Are we ready to be diagnosed by AI doctors?

The WAEH and Artificial Intelligence - Possibilities and Pitfalls

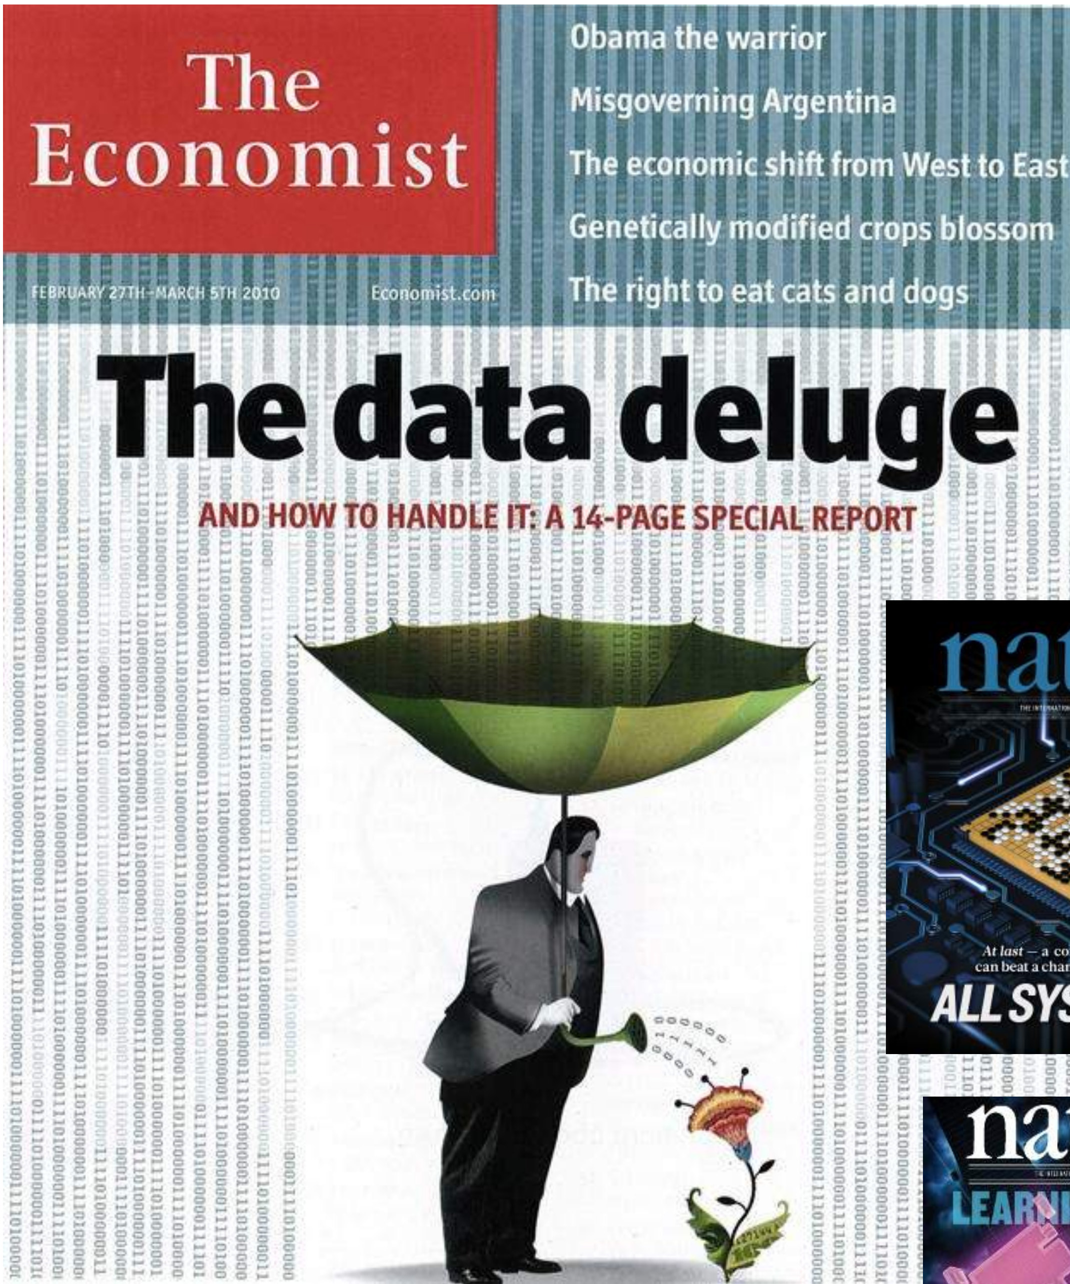

Are we ready to be diagnosed by AI doctors?

The WAEH and Artificial Intelligence - Possibilities and Pitfalls

INTRODUCTION

Within the next decade artificial intelligence (AI) and 'big data' will change the face of medicine. Doctors and patients will be confronted with fundamental changes in the fields of diagnostics and in the organization of medical care. A recent survey by HIMSS (a global organization dedicated to the improvement of health through information technology) showed that half of the world's hospitals expect to introduce AI in the next five years. Another survey, conducted under 11.000 patients worldwide by the accountancy firm PriceWaterhouseCoopers, shows that the majority of patients won't mind that doctors will be supported or replaced by computers or robots. Is this going to be true or do we need to research it ourselves a little bit more in detail? Are patients ready to be diagnosed by AI-doctors (and having their organs 3D printed, or having nanometer-sized robots roaming around in their body to monitor their health)?

source: Gigaom

At this moment the greatest challenge healthcare is facing is the aging of the population, and the surge in age related chronic diseases. Consequently, medical care is gradually shifting from 'repairing' intrinsically healthy patients to monitoring health and supporting and improving quality of life. In recent years, and the coming decades, care workers will give more attention to what patients are still able to do instead of what they can't, and to organizing care outside the hospital: in the community, by members of the family and at home.

Technology will play an important part in this process. The introduction of AI will fundamentally change the way diseases are diagnosed and treated. AI will make diagnostics faster, cheaper and more reliable than present procedures. AI computers, fed with massive amounts of information ('big data'), should be able to equate human experience and expertise within a few years. Now, diagnosis still is one of the four main functions of a hospital (next to advice, treatment and aftercare). But chances are that, thanks to AI, a large part of diagnostics will be automated and that private clinics will take over a significant part of this activity.

On computers, Artificial Intelligence and Deep Learning

On computers, Artificial Intelligence and Deep Learning
To create a computer with artificial intelligence (AI), programmers use a procedure called Deep Learning. They build a neural network, a multilayered grid of connections and nodes, loosely based on the structure of the human brain. These nodes are made in such a way that they can be programmed to relay, alter or stop a signal that passes through. This grid is then fed with data on the input side. From the start this input leads to a certain signal on the other 'output' side. In first instance, the output is totally meaningless but by presenting the network with masses of input data, and at the same time 'telling' the grid at the outside what the output should be, the networks internal settings are changed and the grid 'learns' to give the correct output to new input. When this procedure is repeated hundreds or thousands of times, using a great variety of input data (combined at the output side with the desired output), the neural network can become an 'artificial brain' that operates faster and more accurate than its flesh-and-blood counterpart.

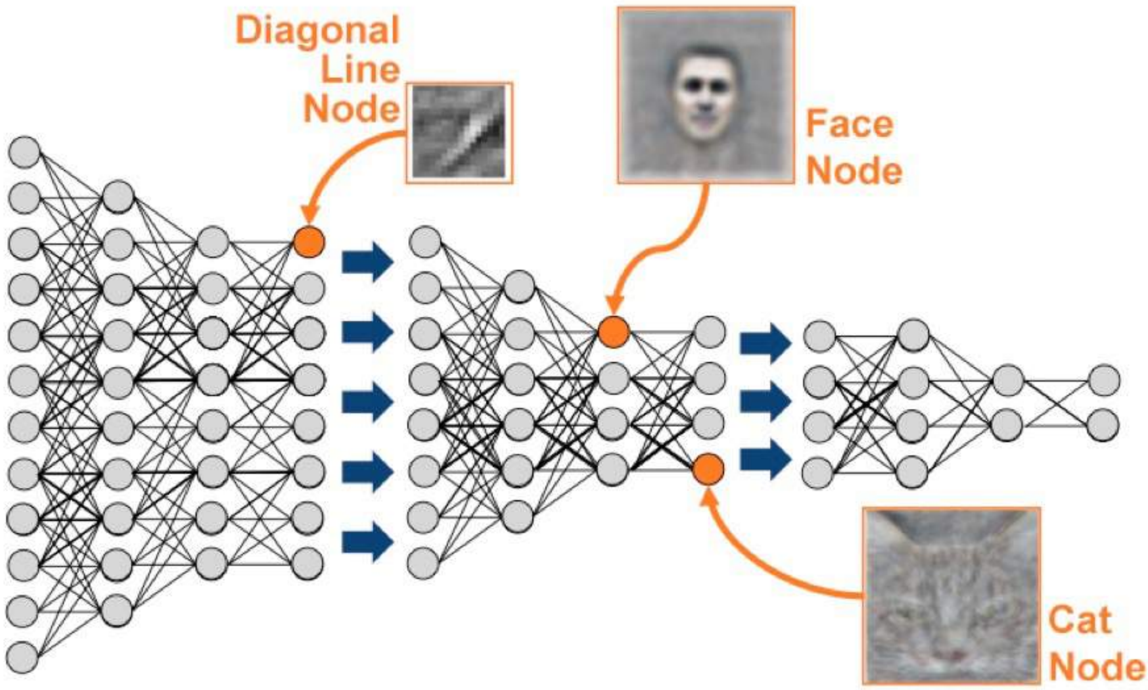

Source: Moorfields Eye Hospital, courtesy of Pearse Kean

WAEH
WORLD
ASSOCIATION
OF EYE HOSPITALS

AI and big data – the Watson Health Cloud

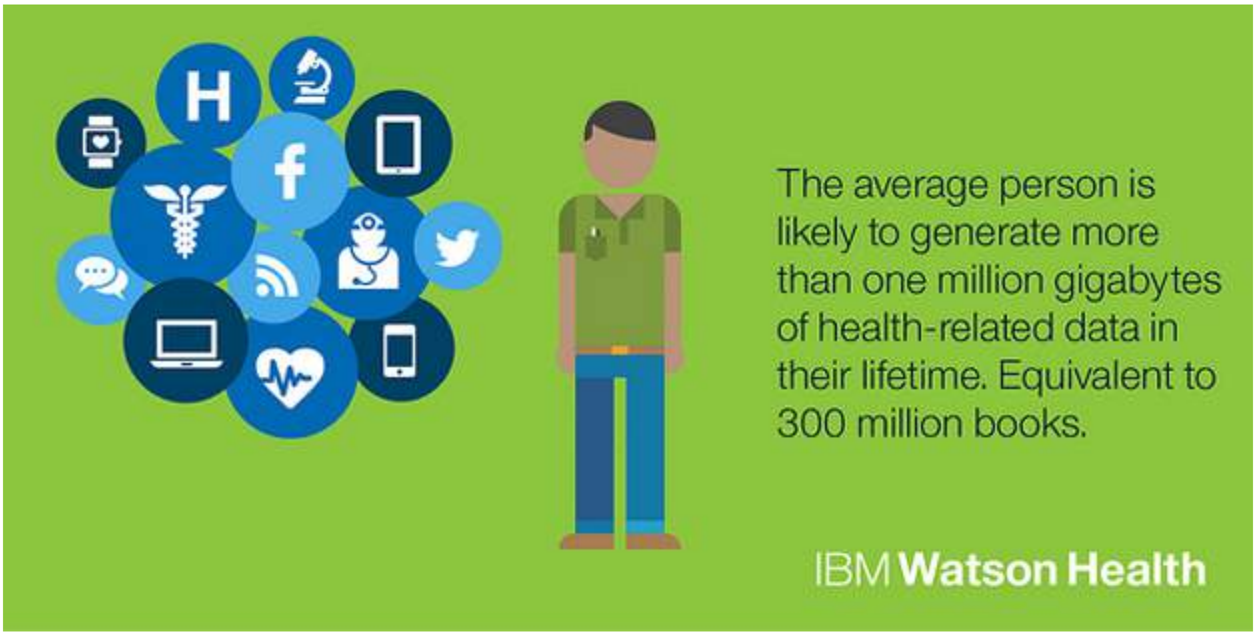

source: IBM Watson

A striking example of combining AI and big data is IBM's Watson Health Cloud. IBM's supercomputer Watson is probably most famous for beating the two successful contestants in the American general knowledge quiz *Jeopardy!* in 2011. But in the next decades Watson will probably have a significant impact on the fields of medical research and medical care. Watson, a computer that can speak like a human, can absorb and analyze an amount of data equivalent to seventy million pages per second. A phenomenal speed – but in fact not enough to keep up with the present growth of medical data. All this is stored in the Watson Health Cloud, which eventually will turn every Watson supercomputer into a useful, and in the long run a superior assistant for the medical professional.

Watson has already shown its value in several medical fields, such as oncology, where it is used for planning treatment for individual patients, based on personal medical records and genomic information, combined with the latest scientific insights. IBM has teamed up with the American firm Quest Diagnostics on tumor genome sequencing for identifying mutations that may have therapeutic significance.

In another project, in collaboration with the American Diabetes Association, IBM launched an app that delivers real time information to diabetics, giving them the possibility to directly see the influence of behavior and food intake on glucose levels. These personal data are collected in the Watson Health Cloud, and they could lead to important new insights in diabetes.

AI having a large impact on blind people's lives

Another example of how Artificial Intelligence can be used, is the OrCam MyEYe: a mini camera, clicked onto people's own glasses, and connected to a small computer that can be placed in a pocket or hung on a belt. If the user points to the text on a menu or for example a product in the supermarket the camera reads the text aloud. OrCam can also learn to recognise faces and to tell the user who is in front of them. It can do this thanks to highly advanced smart software that makes use of artificial intelligence.

source: OrCam

"This is a ground breaking new technology which will ensure that many blind and visually impaired people all over the world, will be able to independently read and live again. They will be able to travel independently by public transport, do their shopping and read a menu in the restaurant. In the world of medical devices for blind and visually impaired people there is currently no other device that allows people to regain their independence like this," says ophthalmologist Dr Joosse from the Netherlands.

CASE STUDY – WAEH
Applying big data – Moorfields Eye Hospital

Turning a mas of raw data into a neat and useful stack of 'big data' can be a problem, as Pearse Keane from Moorfields Eye Hospital, one of our founding members, experienced. Last year, Moorfields announced it would cooperate with DeepMind (part of Google's Alphabet Group, specializing in the application of Artificial Intelligence) on the development of an algorithm for the interpretation of OCT scans. This algorithm would be able to interpret scans and diagnose a range of retinal diseases. Keane soon discovered that to get this cooperation started, he first had to take a deep look at the mass of scans Moorfields has gathered over the years.

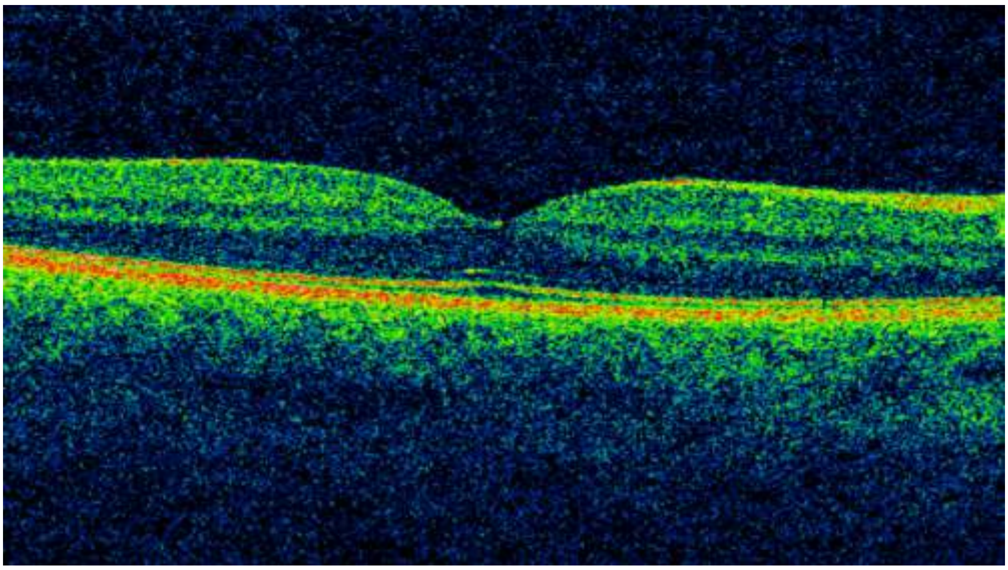

Optical Coherence Tomography - source: Moorfields Eye Hospital

Applying Artificial Intelligence to OCT scans
Applying Artificial Intelligence to medical datasets is currently a hot research topic. Two years ago Nature dedicated its cover illustration to research in which Artificial Intelligence was used to diagnose skin cancer. The algorithm was at least as good as a team of the best specialists and it's clear that dermatology will be confronted with big challenges. For Keane, it was the sheer volume of OCT scans piling up in the archives at Moorfields that gave him the idea of applying AI to these scans. Easier said than done. Keane: 'In July 2016 we formally announced the collaboration between Moorfields and DeepMind. We got much publicity, but at the same time I felt a little bit awkward, because we hadn't done anything yet at that moment.' Twelve months before, he had contacted Mustafa Suleyman from DeepMind, telling him that at Moorfields they treat thousands of patients, and possessed thousands of OCT scans. Keane: 'And the distance between our two organizations is only two Tube Stations. We must cooperate, I said. To my great excitement he replied straight away.'

But that turned out to be the beginning of an arduous process. Keane: 'In the last twelve months me and a big team of colleagues have been working almost every day to get to the point we could announce the cooperation publicly. It was all a question of dataset quantity and quality. When Mustafa and I first met, I told him: We do loads of OCT scans! Three thousand a week! He asked: How many do you have in total? I couldn't tell. And what percentage of these scans have labels on them? And those with labels, what percentage is mild, moderate or severe? I couldn't answer his questions! And then there were problems like: what file format are they in? And lastly, and most importantly, what are the ethical and governance issues of sharing a large amount of NHS data with a commercial partner? That aspect also took a lot of time to figure out.'

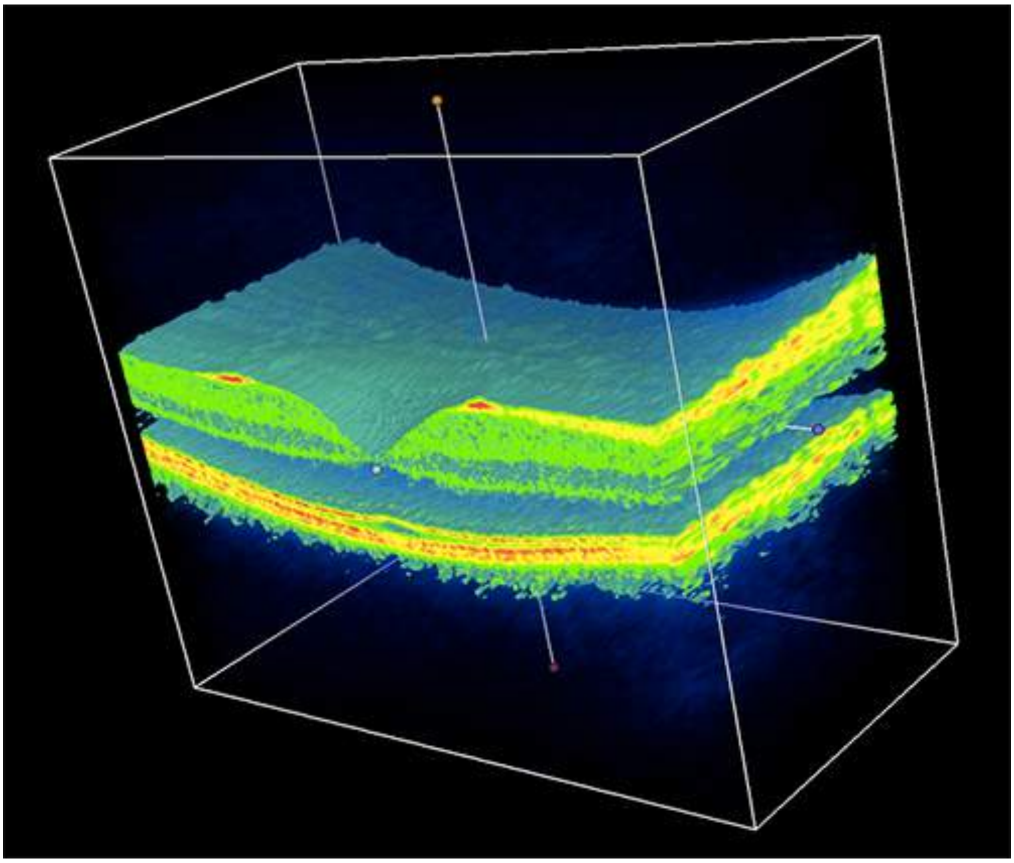

source: Moorfields Eye Hospital, OCT scan, with courtesy of Pearse Kean

Extracting new clinical and scientific insights from big data

Having good, clean datasets turned out to be the bottleneck on applying AI. But when these problems were dealt with, the project was running smoothly. Keane went over to DeepMind on a weekly basis, 'and I could see rapid advances in what they were achieving.' His goal, as he puts it, is 'to bridge the AI chasm.' In the tech world it is often quite easy to do a 'proof of concept' using small, smooth data sets, or a toy problem. The real problem is to transform the small-scale approach into something that can be deployed in a real world setting and in actual clinical practice. Keane: 'What we hope to do is making an algorithm that can diagnose any retinal disease that a retinal specialist can diagnose based on an OCT scan.' In the longer term he hopes to extract new clinical and scientific insights from this data. Algorithms could improve the modelling now based on conventional statistics, or could look at longitudinal changes much more accurately than conventional techniques.

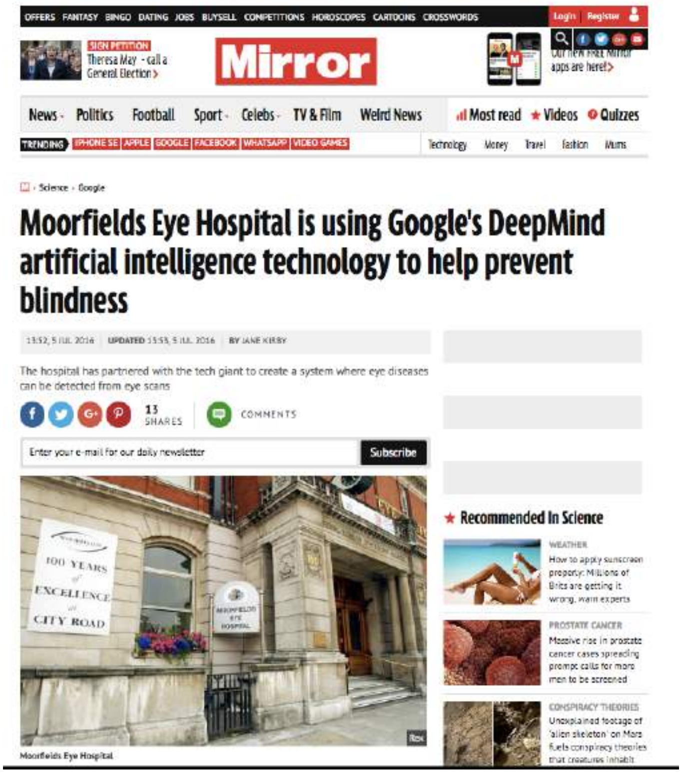
OFFERS FANTASY BINGO DATING JOBS BUYSELL COMPETITIONS HOROSCOPES CARTOONS CROSSWORDS Login | Register

SIGN PETITION Theresa May - call a General Election >

Mirror

apps are here!>

News Politics Football Sport Celebs TV & Film Weird News Most read Videos Quizzes

TRENDING IPHONE SE APPLE GOOGLE FACEBOOK WHATSAPP VIDEO GAMES Technology Money Travel Fashion Mums

Science › Google

Moorfields Eye Hospital is using Google's DeepMind artificial intelligence technology to help prevent blindness

13:52, 5 JUL 2016 | UPDATED 15:55, 5 JUL 2016 | BY JANE KIRBY

The hospital has partnered with the tech giant to create a system where eye diseases can be detected from eye scans

13 SHARES COMMENTS

Enter your e-mail for our daily newsletter Subscribe

Moorfields Eye Hospital

Recommended In Science

WEATHER
How to apply sunscreen properly: Millions of Brits are getting it wrong, warn experts

PROSTATE CANCER
Massive rise in prostate cancer cases spreading prompt calls for more men to be screened

CONSPIRACY THEORIES
Unexplained footage of 'alien skeleton' on Mars fuels conspiracy theories that creatures inhabit

Keane thinks it is likely that the next generation of algorithms will use ever more diverse data. 'And don't forget, in the coming years the amount of clinical data about each patient will grow exponentially. We will get ten different types of images and scans, data from genetics, boleomics, et cetera.' Eventually, every medical specialism that deals with high dimensional data can benefit from applying deep learning. Still, there must be someone who can come up with useful advice for the patient. Keane: 'I strongly DON'T believe that the role of the clinician will be removed. We'll have to stay ahead of developments and update our skills but there are certain aspects that will always require human skills. It must be a human who tells a patient what the diagnosis is. You need a human who has the clinical wisdom and nuanced expertise to give good advice for different kinds of patients. We must make sure we have lots of communication skills. In our research we must be imaginative and creative, and in the clinic we must be experts on teamwork. '

Dr. Pearse Kean: “I believe: 1) we're about to reach a tipping point where artificial intelligence will transform healthcare, and 2) ophthalmology is likely to be medical specialty which leads the way!”

CASE STUDY – Aravind Eye Hospital and Google

In June this year Aravind Eye Hospital – associate member of the WAEH- and Google revealed they had been quietly working together for several years on the development of AI for automated diabetic retinopathy screening. Aravind is the biggest eye hospital in India, with headquarters in Madurai (Tamil Nadu) and several branches throughout the subcontinent. The task ahead is enormous: experts estimate that 70 million Indians are at risk of developing this form of retinopathy. The system, the two partners developed is going through its validation phase at the moment. Chief medical officer Dr. R. Kim has high hopes: 'I have a feeling it could turn out to be a good diagnostic tool soon. Depends on how the sensitivity increases — it could be much more sensitive than human graders.'

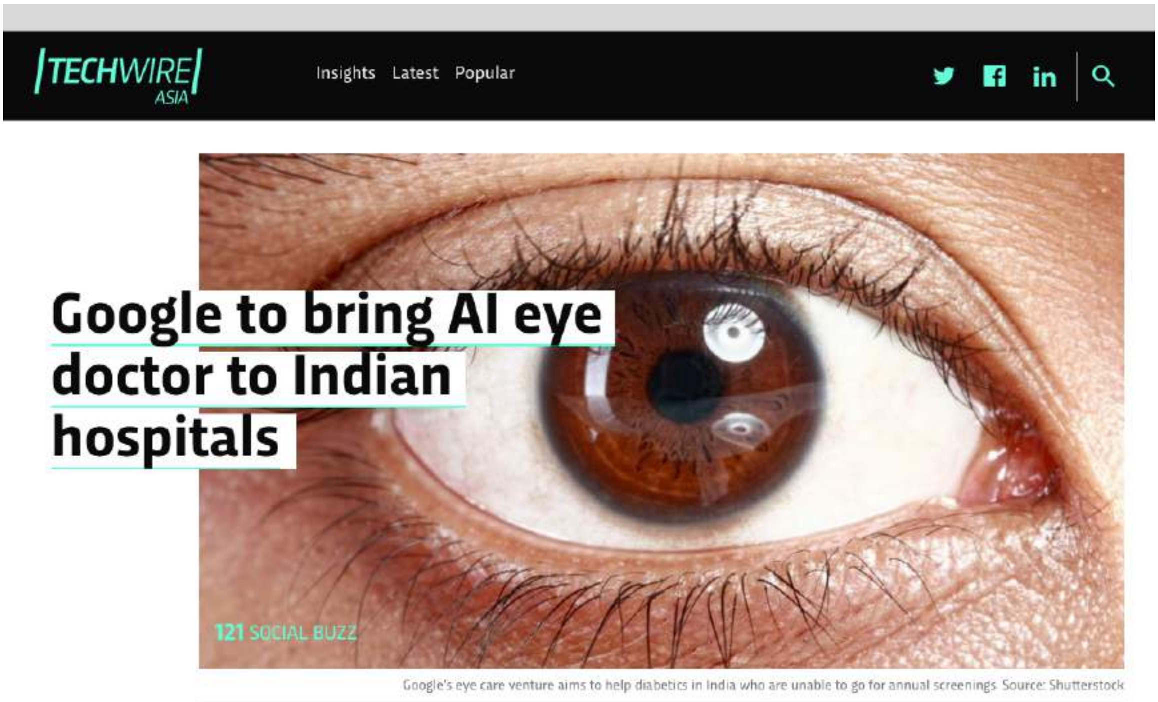

Executive director Mr. R.D. THULASIRAJ expects that with the advent of AI and big data analytics Aravind should be able to move into the "predictive domain": "Till now analytics helped us analyze the past and gave us some new insights and avenues to improve. Though very useful, it still is on the past data. We are now into real-time analytics and alerts that help us give the right prescriptions or manage the processes better as they are happening. We now have the capability to transition into predictive analytics. When we are examining a patient, can we know with adequate certainty that this patient will rapidly develop Diabetic Retinopathy, or that the myopia will progress fast or that the patient is unlikely to comply with the prescribed treatment. Such focal knowledge ahead of time will go a long way in better caring for the patient. This is one of the areas that we would like to pursue.

Dr. R. KIM, Chief Medical Officer and Chief Retina Services, Aravind Eye Hospital, Madurai: "AI in the immediate time will address the difficulty in getting trained resources for grading the retinal images. AI will now make it possible to screen a large number of the diabetic population for diabetic retinopathy, in a much more uniform manner and also much faster. This is just a beginning and all of a sudden it opens a new avenue to screen many other blinding conditions such as glaucoma and AMD."

CASE STUDY - Applying big data – The Rotterdam Eye Hospital

The Rotterdam Eye Hospital – founding member of the WAEH - has a different, ever-growing pool of "big data" consisting of thousands of fundus photos. Koen Vermeer, director of The Rotterdam Ophthalmic Institute (the research institute of the Rotterdam Eye Hospital), worked with PhD student Kedir Adal, ophthalmologist Jose Martinez Ciriano and the Quantitative Imaging group at Delft University of Technology to develop an algorithm which uses these photos not only for automated screening for diabetic retinopathy but also for the detection of diabetic retinopathy progression.

Fundus photos of diabetic patients' retinas are currently taken regularly in Rotterdam, once a year or every two years. The photos show tiny hemorrhages or the development of new veins, which are often very weak and prone to breaking. At present, the photos are still examined by a trained observer. Vermeer and his colleagues want to merge the four photos taken during retinal screening to produce a single superior image. Vermeer: 'The problem is that the four photos are often of poor quality. For example, the quality varies within an image, some parts are sharp but others not, and each photo shows only part of the retina, so you need to merge them in order to obtain a single image. This is further complicated by the fact that the eye is spherical while the images are flat.

So you can't simply merge the photos. The algorithm we are developing needs to correct the photos for such distortions and align them. And that has to be done with very high accuracy due to the small size of some of the lesions.'

The technique works so well that the composite image shows much more than was observed previously. So the question is can the people interpreting the photos "by eye" really see more on the improved images? Vermeer: 'To investigate this, we are working with Moorfields. This institute has a centre for interpreting such photos for international clinical studies. They can have several people look at a photo and give a judgement.'

The Moorfields experts received both the ordinary color photos and the normalized images, and it was clear that they saw more on the latter than on the former. But are the additional things they see clinically relevant lesions, or noise amplified by the procedure. To find out, the team in Rotterdam is currently analyzing the data. Vermeer: 'We go back to the original photos to see if, on closer examination, we can find any evidence in them of damage to the blood vessels.' The team needs to establish whether what the Moorfields experts see on the composite images really exists (i.e. is visible on the original photos).

Only then can they take the next step towards designing an algorithm that makes optimum use of fundus photos.

Lessons learned?

- Obtaining large data sets from current clinical systems is relatively straightforward. Getting a clean data set with annotations (related to the patient, the eye, the disease, the equipment, etc), still requires a lot of effort.
- Current deep-learning approaches work great for very large data sets. In less frequently occurring diseases, such large data sets are not available and new methods need to be developed for automated screening and diagnosis.
- Parts of a newly developed system may be beneficial for other uses as well. In our case, the procedure to normalize fundus photos, which was developed for registration, seems to be advantageous for human observers and also improves the learning rate of our deep-learning system.
- Showing that you can outperform the current gold standard is still a fundamentally difficult challenge because you cannot prove that your algorithm does better – only that it has different outcomes.

CASE STUDY – Deep Learning projects – Singapore National Eye Centre

Several AI projects are currently explored by our full member the Singapore National Eye Centre, both on an operational level as also on a clinical level.

Current topics that are being explored - some of these in close cooperation with IBM Watson: how to improve the patient experience using deep learning, how to make the individualized patient time more efficient, impact of using deep learning in early detection of high risk glaucoma patients, patients with diabetic rethinopathy and patients with macula degeneration.

At least 5 ophthalmologists are currently involved in deep learning projects. All these projects are however still in the Proof of Concept phase and not rolled out yet in the clinic. Dr. Desmond Quek, Director of Medical Imaging at the SNEC: “Deep learning projects will improve the efficiency and accuracy of the doctors, but we don’t have to be afraid that we will be replaced by deep learning. It will make a lot of our decision however less critical and in an earlier stage.” Quek expects that the projects will be rolled out in the clinic in the upcoming 1,5 years.

CASE STUDY – Using AI to make automated AMD analysis possible @Wilmer Eye Institute Johns Hopkins

Age-related macular degeneration (AMD) is the leading cause of vision loss in people over fifty in the US when left untreated and affects millions of people throughout the world. What if you could detect this eye disease in an early stage and start treating it?

"About eight million US individuals have the intermediate stage of AMD that is often asymptomatic with regard to visual deficit. These individuals are at high risk for progressing to the advanced stage where the often treatable choroidal neovascular form of AMD can occur. Careful monitoring to detect the onset and prompt treatment of the neovascular form as well as dietary supplementation can reduce the risk of vision loss from AMD, therefore, preferred practice patterns recommend identifying individuals with the intermediate stage in a timely manner."

Ophthalmologists at the Wilmer Eye Institute have researched the novel application of a machine learning approach using deep learning on AMD analysis and demonstrated in a study the efficacy of machine grading based on deep universal features/transfer learning when applied to automated retinal image analysis. Deep learning is a promising step in providing a pre-screener to identify individuals with intermediate AMD.

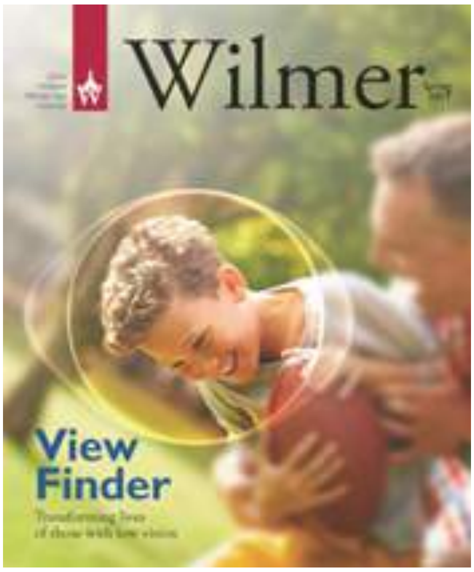

CONCLUSION AND RECOMMENDATIONS

Recent developments in the field of AI and big data are of enormous importance for ophthalmic hospitals. AI-based diagnostic systems will soon be able to work faster and even more accurate than real doctors. As the case studies above show, the results promise to be quite impressive. Still, hospitals who want to venture into this field must be on their guard for several pitfalls, especially when it comes to creating a reliable data stream suitable for this purpose. And as the Rotterdam Case Study shows, AI is not the Holy Grail, other digital techniques still can be very promising.

The use of AI is not only inevitable but also desirable because the number of patients and demand for accurate diagnoses will increase sharply in the coming decades, due to the ageing of the world population. Still, even if AI is able to take over the complete, or a major part of the diagnostic process, most experts think we will still need doctors for a second opinion and for administering medical advice – though there also are experts who dispute this and point out that in the future nurse-practitioners can take care of these tasks.

Automating diagnostics with de aid of AI will likely mean that diagnostics will move outside the hospital. In the not too distant future, customers may be able to have an OCT scan (or similar scans) carried out by their optician, who sends the scan to a company that makes the diagnosis and then forwards the diagnosis to the general practitioner. Treatment could take place at an eye hospital or a private clinic specialized in standard surgical procedures. Eye hospitals are already experiencing fierce competition in this area and this problem will only grow. It demands a repositioning in terms of organization and the services they want to provide. If they are to respond timely and appropriately to these developments, **they need to invest now in data management and relevant partnerships.**

------November 2017----

With great thanks to all the people who were willing to share information ans insights with us.

Marcel Hulspas
Maaike van Zuilen – maaike.vanzuilen@waeh.org

Sources

HIMMS patient survey: http://www.healthcareitnews.com/news/half-hospitals-adopt-artificial-intelligence-within-5-years

PWC survey: https://www.pwc.com/gx/en/industries/healthcare/publications/ai-robotics-new-health/survey-results.html

Moorfields: http://www.moorfields.nhs.uk/news/moorfields-announces-research-partnership

Aravind:
https://factordaily.com/news/google-india-aravind-eye-hospital-ai-diabetic-retinopathy/
https://www.wired.com/2017/06/googles-ai-eye-doctor-gets-ready-go-work-india/

OrCam – AI: http://money.cnn.com/video/technology/2015/07/17/orcam-smart-glasses-for-blind.cnnmoney/

Comparing humans and deep learning performance for grading AMD: A study in using universal deep features and transfer learning for automated AMD analysis - Philippe Burlina[a,b,c], Katia D. Pacheco[d], Neil Joshi[a], David E. Freund[a,□], Neil M. Bressler[b]