laboratory and imaging studies, following which the list of diagnostic possibilities was reconsidered before moving on to additional tests and so on through a series of winnowing procedures, to presumably reach the right answer. This methodical and orderly approach isn't wrong – its just not what most seasoned, experienced clinicians do.

Most of the time, at least in the outpatient clinic, the assessment begins the moment I introduce myself to the patient, listen to their opening remarks and watch how they move and express themselves. Many neurological diseases such as Parkinson's disease, essential tremor, myasthenia gravis, amyotrophic lateral sclerosis and myotonic dystrophy, to name but a few, are strongly suggested by the patient's appearance and behavior. For example in the case of Parkinson's disease, a monotone soft voice, flattened facial expression and failure to swing the arm naturally on one side are strong hints of the disease. These initial impressions are based on experience and pattern recognition and include additional hints about the patient's cognitive status and mood. Initial 'impressions' are correct most of the time but certainly not all the time. For example, several neurodegenerative diseases and some drugs are associated with Parkinson-like features and sometimes it takes repeated assessments and further diagnostic studies to sort matters out. Even then it's sometimes difficult to establish a firm diagnosis. So, while it is true the traditional textbook approach to making clinical diagnoses has a place, so also do experience and pattern recognition play increasingly key roles as clinicians mature.

And it is the latter two – experience from seeing all sorts of similar and related problems and pattern recognition to sort them out – which form the basis for how AI works out its own diagnostic criteria based, in the case of AI, on a much larger data set than any one clinician is familiar with - to solve diagnostic problems, and in some instances even suggest the best management options based the results of peer-reviewed studies garnered from ranked journals worldwide. That's what makes AI so powerful an ally, if not so far, a competitor for Family doctors and specialists, with two exceptions.

One excellent example of the role AI can play in diagnosis is the recent study of AI's potential in the diagnosis of skin lesions and specifically recognizing actinic keratosis and melanoma. Dermatologists often employ the **ABCD** criteria for recognizing skin melanomas, which characteristically tend to be **A**symmetrical, with irregular **B**orders, multi**C**olored and greater than 6 mm in **D**iameter. The guidelines work – just not all the time. The authors of the AI study choose not to use the ABCD criteria explicitly but scanned hundreds of dermatological lesions, including many examples of melanoma gleaned from published data sets, for many of which there was pathological evidence to back up the expert dermatological opinions. What happened was astounding. AI turned out to be every bit as good as the best dermatologists in recognizing melanomas and in some instances performed better. And where AI made mistakes, AI learned not to make the same mistake again. And that's without playing by the **ABCD** rules!

Remember too, that it takes a almost a decade or so from beginning medical school to train a dermatologist, a prep time, which has to be repeated for each and every dermatologist. On the other hand, with AI, once trained, it should be relatively easy and cheap to duplicate, and as the decision making power of AI grows with experience, so will the benefits, which could be shared with other linked devices and professionals. Of course it's not that simple – there's far more to dermatology than spot-checking skin lesions. For now AI is best thought of as a potential aid to dermatologists and possibly extending some aspects of dermatological expertise to regions of the country where there's no easy assess to dermatologists.

In a similar fashion AI has proven to be the equal of or better than the best radiologists in recognizing specific lesions in MRI and CT scans and even plain old X-rays. Again it's a matter of beginning with large data sets – i.e. CT scans of the chest – with known lesions – and allowing AI to develop it's own criteria for recognizing lesions, the criteria for which often extend well beyond what human observers are capable. And as with AI and malignant melanoma, the important thing is that AI simply gets better and better with larger data sets, and learns from its mistakes. Diagnositic radiology is one of the highest paying specialties in medicine, but its not where I would be heading these days, what with AI looking over my shoulder and recent evidence that radiologists, like AI use pattern recognition to pickup and identify specific lesions – its just that AI is proving better than all but the most experienced of radiologists and even they are losing ground to AI.

The latter is the key to the success of AI and why AI has such an advantage over humans in analyzing large data sets – AI learns and has the capacity to handle enormous data sets far beyond what any single human or group of humans can handle. It seems that everyday new applications are found for AI (see references below) - whether analyzing and making decisions in seconds in the financial sector, analyzing big-data in biological research involving complex problems in chemistry and genetics and even designing and supervising new projects. AI increasingly has become a powerful and even essential partner in the worlds of particle physics, astronomy, business and even police work. For example AI driven facial recognition software programs play an increasingly important role surveying footage from London England's thousands of public video cameras to identify and track terror suspects on the loose.

In the near future AI will surely compete with lawyers, financial advisors, real estate agents, and all manner of 'white collar' jobs that involve making sense and keeping track of enormous data sets, which have to be constantly updated. In short AI is here to stay and has huge implications in shaping future job markets. Stay posted! The coming revolution threatens jobs well beyond the blue-collar set President Trump allied with in the 2016 election. AI is coming to you – but not necessarily for you, especially if you're tech savvy and nimble enough to move with the times.

AI and Robots in Industry

Sheelah Kolhatkar (2017) ***Dark Factory: The robotics revolution is changing what machines can do. Where do humans fit in?*** The New Yorker, October 23, Pages 70-81

This well written article highlights the role of robots in warehousing and factories in the United States and China and the resulting impact on the workforce of those countries. One example focuses on ***Symbotics***, a Boston based company, which focuses on robotizing warehouses. And not content with updating existing warehouses, ***Symbotics*** markets fully automated warehouses and looking beyond to warehouses to designing entirely robotic factories. If I worked at the Ford plant or any other automative plant, I would be very worried. But lets look at the warehouses first.

"Simply automating a legacy warehouse, however is a halfway measure as a visit to ***Symbotic*** makes clear. … it sells fully automated warehouse systems to large retail chains, and the new warehouses resemble the old ones about as much as a Telsa resembles a Model T. The company's twenty-thousand-square foot test center is a large cube of interlocking green, yellow, and white steel shelving, tracks, and cages that extend from the floor almost to the celing. There are no isles for lifts to pass through, and no stations for human product pickers. There is no space inside the matrix for people at all."

"Robotic arms unpack pallets of tomato sauce, toilet paper, and soda, and place them on a blue conveyer belt, where they are carried deep into the storage cage. A fleet of little green robots that look like race cars in a Pixar film come to life and zoom inside the cage on dedicated tracks, emitting high-pitched whirring sounds. They collect the cases of products and stash them on shelves until they are needed. Then, an algorithm directs the little car-bots to go back in and bring the desired products out."

These fully automated warehouses, "store more products in a smaller space closer to their retail outlets, requiring less trucking. The robots don't need light to operate" … while requiring a third less energy … and "reducing labor costs by eighty percent … and could run twenty-four hours a day," every day with a minimal amount of human intervention for maintenance. Firms such as ***Walmart*** and ***Coke*** use Symbotic warehouses now, with Pepsi and other large companies, soon to follow.

If fully robotic warehouses work, why not fully robotic manufacturing plants? That's exactly the direction China is taking these days, with a heavy push and investment from the government and not much concern about the workers, except to say that most if not all will find jobs in the service industry. The Chinese attitude is based on the fact that their competitive advantage for decades was based on the low wages; however these days labor costs have risen substantially, and if they hope to continue to compete and win in the markets, they need to drastically reduce labor costs and fully automate using AI driven robots to do the work. Given the size of China and it's growing clout, whatever our Eastern giant does, so must we – or do we?

October 19, 2017 From the "National Post" and Nature

Go is a extremely complex game to master – far tougher than chess. Top human players characteristically rely heavily on intuition and feelings to choose among a nearly infinite number of board positions, making the game extremely challenging for the AI community. Two years ago ***Google DeepMind***, a division of the company devoted to developing AI, created ***AlphaGo***, which was programmed with millions of moves based on past 'masters', which managed to beat the best human players including ***Lee Sedol.***

Then more recently ***Google DeepMind*** created the next generation of ***AlphaGo***, ***AlphaGo Zero***, which managed to teach itself the game by playing against itself millions of times and figuring out within 40 days the underlying principals guiding the moves in the game, something that took humans *hundreds* of years to work out, and went on to exhibit what its creators called, '*genuine moments of creativity''*.

Demis Hassabis, co-founder and chief executive of ***DeepMind***, said the program was so powerful because it was "*no longer constrained by the limits of human knowledge*". He believes that if applied to big health problems such as defeating ***Alzhiemer's*** it could, in a matter of weeks, come up with cures that would take humans hundreds of years to find. "Ultimately we want to harness algorithmic breakthroughs like this to help solve all sorts of pressing real world problems," he said.

"If similar techniques can be applied to other structured problems, such as protein folding, reducing energy consumption, or searching for revolutionary new materials, the resulting breakthroughs have the potential to drive forward human understanding and positively impact all of our lives."

DeepMind has already begun using ***AlphaGo Zero*** to study protein folding and has promised it will soon publish new findings. Misfolded proteins (prions) are responsible for several devastating diseases, including ***Alzheimer's*** and ***Parkinson's***.

Technology companies are increasingly moving into health. Last year, ***Microsoft*** announced it planned to beat cancer within 10 years after launching several projects to "hack" the body. ***Calico***, a secretive Google arm, is also investigating ways to extend human life and even stop ageing altogether.

Details of the AlphaGo Zero test were published in the journal Nature.

AI: Sensate or Not

One of the defining moments in the Star Trek series was the episode when 'Data', the sole robotic member of the crew on the starship Enterprise was subjected to a trial to determine whether he was sensate and therefore entitled to all the rights and privileges extended to his human crew mates or simply a silicon-based robot, highly intelligent though he was, but in the end a robot and not entitled to be treated as a human would be. For those of you who don't know the answer to this challenge, be patient, the answer comes at the end of this essay.

With AI poised to take on more and more big data and more importantly increasingly capable of autonomous learning and problem solving, AI tests the limits of silicon-based intelligence and the limits of what it means to be human. AI is in its infancy and already demonstrated an astounding talent for reading human emotions and intentions, autonomously taking on tasks hitherto the province of humans, such as driving cars, flying airplanes, operating drones, devising, carrying out and analyzing research projects, composing music in the style of major composers of the past and present, creating entirely novel music, and topping off AI's many talents was the writing of poetry indistinguishable from poems created by human writers.

This talent for self-learning or what some in the field call 'deep learning' raises serious, sometimes hostile and even amusing questions about the borderlines between carbon-based intelligence – that's us – and silicon-based intelligence (AI). The arguments and counter arguments on display in the fictional court in Star Trek centered on questions about what's uniquely human about humans. Clearly the borders are fraying.

It wasn't so long ago that many of our species, gave little credence to the idea that animals possessed cognitive talents and a sense of altruism akin in nature, if not degree, to our own. However recent work by primatologists make clear that the great apes including chimpanzees are more than capable of altruistic behavior and to wit, provide long-term-care for members in their band. And these talents aren't limited to apes – they extend to whales, porpoises, elephants and even crows to name a few of our fellow species. Curiously enough, American courts recently extended rights somewhat akin to those of humans, on chimpanzees as a means of bringing to an end the sad saga of sometimes appalling medical experiments carried out on chimpanzees by the scientific and medical establishment in the United States, and I might say elsewhere as well. Our reluctance to admit that moral behavior might extend to other animals is more than mirrored by similar ignorance and prejudice toward members of our own species, including a reluctance to extend the rights and privileges we enjoy, to those of differing color, culture and beliefs, women, the underprivileged and those with differing sexual orientations. The list goes on.

Of course it's a whole other matter when it comes to inanimate AI. Few would argue that robots and AI warrant human rights – yet. Among the reasons for not doing so, even in a future when AI might well rival the comprehensive array of cognitive powers possessed by humans, is the human aversion for accepting silicon-based machines - without a soul or consciousness - as equals!

It's a specious argument like so many philosophical arguments, which lead nowhere. Speaking as a physiologist, when the heart stops, whatever soul might be lurking in me, and any awareness I might have of the moment is surely gone within minutes when my brain stops working. So are humans' conscious beings? Of course we're conscious when we're awake – just not so much as when we're asleep. But even when we're fully awake we're aware of only the tiniest fraction of what floods our sensory system or the inner workings of our brains. As I type this document I'm aware of the sound of the keyboard, the fleeting changes in the text on my computer screen and constant checking to see whether what I'm typing makes sense, but little else except the faint hum of a fan in the background. And thank goodness, I'm not aware of all that's going on for I would surely be overwhelmed and unable to type this manuscript.

Now lets return to the question of whether AI might qualify as a sensate being, if not now, then in a more fully developed form in the future. I would argue that questions about whether AI possesses '*consciousness*' and '*awareness*' are specious questions too. What surely matters is whether AI is capable of mastering the nuances of symbolic language, elements of social intelligence such as reading human intentions and feelings, the complexity of human relationships and developing a moral sense. If AI acquired these traits some time in the future – then whether the machine is conscious or not is surely irrelevant – AI will have achieved the prime attributes of sensate humans. So far Google AI is capable of translating and mastering many of the nuances of symbolic language, reading human feelings and recognizing human faces and the trajectory is that much of the rest is within Google's AI grasp within a generation or two. That really is scary.

There's another issue here too. Complex carbon-based life with intelligence the equivalent of our own might well emerge elsewhere in the universe, as is likely to be the case, given the billions of planets potentially capable of supporting carbon-based life. But what if life elsewhere in the universe turns out not to be carbon, but silicon-based. And if so, what will we make of it? Will we accept them as equals? That's really the question we should be asking. Maybe AI is a life form evolving in our midst, admittedly a machine, and created by us, but which might become our equal or even surpass us in the future.

On a more humorous note Richard Yonek author of '***Hearth of the Machine: Our Future in a World of Artificial Emotional Intelligence***, (2017) imagines a moment in the future when a father is confronted with a daughter who's determined to marry an AI robot.

Something to think about don't you think?
William F. Brown August 2, 2017

PS

After vigorous debate - a debate in the 1980's, which anticipated in so many ways the current fears surrounding AI today - the court decided unanimously that Data was indeed 'sensate'. And as such his inner workings could not be subjected to probing without his full consent. Smart court, wouldn't you agree? It's a debate that also played out in American courts recently over the competing rights of investigators to use chimpanzees and other primates for experimental reasons. Here too, the chimps won out and are now in retirement homes, but no longer subject to invasive surgery and studies of infectious diseases such as various retroviruses and hepatitis.

Readers shouldn't be put off by my references to science fiction. Many physicists are big fans of science fiction including Stephen Hawking who among others, points out that what was science fiction yesterday, all too often turns into reality today or in the future. Remember too, that the writers of Star Trek consulted with many physicists before writing their material.

Quotes from Yuval Noah Harari (2015) ***Homo Deus: A Brief history of tomorr***ow, McClelland & Stewart, Pages 324-325

“David Cope is a musicology professor at the University of California in Santa Cruz. He is also one of the most controversial figures in the world of classical music. Cope has written programs that compose concertos, chorales, symphonies and operas. His first creation was EMI (Experiments in Musical Intelligence), which specialized in imitating the style of Johann Sebastian Bach. It took seven years to create the program, but once the work was done, EMI composed 5000 chorales a al Bach in a single day. Cope arranged a performance of a few select chorales in a music festival at Santa Cruz. Enthusiastic members of the audience praised the wonderful performance, and explained excitedly how the music touched the innermost being. They didn’t know it was composed by EMI rather than Bach, and when the truth was revealed, some reacted with glum silence, while others shouted in anger.”

“EMI continued to improve, and learned to imitate Beethoven, Chopin, Rachmaninoff and Stravinsky. Cope got EMI a contract, and its first album – *Classical Music Composed by Computer* – sold surprisingly well. Publicity brought increasing hostility from classical-music buffs. Professor Steve Larson from the University of Oregon sent Cope a challenge for a musical showdown. Larson suggested that professional pianists play three pieces one after the other; one by Bach, one by EMI, and one by Larson himself. The audience would be asked to vote who composed which piece. Larson was convinced people would easily tell the difference between soulful human compositions, and the lifeless artifacts of a machine. Cope accepted the challenge. On the appointed date, hundreds of lecturers, students and music fans assembled in the University of Oregon’s concert hall. At the end of the performance, a vote was taken. The result? The audience thought that EMI’s piece was genuine Bach, that Bach’s piece was composed by Larson, and that Larson’s piece was produced by a computer.”

“Critics continued to argue that EMI’s music is technically excellent, but that it lacks something. It is too accurate. It has no depth. It has no soul. Yet when people heard EMI’s composition without being informed of their provenance, they frequently praised them precisely for their soulfulness and emotional resonance. "

"Following EMI's successes, Cope created newer and even more sophisticated programs. His crowning achievement was Annie. Whereas EMI composed music according to predetermined rules, Annie is based on Machine Learning. Its musical style constantly changes and develops in reaction to new inputs (much as Jazz improvises?) from the outside world. Cope has no idea what Annie is going to compose next. Indeed Annie does not restrict itself to music composition but also explores other art forms such as haiku poetry. In 2011, Cope published *Comes the Fiery Night: 2,000 Haiku by Man and Machine*. Of the 2,000 haikus in the book, some are written by Annie, and the rest by organic poets. The book does not disclose which are which. If you think you can tell the difference between human creativity and machine output, your are welcome to test your claim."

Pages 325-326

"In September 2013 two Oxford researchers, Carl Benedikt Frey and Michael A. Osborne, published '*The Future of Employment'*, in which they surveyed the likelihood of different professions being taken over by computer algorithms within the next twenty years. The algorithm developed by Frey and Osborne to do the calculations, estimated that 47 percent of US jobs are at high risk."

They calculated the probabilities of losing their jobs, for several vocations by the year 2033 (page 326)

Vocation	**% probability for losing their jobs**
Telemarketers and insurance Underwriters	99
Sports referees	98
Cashiers	97
Chefs	96

Waiters	94
Paralegal assistants	94
Tour guides	91
Bakers	88
Bus drivers	89
Construction laborers	88
Veterinary assistants	86
Security guards	84
Sailors	83
Bartenders	77
Archivists	76
Carpenters	72
Lifeguards	67
Archeologists	0.7

Addendum

AI Predicts Heart Attacks and Strokes More Accurately Than Standard Doctor's Method

By Eliza Strickland
Posted 1 May 2017 | 17:00 GMT

Here at The Human OS, we are slightly obsessed with matchups between artificial intelligence and doctors.

In many experiments (though not yet in many clinics), AI systems are showing great promise in diagnosing diseases, analyzing medical images, and predicting health outcomes. They've even performed better than human doctors in certain tasks like surgical stitching and diagnosing autism in infants.

Now, in the latest win for AI medicine, researchers at the University of Nottingham in the UK created a system that scanned patients' routine medical data and predicted which of them would have heart attacks or strokes within 10 years. When compared to the standard method of prediction, the AI system correctly predicted the fates of 355 more patients.

Predicting these cardiovascular events is a notoriously difficult task. In a recent paper, published in the journal *PLOS One*, the researchers note that about half of all heart attacks and strokes occur in people who haven't been flagged as "at risk."

"The leap from research studies to applications in clinical care will happen over the next five years."
—Stephen Weng, University of Nottingham

Currently, the standard way of assessing a patient's risk relies on guidelines developed by the American Heart Association and American College of Cardiology. Doctors use these guidelines, which focus on well-established risk factors such as high blood pressure, cholesterol, age, smoking, and diabetes, to shape their counsel and treatment for their patients.

To make a system that could do better, researcher Stephen Weng and his colleagues tested several different machine learning tools on medical records from 378,256 patients across the UK. These records tracked the patients and their health outcomes from 2005 to 2015, and contained information on demographics, medical conditions, prescription drugs, hospital visits, lab results, and more.

The researchers took 75 percent of the medical records and fed them into their machine learning models, which set out to find the distinguishing characteristics of those patients who experienced heart attacks or strokes within the 10-year span. Then Weng's group tested the models on the other 25 percent of the records to see how accurately they'd predict heart attacks and strokes. They also tested the standard guidelines on that subset of records.

Using a statistic in which a score of 1.0 signifies 100 percent accuracy, the standard guidelines got a score of 0.728. The machine learning models ranged from 0.745 to 0.764, with the best score coming from a type of machine learning model called a neural network.

While the machine scores may not sound like a resounding triumph, when translated into human terms the significance becomes clear: The neural network model predicted 4,998 patients who went on to have a heart attack or stroke out of 7,404 actual cases—355 more than the standard method. With those predictions in hand, doctors could have taken preventative measures such as prescribing drugs to lower cholesterol.

Weng says the AI medical tools being tested in labs today will soon boost clinicians' accuracy in both diagnosis and prognosis. "The leap from research studies to applications in clinical care will happen over the next five years," he says.

What might that look like in practice? Weng pictures busy primary care doctors using AI tools that have been trained to recognize patterns. "Then the algorithm can look through the entire patient list, flag this up, and bring this to the attention of the doctor," he says. "This could be done with the patient sitting in front of them during a routine appointment, or in a systematic screen of the entire list." While Weng notes that similar clinical decision support software already exists, he says those systems don't make use of AI pattern recognition, which could provide far more accurate results.

Before AI comes to your doctor's office, however, the technology will have to get past major regulatory hurdles. "The key barrier to implementation will be managing privacy and patient confidentiality issues, with computer algorithms trawling through vast amounts of patient data which contain confidential and sensitive medical information," Weng says.

In addition to coping with those privacy concerns, any AI technology will have to deal with regulators' wariness of medical machines that make their own decisions. With all that red tape looming, one wonders: What would a machine learning tool predict about its own chances of gaining approval?

Can Deep Learning Help Clinicians Predict Alzheimer's Disease?

By Amy Nordrum
Posted 15 Jun 2016 | 16:00 GMT

There's no clinical test for Alzheimer's disease, so physicians diagnose it by conducting assessments of patients' cognitive decline. But it's particularly difficult for them to identify mild cognitive impairment (MCI), an early stage of dementia when symptoms are less obvious. And it's even harder to predict which MCI patients will develop Alzheimer's disease (not all of them do).

So it makes sense that researchers might try to apply deep learning to this challenge. "There's a lot of interest in some type of test that would say, 'This person will go on to develop Alzheimer's and this person will not,'" says Pamela Greenwood, a psychology professor at George Mason University, in Fairfax, Va.

Recently, collaborators from Harvard University, Massachusetts General Hospital, and China's Huazhong University of Science and Technology designed a program that combines fMRI brain scans with clinical data to make this prediction. They presented the work, which has not yet been published, in May at the IEEE International Conference on Communications in Kuala Lumpur, Malaysia.

"We try to find the disease at its very early stage," says Quanzheng Li, principle investigator from Massachusetts General Hospital's Center for Clinical Data Science. "A lot of people try to use traditional machine learning to do this, but the result is not that great because it's a very difficult problem to solve."

After initial tests, they say their deep learning program, when paired with a special fMRI dataset, is about 20 percent more accurate than other classification methods using a more basic dataset. However, when those traditional classifiers also used the special dataset, they showed similar gains in accuracy.

Javier Escudero, a biomedical engineer at the University of Edinburgh, says that means the new program may not be all that much better than the old ones; it may have just been working with better data.

If that's true, then other experts who want to apply deep learning to diagnosing Alzheimer's disease may want to take a close look at the data they incorporate into their analyses. According to this latest work, fMRI scans that are processed to show relationships between areas of the brain provide a more nuanced view of the condition than those which merely record measurements over time.

For now, the Harvard-led team is among the first to try to combine fMRI scans and deep learning into a program that could predict an MCI patient's chance of developing Alzheimer's disease. The fMRI scans used in their analysis were taken when patients were at rest. As with any fMRI scans, they reveal where electrical signals are flashing in the brain and how these areas relate to one another.

The term for this relationship is functional connectivity, and it changes as patients develop MCI. This is because signals rely on the flow of oxygen to neurons, but the accumulation of tau proteins in Alzheimer's disease patients strangles these neurons, causing regions of the brain to atrophy.

The group wanted to see if they could use these changes in functional connectivity to predict Alzheimer's disease. They began with data from 93 MCI patients and 101 normal patients provided by the Alzheimer's Disease Neuroimaging Initiative. Based on a time series of 130 fMRI measurements taken from each of 90 regions within participants' brains, the researchers could tell where signals flashed over a period of time.

Next, in a crucial step, the group processed this dataset to create a secondary measurement of the strength of these signals in brain regions relative to each other. In other words, they constructed a map of functional connectivity that showed which areas and signals were most closely related to each other.

Lastly, the team built a deep learning program that could interpret the strength of these patterns, and, when combined with clinical data on age, gender, and genetic risk factors, predict whether a person would develop Alzheimer's disease.

In the end, the team says, its program—using the specially processed dataset of functional connectivity—could predict whether the patients in their cohort would progress to Alzheimer's disease with nearly 90 percent accuracy.

Li says the program is almost accurate enough to be helpful in a clinical setting. "When the method reaches about 90 percent, it's very useful," he says. "We are not there yet but we are very close."

Even with laboratory techniques such as testing for too many proteins in cerebrospinal fluid, experts' predictions of which MCI patients will develop Alzheimer's disease are only about 65 percent accurate. That means some people go undiagnosed, while others worry needlessly about developing the disease.

But after reviewing a draft version of the Harvard team's research, Dinggang Shen, who has done similar work on cognitive computing at the University of North Carolina at Chapel Hill, is skeptical.

"Nobody in the field can get to 80 to 90 percent," he says. "That's impossible, just based on this method." (The authors admit that there were several typos in an early draft shared with Shen, but insist that the accuracy is correct).

The results represent a roughly 20 percent improvement over other classifiers they tested, which used only time series fMRI data, and not the functional connectivity map. However, the accuracy of those other classifiers improved by about 16 percent when they also used the functional connectivity map.

The takeaway, Escudero says, should be that programs which assess functional connectivity, or the strengths of signal associations within the brain, are much better at predicting patients' likelihood of developing

Alzheimer's disease than those which only measure brain signals over time. "It seems clear that the biggest gain comes from the consideration of the connectivity data," he says.

This newest experiment is part of a wide range of efforts to apply deep learning or artificial intelligence to assist clinicians with complex decisions. Perhaps most famously, IBM's Watson aims to guide doctors faced with mountains of medical records and many volumes of literature.

George Mason's Greenwood provides important perspective, pointing out that because there's still no cure for Alzheimer's disease, the usefulness of any such prediction tool is limited for this particular disease. The new program would also need to undergo thorough peer review and far more testing before it could ever be used in a clinical diagnosis.

IBM Watson Makes a Treatment Plan for Brain-Cancer Patient in 10 Minutes; Doctors Take 160 Hours

By Eliza Strickland
Posted 11 Aug 2017 | 15:00 GMT

In treating brain cancer, time is of the essence.

A new study, in which IBM Watson took just 10 minutes to analyze a brain-cancer patient's genome and suggest a treatment plan, demonstrates the potential of artificially intelligent medicine to improve patient care. But although human experts took 160 hours to make a comparable plan, the study's results weren't a total victory of machine over humans.

The patient in question was a 76-year-old man who went to his doctor complaining of a headache and difficulty walking. A brain scan revealed a nasty glioblastoma tumor, which surgeons quickly operated on; the man then got three weeks of radiation therapy and started on a long course of chemotherapy. Despite the best care, he was dead within a year. While both Watson and the doctors analyzed the patient's genome to suggest a treatment plan, by the time tissue samples from his surgery had been sequenced the patient had declined too far.

IBM has been outfitting Watson, its "cognitive computing" platform, to tackle multiple challenges in health care, including an effort to speed up drug discovery and several ways to help doctors with patient care. In this study, a collaboration with the New York Genome Center (NYGC), researchers employed a beta version of IBM Watson for Genomics.

IBM Watson's key feature is its natural-language-processing abilities. This means Watson for Genomics can go through the 23 million journal articles currently in the medical literature, government listings of clinical trials, and other existing data sources without requiring someone to reformat the information and make it digestible. Other Watson initiatives have also given the system access to patients' electronic health records, but those records weren't included in this study.

Laxmi Parida, who leads the Watson for Genomics science team, explains that most cancer patients don't have their entire genome (consisting of 3 billion units of DNA) scanned for mutations. Instead they typically do a "panel" test that looks only at a subset of genes that are known to play a role in cancer.

The new study, published in the journal *Neurology Genetics*, used the 76-year-old man's case to answer two questions. First, the researchers wanted to know if scanning a patient's whole genome, which is more expensive and time consuming than running a panel, provides information that is truly useful to doctors devising a treatment plan. "We were trying to answer the question, Is more really more?" says Parida.

The answer to that question was a resounding yes. Both the NYGC clinicians and Watson identified mutations in genes that weren't checked in the panel test but which nonetheless suggested potentially beneficial drugs and clinical trials.

Secondly, the researchers wanted to compare the genomic analysis performed by IBM Watson to one done by NYGC's team of medical experts, which included the treating oncologist, a neuro-oncologist, and bioinformaticians.

Both Watson and the expert team received the patient's genome information and identified genes that showed mutations, went through the medical literature to see if those mutations had figured in other cancer cases, looked for reports of successful treatment with drugs, and checked for clinical trials that the patient might be eligible for. It took the humans "160 person hours" to formulate recommendations, while Watson got there in 10 minutes.

However, while Watson's solution was first, it might not have been best. The NYGC clinicians identified mutations in two genes that, when considered together, led the doctors to recommend the patient be enrolled in a clinical trial that targeted both with a combinatorial drug therapy. If the patient had still been healthy enough, he would have been enrolled in this trial as his best chance of survival. But Watson didn't synthesize the information together this way, and therefore didn't recommend that clinical trial.

While it's tempting to view the study as a competition between human and artificial intelligence, Robert Darnell, director of the NYGC and a lead researcher on the study, says he doesn't see it that way. "NYGC provided clinical input from oncologists and biologists," he writes in an email. "Watson provided annotation that made the analysis faster. Given that each team addressed different issues, this comparison is apples to oranges."

IBM's Parida notes that the cost of sequencing an entire genome has plummeted in recent years, opening up the possibility that whole-genome sequencing will soon be a routine part of cancer care. If IBM Watson, or AI systems like it, are given swift access to this data, there's a chance they could provide treatment recommendations in time to save the lives of people like the brain-cancer patient in this study.

Darnell says he hopes IBM Watson will become a routine part of cancer care because the amount of data that clinicians are dealing with is already overwhelming. "In my view, having doctors cope with the avalanche of data that is here today, and will get bigger tomorrow, is not a viable option," he says. "Time is a key variable for patients, and machine learning and natural-language-processing tools offer the possibility of adding something qualitatively different than what is currently available."

This study was part of a collaboration between IBM and the NYGC announced in 2014, which set out to study the genomics of a few dozen brain-cancer patients. Darnell says the team is now working on a paper about the outcomes for 30 patients enrolled as part of that larger study.

It's worth noting that not everyone is sold on the value of IBM Watson for health care: A recent Wall Street analyst report declared that the Watson effort is unlikely to pay off for shareholders. Even though it called Watson "one of the more mature cognitive computing platforms available today," the report argued that Watson's potential customers will balk at the cost and complications of integrating the AI into their existing systems.

The report also called attention to a fiasco at the MD Anderson Cancer

Center in Texas, in which an IBM Watson product for oncology was shelved—after the hospital had spent US $62 million on it.

Autonomous Robot Surgeon Bests Humans in World First

By Eliza Strickland
Posted 4 May 2016 | 18:17 GMT

In a robotic surgery breakthrough, a bot stitched up a pig's small intestines using its own vision, tools, and intelligence to carry out the procedure. What's more, the Smart Tissue Autonomous Robot (STAR) did a better job on the operation than human surgeons who were given the same task.

STAR's inventors don't claim that robots can replace humans in the operating room anytime soon. Instead they see the accomplishment as a proof of concept—both for the specific technologies used and for the general concept of "supervised autonomy" in the OR.

Pediatric surgeon Peter Kim, one of the researchers, didn't sound threatened when he spoke to reporters in a press call yesterday. "Even though we surgeons take pride in our craft at doing procedures, to have a machine that works with us to improve outcomes and safety would be a tremendous benefit," he said.

For this study, published today in the journal *Science Translational Medicine*, researchers programmed their robot to carry out a procedure called intestinal anastomosis, in which a piece of intestine that's been cut through is stitched back together. It's like repairing a garden hose, said Ryan Decker, the senior engineer on the team, in that the sutures must be tight and regularly spaced to prevent leaks. STAR performed this task both on *ex vivo* tissue in the lab and on *in vivo* tissue in an anesthetized pig, and experienced human surgeons were given the same tasks. When the resulting sutures were compared, STAR's stitches were more consistent and more resistant to leaks.

The robot did have a little help. In about 40 percent of its trials, the researchers intervened to offer guidance of some sort—as in the GIF above, where a human hand is seen holding the thread. In the other 60 percent of trials, STAR did the job completely on its own.

The researchers don't think these assists invalidate their claim of autonomy; instead they see the setup as representative of shared control setups that would be appropriate for real ORs. Human surgeons could supervise procedures or even trade off tasks with the robot, letting the machine do more routine or tedious parts of an operation. “You can imagine that if something critical is happening, that would be a point where the surgeon is going to be closely monitoring the robot,” Decker said. “I’m sure they wouldn’t feel comfortable just letting it run and going to take a coffee break.”

Today, some surgical procedures already incorporate smart machines. Robots routinely carry out the crucial steps in some procedures including orthopedic knee replacements, Lasik eye surgery, and hair transplants. What these types of surgery have in common, though, is the fixed nature of their targets, as leg bones, eyes, and heads can be held in place during the procedure. Soft tissue surgeries are much messier and more difficult to automate, because all the slippery pink parts of the body shift around and are hard to track.

The current state-of-the-art robot for soft tissue surgery is the da Vinci system from Intuitive Surgical, but it’s not automated at all. The da Vinci is a teleoperated system, in which the surgeon sits at a console and manipulates controls in dexterous maneuvers that are mimicked by tiny tools inside the patient’s body.

STAR solved the soft tissue challenge by integrating a few different technologies. Its vision system relied on near-infrared fluorescent (NIRF) tags placed in the intestinal tissue; a specialized NIRF camera tracked those markers while a 3D camera recorded images of the entire surgical field. Combining all this data allowed STAR to keep its focus on its target. The robot made its own plan for the suturing job, and it adjusted that plan as tissues moved during the operation.

The researchers trained STAR only on how to perform this particular intestinal suturing procedure. “We programmed the best surgeon’s techniques, based on consensus and physics, into the machine,” Kim said.

An outside expert in the field of surgical robotics called this study a breakthrough, but also said its limitations show that autonomous robots "will not come to the OR soon." Blake Hannaford, a pioneer of autonomous surgical robotics at the University of Washington, noted that the NIRF tags that the robot relied on were placed by humans.

Hannaford also questioned the clinical significance of the task that STAR performed. "While in a technical sense, semi-autonomous suturing is a 'grand challenge' problem of surgical robotics, clinically much suturing and bowel anastomosis is done by staplers which can do the whole thing in seconds," he wrote in an email. "Clearly the task they chose does not justify the elaborate equipment they used."

The STAR team said this task was simply intended as proof that autonomous robots could meet the challenge of soft tissue surgery. While the robot may not be ready to take over the OR, Kim said he hopes his technology will be integrated into commercial devices in the next few years. If robotic systems are shown to improve safety and patient outcomes, he said, medicine may go the way of the auto industry.

"Now driverless cars are coming into our lives," Kim said. "It started with self-parking, then a technology that tells you not to go into the wrong lane. Soon you have a car that can drive by itself." Similarly, he said, surgical robots could start by giving human surgeons a helping hand. And maybe one day they'll take over.

Ophthalmologists vs. AI: It's a Tie

By Megan Scudellari
Posted 1 Feb 2017 | 15:00 GMT

Last week, we reported on an algorithm that recognizes skin cancer as well as the world's best dermatologists. That computer program was trained using 130,000 images from more than 2,000 diseases. It, like most artificial intelligence (AI) breakthroughs, relied on big data.

Now, a team in China has demonstrated that AI also has the potential to aid medical diagnoses in situations where there is limited high-quality data available. An AI program trained with just 410 images of congenital cataracts (a rare disease that causes irreversible blindness), plus 476 images of disease-free eyes performed as well as doctors in diagnosing the condition, recognizing its severity, and offering a treatment suggestion.

Inspired by a 2015 research paper from Google's DeepMind artificial-intelligence company, which described how a machine-learning algorithm beat professional human players at a series of arcade games based on minimal starting information, Haotian Lin, an ophthalmologist at Sun Yat-sen University in China, and colleagues had the idea of creating an AI agent to mine their clinical database on childhood cataracts.

Working with Xiyang Liu's team at Xidian University, they created CC-Cruiser, an AI program able to diagnose congenital cataracts, predict the severity of the disease, and suggest treatment decisions. The program was created using deep-learning algorithms trained with the aforementioned images of affected and control eyes from children.

Then, the researchers put CC-Cruiser to five tests. First, in a computer simulation, the AI program was able to distinguish patients and healthy individuals with 98.87 percent accuracy. It also achieved above 93 percent accuracy in estimating each of three indicators of the disease's severity: lens opacity areas, density, and location. The program also provided treatment suggestions with 97.56 percent accuracy.

Next, the team conducted a clinical trial using 57 images of children's eyes from three collaborating hospitals in China. None of the chosen hospitals specialize in diagnosing or treating this condition, because the team hopes the platform will eventually be most useful in helping hospitals like these, which lack on-site specialists. Again, CC-Cruiser performed well: 98.25 percent identification accuracy; over 92 percent on all three severity factors; and 92.86 percent accuracy in treatment suggestions.

In yet another test of the AI's capabilities, the program was presented with 53 low-quality cases mined from the Internet. It handled them with a high level of accuracy. But the researchers still weren't done. After that, the program successfully identified three needle-in-a-haystack cases, correctly pointing out three cataract cases in a data set with 300 normal cases.

Finally, in an effort to simulate real-world use, they pitted the program directly against individual ophthalmologists. CC-Cruiser and three ophthalmologists—an expert, competent, and a novice—went head-to-head diagnosing 50 clinical cases. The computer and doctors performed comparably.

The program did incorrectly flag a few cases in the hospital trial, and Lin hopes that a larger dataset could improve its performance. The team plans to build a collaborative cloud platform to do so, but Lin emphasize that the technology is "insufficient" to determine the best course of treatment with 100 percent accuracy. "Doctors should therefore make good use of the machine's suggestion to identify and prevent the potential misclassification and complement their own judgment," Lin told *IEEE Spectrum* in an email.

So, it's unlikely that CC-Cruiser will make ophthalmologists obsolete anytime soon. Especially because there is one key skill that it cannot do as well as doctors: "The human ability to communicate and interact affectively is indispensable for medical treatment," said Lin. "The simulation of human emotion is very challenging for machine[s]. The face-to-face interaction between doctors and patients will be one of the last bastions of human intelligence."

The team hopes that, with further clinical trials, doctors in non-specialized hospitals could use the program to identify the condition and send patients to specialized centers. Patients may eventually use it themselves and see out a specialist if there were a concern, added Lin. “The ultimate goal of artificial intelligence is to leverage it, in combination with human abilities, to make the world a better place.”

Sentient Technologies Manages Investments With Artificial Intelligence

IEEE Fellow helped develop software based on evolutionary computation to make smarter stock trades

By AMANDA DAVIS 10 October 2017

Stock traders are fallible. Sometimes they act based on gut instinct rather than sound logic. That's why Sentient Technologies (known as *Sentient.AI*) of San Francisco is working on Sentient Investment Management, an artificial-intelligence program that decides which stocks to buy and sell.

"Our AI system can be more consistent and reliable," says IEEE Fellow Risto Miikkulainen, the company's vice president of research. "It buys and sells stocks based on an entire history of data rather than, for example, on a single unreliable piece of information that a stockbroker might fixate on."

The company's researchers have been developing the program for the past 10 years to manage its investments. It relies on evolutionary computation, in which computers are fed massive amounts of data and, after hundreds of thousands of trials at making decisions, the machine evolves to make better ones. With Sentient's program, that data includes a history of changes in stock prices and past returns on its investments. The program, which learns to make better trades as it goes, now makes all the trading decisions for Sentient's investments.

YEARS OF TRAINING

Evolutionary computation is a form of machine learning, Miikkulainen notes. With standard machine learning, computers are fed a large data set containing hypothetical situations and told what the right answers should be. "Evolutionary computation is more like learning to play a video game," he says. "We didn't know what the right actions were at the beginning. Instead we learn by exploring and trying out different things."

Evolution uncovers what Miikkulainen calls *candidates.* Different candidates emerge as the system is tested, and the candidates learn by trial and error. If they make bad trading decisions—making small gains or, worse, losing money—they're discontinued. But if they make smart, profitable choices, they live to make another simulated trade.

At its peak, Miikkulainen says, Sentient's system was testing 40 trillion candidates per year. Each one is tested hundreds of thousands of times, and if it passes them—essentially, if it makes mostly intelligent, profitable decisions—then it is deployed to make actual trades.

"These candidates wake up each morning and buy and sell stocks, just like day traders, only there's no human oversight," he says. "They are completely autonomous."

HUMANS STILL NEEDED

AI lacks at least one important quality: common sense. "Automated stock traders are basically idiot savants," Miikkulainen says, meaning that they can do only the thing they've been trained to do. Sometimes they don't recognize that something has gone wrong.

During a major event that could negatively affect stock prices, there always will be a "panic button" that a trader can press to shut down the automatic system, Miikkulainen says. That might include the outbreak of a war or the unexpected results of an election. Because such events were not part of the AI program's training, it could not factor them into its trading strategies.

Sentient Investment Management is getting smarter by the day, Miikkulainen says. For now, the product is not available to customers; it is still being tested internally. It could eventually be applicable elsewhere, though. "After proving the technology, we'll expand to other markets," Miikkulainen says.

No matter what, minders with expertise in engineering and finance will still be required. "For the foreseeable future, these systems will need to have a human companion, providing the common sense," he says. "AI is simply an advanced decision-making tool."

References

*Judith H. Dobrzynski (2017) ***Artificial Intelligence as a Bridge for Art and Reality***, The New York Times, October 25

Cade Metz (2017) ***How A.I. Is Creating Building Blocks to Reshape Music and Art***, The New York Times, August 14

Gideon Lewis-Kraus (2016) ***The Great A.I. Awakening: How Google used artificial intelligence to transform Google Translate, one of its most popular services – and how machine learning is poised to reinvent computing itself, The New York Times Magazine, December 14

Luke Dormehl (2017) ***Thinking Machines: The Quest for Artificial Intelligence – and Where It's Taking Us***, TarcherPerigge

Richard Yonek (2017) ***Heart of the Machine: Our Future in a World of Artificial Emotional Intelligence***, Arcade Publishing,

****N. Bostrom (2014) ***Superintelligence, Paths, Dangers, Strategies***, Oxford University Press

Martin Service (2015) ***Rise of the Robots: Technology and the threat of a jobless future***, Basic Books, A member of the Perseus Group, New York

Yuval Noah Harari (2015) ***Homo Deus: A Brief History of Tomorrow, Signal, McClelland & Stewart

William F. Brown (2016) ***Perspectives: The Evolution of the Cosmos, Life, Humans, Culture and Religion and a Look into the Future***, Friesen Press, Pages 299-302

Luke Dormehl (2017) ***Thinking Machines: The quest for artificial intelligence and where it's taking us next***, A TarcherPerigee Book

Max Tegmark (2017) ***Life 3.0: Being Human in the Age of Artificial Intelligence***, Alfred K. Knopf

Sheelah Kolhatkar (2017) ***Dark Factory: The robotics revolution is changing what machines can do. Where do humans fit in?*** The New Yorker, October 23, Pages 70-81

Journals

(Ziad Obermeyer and Thomas H. Lee (2017) ***Lost in Thought – The Limits of the Human Mind and the Future of Medicine, N. Engl. J. Med. October 4

Nichola Jones (2017) ***Machine learning tapped to improve climate forecasts***, Nature, 548, 24 August, Pages 379-380

Eliza Strickland (2014) ***IBM Watson on the Genetics of Brain Cancer**,** IEEE Spectrum, 19 March

Eliza Strickland (2017) ***IBM Watson Makes a Treatment Plan for Brain-Cancer Patient in 10 minutes; Doctors Take 160 hours, IEEE Spectrum, 11 August

Sarah Lyall (2017) ***Enjoy Those U.S. Open Highlights. A Computer Picked Them for You***. The New York Times, September 5

IEEE Spectrum (2017) Special Issue, ***Can We Copy the Brain? Intensive efforts to understand and re-create human cognition will spark a new age of maching intelligence and transform the way we work, learn and play, June

Fred Rothganger, ***The Dawn of the Real Thinking Machine***, Pages 23-45
Karlheinz Meier, ***the Brain as Computer***, Pages 29-33

Jeff Hawkins, ***What intelligent Machines Need to Learn From the Neocortex***, Pages 35-39, 70-
Eliza Strickland, ***From Animal Intelligence to Artificial Intelligence***, Pages 41-45
Jennifer Hasler, ***A Road Map for the Artificial Brain***, Pages 47-50

Jean Kumagai, ***Navigate Like a Rat***, Pages 59-63
Christof Koch and Giulio Tononi, ***Can we Quantify Machine Conscious-ness?***, Pages 65-71

Nature Insight: Quantum Software (2017) 549, 14 September

Earle T. Campbell, Barbara M. Terhal & Christophe Vuillot, ***Roads toward fault-tolerant universal quantum computation***, Pages 172-179
Frederic T. Chong, Diana Franklin & Margaret Martonosi, ***Programming languages and compiler design for realistic quantum hardware***, Pages 180-187
Daniel Bernstein & Tanja Lange, ***Post quantum cryptography***, Pages 188-193
Jacob Blamonte et al. ***Quantum machine learning***, Pages 195-202
Aram W Harrow & Ashley Montanaro, ***Quantum computational supremacy***, Pages 2013-209

(2017) ***The AI composer: Q&A with Luc Steels*, Nature 549, 14 September, Page 157

Will Zeng (2017) ***First quantum computers need smart software***, Nature, 549, 14 September, Pages 149-150

Dina Fine Maron (2017) ***The Computer Will See You Now***, Scientific American, August

Elizabeth Gieney (2016) ***Google masters Go: Deep-learning software excels at complex ancient board game, Nature, 329, 28 January, Pages 445-446

Omid Veisen and Robert Langer (2015) ***A smart insulin patch***, Nature, 524, 6 August, Pages 39-40

David Pogue (2017)) ***Computer-Aided Creativity: How much should an artist reveal about letting technology make some choices?*** Scientific American, July, Page 28

Andrea Esteva et al (2017) ***Dermatologist – Level classification of skin cancer with deep neural networks Nature 542, 2 February, Pages 115-118

Siddhartha Mukherjee (2017) ***The Algorithm Will See You Now: When It comes to diagnosis, will A.I. replace the M.D.? The New Yorker April 3, Pages 46-53

D. Silver et al. (2016) ***Mastering the game of Go with deep neural networks and tree search Nature 529, 29 January, Pages 484-489

Gary Marcus (2017) ***Artificial Intelligence Is Stuck. Here's How To Move It Forward***, The New York Times, July 29

Elizabeth Kolbert (2016) ***Rage Against the Machine: Will robots take your job? The New Yorker, December 19 &26

Davide Castelvecchi (2015) ***Artificial intelligence called in to tackle LHC data deluge***, Nature, 528, 3 December, Pages 18-19

Mnih V et al. (2015) ***Human level control through deep reinforcement learning***, Nature, 518, 26 February, Pages 529-533

Podolny S. (2015) ***If an algorithm wrote this, How would you even know***?, The New York Times, 8 March

Miller C.C. (2015) ***Can algorithms hire better than a human?* The New York Times, 28 June, Page SR4

Bohannon J. (2014) ***Meet your new co-worker*, Science, 346, 10 October, Pages 180-181

Service R.F. (2014) ***Minds of their own***, Science, 346, 10 October, Pages 182-183

Normile D. (2014) ***In our own image***. Science, 346, 10 October, Pages 188-189

Tim Appenzeller (2017) ***The Scientist's Apprentice***, Science, 357, July 7, Pages 16-17

John Bohamnon (2017) ***The Cyberscientist: Artificial intelligence isn't just a tool. In some labs it conceives and carries out experiments – and then interprets the results, Science, 357, July 7, Pages 18-21

Paul Voosen (2017) ***The AI detectives: As neural networks push into science, researchers probe back, Science, 357, July 7, Pages 22-23

Matthew Hutson (2017) ***How algorithms can analyze the mood of masses, Science,*** 357, July 7, Page 23
Elizabeth Pennisi (2017) ***Combing the genome for the roots of autism***, 357, July 7, Pages 25-27
Joshua Sokoi (2017) ***Machines that make sense of the sky***, Science, 357, July 7, Page 26

Robert F. Service (2017) ***Neural networks learn the art of chemical synthesis***, Science, 357, July 7, Page 27
Jennifer Sills (2017) ***Artificial intelligence in research: A survey of scientists***, Science, 357, July 7, Pages 28-29

Eric Horvitz (2017) ***AI, people and society***, Science, 357, 7 July, Page 7

Siddhartha Mukherjee (2017) ***The algorithm will see you now***, The New Yorker, April 3

Gerome Groopman (2017) ***What's The Trouble? How doctors think***, The New Yorker, January 29 issue

Ryan Cross (2017) ***Artificial Intelligence helps scientists map behavior in the fruit fly brain***, Posted in: Brain & Behavior Plants & Animals

Jelena Stajie, Richard Stone, Gilbert Chin and Brad Whible (2015) ***Rise of the Machines,*** Science 349, 17 July 27, 2017

John Bohannon, ***The synthetic therapist: Some people prefer to bare their souls to computers rather than to fellow humans***, Pages 250-251

John Bohannon, ***Fears of an AI pioneer: Stuart Russell argues that AI is as dangerous as nuclear weapons***, Page 252

Eric Horvitz and Deirdre Mulligan, ***Data, privacy, and the greater good***, Pages 253-254

M. I. Jordan and T. M. Mitchell, ***Machine learning: Trends, perspectives, and prospects***, Pages 255-260

Julia Hirschberg and Christopher D. Manning, ***Advances in natural language processing***, Pages 261-2266

David C. Parkes and Michael P. Weilman, ***Economic reasoning and artificial intelligence***, Pages 267-272

Samuel Gershmann, Eric J. Horvitz and Joshua H. Tenenbaum, ***Computational rationality: A converging paradigm for intelligence in brains, minds, and machines***, Pages 273-278

Judith H. Dobrzynski (2016) ***Artificial Intelligence as a Bridge for Art and Reality***, The New York Times, October 25

Cade Metz (2017) ***How A.I. Is Creating Building Blocks to Reshape Music and Art***, The New York Times, August 14

Electronic version of Science

Science News Staff (2017) ***AI is changing how we do science. Get a glimpse***, July 5

Adrian Cho - ***AI's early proving ground: the hunt for new particles***

How algorithms can analyze the mood of the masses

Matthew Hutson – ***A new breed of scientist, with brains of silicon***

Elizabeth Pennisi – ***Combing the genome for the roots of autism***

Joshua Sokol – ***Machines that make sense of the sky***

Robert Service - ***Neural networks learn the art of chemical synthesis***