Ulster Medical Society

14 November 2019

Joint Meeting With
The Belfast City Hospital Medical Staff

Artificial Intelligence Will Make Doctors Obsolete

Professor Jörg Goldhahn

Professor Morrison:

Ladies and gentlemen, thank you very much for coming on this dark evening. I'm glad that the weather's dry. Welcome to the Annual City Hospital Lecture. We've run this lecture for several years, and for the last few years it's been jointly with the Ulster Medical Society. I'm the chairman of Medical Staff, so it's a pleasure to introduce you to the City Hospital for guests tonight. We've always had a tradition of the City Hospital lecture, of having people at the top of their game on the outside of the curve, that's the good bit of the curve, as they say, so we're delighted to have an excellent speaker tonight who's flown in specially from Zurich to chat to us. So I'm not going to do any further introduction, but I'm going to ask the President of the Ulster Medical Society to chair this evening's meeting and introduce our speaker, so Mary Frances, thank you very much for chairing.

Professor McMullin:

So thank you very much, Patrick. I am delighted to introduce our speaker for tonight, Professor Jörg Goldhahn from Zurich. Dr Goldhahn is the Medical Director of the Swiss Federal Institute of Technology in Switzerland. However, he received his MD in 1997 in Jena in Germany. He did various postgraduate courses and research on teaching, and has done a lot of things in between, but now would describe himself, I think, as a translational medicine specialist. He's had a long research career, and this is interesting because it's mostly in orthopaedics and muscular skeletal disease. He has well over 100 publications and a big research profile.

However, when I was asked to become President of the Ulster Medical Society, and developed a theme, my theme was Diagnostics in the Future, and when I decided on this, I then looked and thought, "Well, if you're going to talk about diagnostics in the future, you're going to have to talk something about artificial intelligence", so I was thinking about this and thinking how I was going to do that, and I did what you usually do on a Friday night, I opened the BMJ, and in the BMJ there was a debate about, "Will artificial intelligence make us doctors obsolete?"—with somebody speaking for and somebody speaking against. So I thought, well this is what we need, so out of the blue I emailed him, the person supporting this opinion. He'd never met me before but I am absolutely delighted to say that he accepted immediately to come here and talk to us tonight about, *Artificial Intelligence Will Make Doctors Obsolete*, so Professor Goldhahn, thank you.

Professor Goldhahn:

Good evening, a very warm welcome from me as well. I'm absolutely delighted to be here. In a previous talk I just found out that my university and society are the same age, so we share some history already, and I want to take you to a journey into the future, which we all don't know. It's all speculation, but I want to take you a bit into some things that may happen in the future. I also have to give you some background information, how I arrived at this topic.

The university was founded in 1855, and it was founded especially to solve burning problems of society. At that age, it was building bridges, tunnels, engineering, so no wonder it became an engineering school. Today we have other challenges, and this is why the university decided to shift some priorities, for instance, in medicine, climate, energy, and just a few numbers to give you an idea about the university itself. It's not a big one, but they consider itself as quite progressive, and in fact they really react on actual development with new competence centres, new professorships or the new directions, and this was the reason that I got a chance to set up a new medical curriculum there, in a technical university, which is a challenge in itself, but also a very nice job, and I got a very interesting job description from my president. He gave me two sentences. He said, "Make it different." "Make it successful—but no pressure!" So that was kind of my job description, and since this engineering or technical school, they always had some kind of medical research, even some Nobel prizes were awarded for research in medicine, and in 2012 the leadership decided to put all those units with some kind of bio in the name, like biomechanics, biomaterials, biosciences, into a new department called Health, Science & Technology, and this department has now grown for quite a while, and in number but also in specialities, and again one of my jobs was to think about a new medical curriculum. And this was also the starting point to think about future technologies, because if you have a chance to design such a [new curriculum?] from scratch, you also have the chance to think about the future, so we were kind of debating, how will the future of medicine look like for our students that graduate in six years from now? And if you now extrapolate a little bit what has happened in the field, in the digital field, in the last six years, I mean, it has turned some fields upside-down. Some simple apps have turned fields upside-down, like Spotify, like Airbnb or Uber. They have just turned the whole field upside-down, and we were kind of debating whether the same will happen to medicine with some new applications, and this is how we arrived into our discussion about artificial intelligence in medicine.

The whole university is embedded in a whole network of partners, and so just to give you a bit of an

idea how we designed our research and our teaching, so the components were all there—we had to put them together. Now, just a few more words about teaching before we enter into artificial intelligence. One of the challenges in teaching is that we all teach outdated historic material, because the normal way is, only things that have been established in clinic will find its way into the textbook, and even then it's very hard to come up with new ideas, and I just put up this very nice reading, [slide] how hard it was to get body temperature established in medicine. It took a physician who was writing a book with more one million observations, a full book, to convince colleagues that body temperature is something we should care about in medicine. So medicine is not known as the most progressive field taking up things early. They are not early adopters for many reasons, and then stuff ends up very late in the textbook. Now, if you want to anticipate what will happen, we have to do it the other way round, so it means we cannot wait until stuff has been established, we have to anticipate what will come, and this is the challenge, and how I went into this field of artificial intelligence, because my colleagues and I, we are convinced that this will play a role in the future. At this point I have to make a disclaimer—I'm not the one developing it, artificial intelligence, but I'm surrounded by engineers, freaks and nerds, who do this!—so I'm embedded in the whole community, and it's made me thinking about it. You may have followed the press and have seen some examples of artificial intelligence. You can almost, every day you can read another headline, especially in the field of diagnostics, imaging, dermatology, in radiology, you can read about studies where technology, especially AI-driven technology, performs the same skill, or even outperforms healthcare workers, and this is also true for the field of big data. It means we have already apps that investigate large amounts of data, or we have virtual chatbots. We can also download this already as an app, and they are used already in treatment to pain, patients in chronic diseases. You can change, I want to have another face, you can just change it with your mouse click. We also learned how we can predict in this case mortality, even track development, benefits from artificial intelligence, and we have already the first FDA-approved artificial intelligence system, so it is happening right now. Last year, an artificial intelligence system passed the final exam in China, so it's really up to a similar level, and here's the list of approved, or companies with an approved indication on the right side, you'll see examples where it is used for. And of course this attracts also a lot of money. All the big players in the field, like Apple, Google, IBM, invest a lot of money there, because they are waiting for the big business to come, so we have, right now a situation which is a bit of, let's say, the Wild West. Everybody's testing what's possible. Can we get reimbursement, can we find an indication, so at the moment, rules are not established yet, so companies try out what's possible, how can we develop? Sorry for a German one, and this is a summary from Eric Topol here, where he summarizes how we can use this vast amount of data that we generate every day in the hospital with our own devices, with genomic characterization; process them with deep learning as depicted here, for instance, get virtual health guidance, so there's a lot of potential and he also exemplified this from embryo to mortality, so kind of a lifelong help by different systems, different artificial intelligent systems, so there's a big promise what we can do with those types of system. One of the areas that I'm interested in is this so-called digital biomarker area, so we can explore tonnes of data that we generate by artificial means. So if you take your fitness app or your watch and read out the data, and look for patterns, then we can generate a lot of additional information out of data that we generate anyway. So to give you an example, if you walk and you measure how you walk with one of the other sensors, you can, for instance, collect balance data, so your sway, left and right. We know this is predictive for falls, as the more you sway, higher is the risk that you fall. It's known, I mean, when you think about it, the drunken seaman. This is now possible with the amount of sensors that we have in all our devices, and this also leads to one of the very important findings—why it is now a topic, because now we have the sensors, now we have the computer power, and the storage of big data to make use of the data. Some years ago, it was simply not possible, because we couldn't store and could not process this large amount of data which we call big data now. And on the right side, this is a figure out of Topol's review for the UK health system, about technology in the future. You see what types of applications he and his team sees in the next few years to come, so it's a very specific expectation or extrapolation, when may we see the first effects of those systems, and you see here, how to make the image interpretation pretty soon to get into medicine, predictive analytics using AI, but later to come. What you see there are very specific assumptions, when we can already expect first applications and when it may enter the system; I would not say the market, but the healthcare in general. And this is what [?], my host, was talking about, so it was a very interesting experience for me, because I gave a talk about modern medicine, and afterwards I was approached by a young philosopher, and she was asking me, "Can we run a project together?" I said, "Yeah, why not?"—and the project was a thought experiment, so for me, it was a very interesting experience. I thought about a study, a biomechanical experiment, and she said, "Let's do a thought experiment." The thought experiment is, "If you realize that a computer can do the job better than yourself, wouldn't you have the obligation to replace yourself?"—because as MDs, we should always offer the best treatment to our patients. It means we should not withhold treatment where we know that this is better than the existing one, and this would lead to a philosophical question. If a computer, or an artificial intelligence system, can make a better diagnosis, than you do as an MD, then you would have to have the obligation to replace yourself. It's a kind of, you see the dimension of the

problem, and then we went into this, and so the project was that we met once a month, and had a coffee together, and discussed this problem up and down, what speaks for the AI system and what speaks for the doctor? Is there a role left for the doctor, or is it just, do we make ourselves obsolete? You know, it came to a certain status, and tried to get it public. We went to the BMJ, and said, okay, "This is what we think about the pros and cons", and had kind of an essay written, and they said, "Okay, no, that's not [?]. We have to make the pro and con debate. We need to put a spark in the community, so set up the fire." That was this only question, who is the bad guy, and I said, "Okay, I work for the bad guy. I know what will come", so I wrote it for you. Artificial intelligence could make doctors obsolete, and so, the original quote that I made was, "Doctors as we know them today, as we know them, will become obsolete eventually." I can sign it. But you may know the yellow press, I heard in the UK, it's not always that accurate, as you would assume, and this is what came out, and this is when I got calls from medical societies, from health insurance companies, questioning what we have done. Just to give you the idea that it was the original sentence, and this how it came out. So I was interviewed by an Austrian radio station, and they said, "Do you really think this?"—and then, they had to ask other experts, and all they said "No, this will never happen, there was always a role for doctors", and I went even further and said, "Okay, and those who are the highest-paid doctors, based on their expert knowledge, these are the ones who we will replace first!" In real trouble, as you know, if you run into real trouble, then you know you have a point.

So just to explain a little bit how we arrived into this discussion, but I will explain a little bit more in depth what we mean and what are the consequences, so you can make your own conclusion afterwards. First, I am putting a definition of artificial intelligence, and this is quite simple, and if you look into the web, find tonnes of definitions. It says simply that the ability of a digital computer to perform tasks commonly associated with intelligent beings. And there are three main criteria which we associate with artificial intelligence, that's the ability to learn, the reasoning and the self-correction—this is what we also would claim for ourselves, and this is what we call artificial intelligence, but important, another important aspect is there are two different types of artificial intelligence, a so-called weak or narrow artificial intelligence, and a strong one, and all what we see today falls into the category of weak or narrow artificial intelligence. It is a very well-defined task, very narrow field of application, so if you do an image analysis that's a very narrow field. It's just, this image analysis. What we know from fantasy films is the opposite, so it's this strong artificial intelligence, so this whatever, your robots who decide upon your life and what to do with it—this is strong intelligence. The confusion comes from mixing these two terms. I'll come back to this at a later stage, but it's very important to put some facts into the discussion.

And then another term which is very interesting is the machine-learning, so it means the learning aspect is one of the fundamental things leading to artificial intelligence, and the vast amount of learning. You see in the background, this is the ETH super-computer. It's one of the ten fastest in the world. I could talk about petabytes and petaflops. I don't even know how many zeroes these are, this is kind of unbelievable computing power, and that's, of course, the basics for the very fast learning systems, and the question is, how does it learn? I want to give you a very simple example, so that you have an understanding how those systems work, let's see whether it works. We have a very interesting challenge. You can see this if you come to the main building at the right time:

(music plays)

So this is Duckietown, and those ducks, they take cars and they drive autonomously, and this is the challenge for students, they have to develop the system and programme the algorithms. So they give an input and tell the cars, you should not leave the road—that's the input they make, and the rest is done by learning, so every time a car goes out of the street, it gets the message, you shouldn't do this, and all the other rules, the cars learn by themselves, by a reinforcement. Just to give an example, what machine learning is, one of the really important things is, you don't programme every step and you don't correct every step—you provide some learning algorithms. In this case, you should not leave the road, and you should not crash the other car—those are kind of the two pre-conditions, and every time something like this happens, the system has to learn, okay, I did something wrong, I have to do it a different way, and then this gets programmed. So the machine itself is learning and not by reprogramming. In the old days, we would reprogramme it, take the car and say hey, a new course, a new route, and we would programme. In machine learning, the machine does it by itself, so that's the difference. And this explains certain different methods how to do this, this huge amount of learning, and those new methods of learning. So you may have heard about deep-learning, so that's a multilayer neural network based on this neuronal networks that you may be familiar with, some years ago, now much larger scale than the machine learning that I explained already. So this capacity to learn, and based on this, you have then this artificial intelligent system, but it requires an enormous amount of learning.

And this explains some of the key features of artificial intelligence, this almost unlimited capacity to learn. We cannot compete with it, there's no way. Some decades ago, you may have heard about the first chess computer who was better than the champion. Now two or three years ago, the same applies to Go, which is much more, a much more decrease of freedom, but that's something you cannot, or we cannot keep up. So the big data handling is another key feature and, as you know, if you go into rare diseases, there is about 50,000 diseases that we know. I don't know all of them, but a machine does not forget a sin-

gle one of them, so there is no data loss unless you have a crash. So it will not forget about any single detail, and it's available around the clock, so we have shifts, we have to be available to provide our knowledge. As long as the computer runs, it's there, and it's remote, and this is a very interesting feature for developing countries. In our system, we have quite good access to the healthcare system. In other areas, remote areas, this is not the case, so it's very interesting, this aspect, to have remote access to knowledge, and it's scalable. It doesn't matter whether five patients or 500 patients ask a question at the same time, so you don't need ten times more doctors, you just have one central computer. Finally, at the end of the day, it becomes cheap. It's not cheap in the beginning, because you have to invest money, but if you run the system it is cheap, because if you save a lot of manpower, it automatically becomes cheap. So what's the question, what prevents us from using it on a larger scale? One of the big questions that we ask is black box—we don't know exactly what happens, and this frightens us. We don't know whether these systems have any kind of intuition left. This is something what especially my colleagues in the article have emphasized, diagnostics is a type of an art, so we associate things. We know things just by referring to some other instance, so not everything is hard coded, and this is a very important argument. Of course, one of the biggest arguments is empathy. I argued in the article, empathy can also be something which might be biased. We all have people that we like more than others, so this may influence your diagnostics. It's good and bad. The machine is not biased per se. It doesn't look at the person in front and say, hey, I like you, I don't like you. Then we have, of course, the number of ethical implications, and I think that's, at the moment, the biggest concern, if you come into a conflict situation, and we know this from self-driving cars—what decision to make? Here, a lot of work is required to come up with the framework for those systems, responsibility, if I get diagnostics made by a computer, who is responsible for the diagnostics, for the decision made? Is it the company manufacturing the device?—or is it the doctor using the device?—not solved yet. Safety is, of course, an issue—who has access to the data? We see it right now in China, quite a lot of observations, so privacy is, of course, linked to data safety, so those are different issues which worry us, and I think this is something what goes into discussion around artificial intelligence.

If we go back to medical doctors, this is a common representation of the roles of medical doctors, so the so-called (?? 0:27:22) roles, where we say, okay, based on the function, as a medical doctor, we have other roles in contact with patients, so we have to communicate with patients quite clear, we have to make decisions together. We have to work together in teams, so act as a professional. We have to learn, lifelong learning. Sometimes we also have to behave as a health advocate. We have to order the wheelchair for the patients, so we have to battle with the insurance company, for instance. We have to lead teams, private practice, or even a hospital, and, I mentioned as well, you collaborate.

Now, AI goes into the middle of this role, so the expert role is the one in danger, and this was the reason that I claimed that the ones who were on most of the money, based on their expert role, are the ones who have to lose most of all in this. You saw this already, so the question is, how do we handle this? You can neglect it and say, okay, we don't care, it doesn't affect us, or we can deal with it, and this refers to my sentence, 'Doctors, as we know them today, might become obsolete', so my conclusion is, not that doctors all become obsolete, but the expert role may shrink, or will be affected, whereas the other roles might be even more important. If the system tells you that the patient has a 30.1358% chance of getting a certain cancer, what does it mean to you and to the patient? How do we translate this to a patient? What do we do with the knowledge generated by an artificial intelligence system?—and so I think that the communication role, or the collaborator role, is even more important. But this led us to a follow-up project [?], together with the philosopher Vanessa, she's now at McGill, so we formulated a few consequences for those different roles, so one thing is, if an artificial intelligence system is available, all our other professional partners may also have access, so this diminished a little bit this distance, this professional distance, between different professions. We have to learn how to handle this, that either patients or other disciplines may have almost similar knowledge to we have. As leader, we are responsible for introducing the systems. We can either push them out and say, we don't care, or we can do a proactive role. As an advocate, we have to care about social justice. If a system said, okay, it's not worse [?] there, this patient will not get this or that therapy. We have to reconsider and say, okay, in this article, this is ethical, how do we handle this?—and as a scholar, it was quite obvious. We knew to learn by ourselves about new systems, and this is quite a challenge. If I want to know something about my iPhone, I ask my daughter, because your younger generations are so-called digital natives. They have a completely different approach to technology, or at least I have. Again, it has an impact on all of those different roles, and we especially focus also on the educated patient. Many patients Google their diagnosis, or they Google things they have heard about. It's not always helpful, sometimes it is, but we have to simply handle this, so we have a lot of impact in our roles as a medical doctor.

Now I come back to teaching how we integrate this, and how we want to prepare students for this future as we see it, and this is the review I mentioned already, so the team around Eric Topol, they prepared a report about the future medicine for all healthcare professions. I can only recommend this, to read. It's freely available on the internet. You can download it as we did, the German article, where a large group, you see this was a really large group, to say, okay, what does it mean for teaching? How can we best prepare our students for it?—and it's not an easy

thing. Now, in the curriculum, our students, they have quite some maths as a basic, so they learn about algorithms, they have some statistics, informatics, so in the third year, we go with them into machine learning. We want to tell them what it is, how it works, so that they can make use of it and can potentially join the development of those directions, and at the same time, we also have to teach them the difference between something in technology-driven and a real patient, so I don't want to leave the impression that I want to replace doctors, and especially not all doctors. I want to replace some functions where I'm convinced that machines can do a better job, and this is usually all what is related to very repetitive functions, and challenges associated with large amounts of data with a lot of learning in it, so this is where, in the future, we will not compete with machines. However, the medical profession is much more than just repetitive, a large amount of data, and this is still left to medical doctors, so to go back to my initial question, no, I don't think that doctors will be obsolete. Certain functions, certain types will be obsolete, and this will change, of course, the profession as we know it, so if you have an image analysers in a remote way, the question then is, do we need a radiologist at every hospital, or can we centralize it, or can we outsource things? This will definitely have an impact, so this is what I'm convinced about, but not that we make the whole profession obsolete. And to train it with our students, to compensate a little bit for the technology-driven approach, we also introduce very early patient contact. And it's very interesting to see, they should take a case history, and you can already see from the expression of the faces, how difficult it is for very young medical students to extract the knowledge that they need for decision-making, for the history, and then they realize, okay, one thing is a textbook, and one thing is a kind of virtual reality environment, and the other thing is having a real patient talking about whatever he or she wants, and you have to extract the information that you need for treatment. So we are fully aware of this kind of two worlds right now, and we have to teach both, we have to prepare students for both of it, and the question is, which direction it will go. I would be extremely happy to review the whole story, then in six years, and say, were we right or wrong?—and what holds the future for us basically. So I want to close here, and would like to open it for discussion, and I can imagine there are a few maybe questions from your side. Thank you for your attention.

Professor McMullin:

Thank you very much. So hopefully that has produced some thoughts for people. Has anybody any questions to start with?—yes.

Dr S McAleer:

I suppose my view of AI is coloured by 2001, the film. I'm just wondering, how do you stop computers acting out of self-interest, rather than for the benefit of the patient?—holding a role for themselves, rather than solving all the problems.

Professor Goldhahn:

With machine learning, you set the initial assumptions, and then the system starts optimizing itself, and it clearly depends on the initial assumptions and the frame you provide, that this should not happen. If you would say, the computer should self-optimise with respect to energy consumption or any variable you set, any computer, of course, will go into this direction, but if you set a kind of ethical framework, for instance, then this should not happen. At the moment, one of the hottest debates of all, how do we define those frameworks?—because some research has already, on purpose, developed bad artificial intelligence by setting an immoral framework, and this is one of the dangers that we are facing. You may see some films in this direction, that you can also misuse the whole story, depending on your input variables basically.

Professor S Elborn:

Thank you, Jörg, it was an excellent talk. The story is very compelling, and very exciting, but what do you see as the bottlenecks in developing this, and making the transition from [?] and maths into GPs looking after patients in the community?

Professor Goldhahn:

A very good question. Again, what we see right now are very specific and very narrow use cases, so people have shown it on very specific indication. If you want to extend it to, let's say, a broader application, first of all we have to have good data to use, and as you know, crap in, crap out, and this is one of the big challenges right now in many different countries and systems, how do you get the data to use?—so that's one thing. The other one is the ethical question about privacy and data use. I personally think we can only see a major development if we are able to resolve this ethical framework, otherwise we are building very limited, on very specific indications, and cannot extrapolate it to larger use, and here, it's going to be very important, we have two options. Either we wait until the big companies have resolved it. You know Google, they already tried to establish an AI ethical board, which finally did not work, or we engage ourselves in the discussion, or we contribute to this and say, hey, we know it will come, we have to contribute, we have to shape it. That's for me, a very important point, one part of the motivation why I went into this topic basically.

Audience member:

Two small things—one is, you mentioned the ethical difficulty, for example, the car, the self-driving car crashing into a bunch of school children in order to protect the occupants of the car, because that's what it's programmed to do, so we haven't overcome that problem, and the second is translating work from big industry that use this already, to medicine, which is

dangerous. Now, you would imagine the very first line of any code in an aircraft, a modern aircraft, would be, "Please do not crash this, a perfectly serviceable plane, into the ground", which is exactly what the Boeing 737 Max did, twice, so if we're going to depend on other super-duper people to give us a steer on how to bring this into medicine, we haven't exactly covered ourselves on the way so far.

Professor Goldhahn:

And this is one of the challenges that we have to overcome, what is the role of the AI system? Is it kind of to find a final say, or is it for us more like the autopilot in the airplane, where I say, okay, I can outsource some of the more repetitive things, but I take the real, the decisions that really matter, I take still by myself. So it's kind of, what's the interplay between the AI system and the physician, which is not defined at all. We discussed right now a project, or we want to look into a physician, AI, and patient, and this triangle, do the patients trust more the physicians, do they trust more the system? Do the physicians see it as a competition, or as something synergistic?—we don't know, right? I'm not aware of any rules that we could apply, so this is part of the whole process, of the whole landscape. My colleagues at ETH, they are mainly engineers. They go for the technological side, but I think the make-or-break is exactly what you say, how can we embed it? If we cannot resolve this, we either run into big problems, or we will never use it, and I think you're completely right with the example you make.

Audience member:

So it's just an observation, and I'm a neuroradiologist, so what we've noticed in radiology, in the early PAC systems which were clearly designed by the background nerds that you describe, of drawing the algorithms, a lot of it was not very user-friendly, and this day, it's still the case with our PAC systems. So the important thing is to have integration between smart players in the background, but actually the [?] and if the systems are designed with a true medic at the forefront, rather than standing in the background, when they deliver this suite that allegedly performs x, y and z, I think that's really important, but just for younger folks, it's already here. AI is already here, it's in my field, and it's rapidly increasing in my field. There's already been two clinical trials, but in the background. It's called Rapid Software, and it's for perfusion and stroke, and it's FDA approved. It's in the UK now, so it's already here, but it is niche market, as you mentioned, but we can't kid ourselves, it's not going to be insignificant.

Professor Goldhahn:

I think what we are lacking right now is the clinical evidence that we can really change treatment, so we can avoid the complication, or we can make it faster, cheaper, better. We are lacking this clinical evidence right now. We know that we have technological solutions, but we need to demonstrate it was safe for patients, because we have avoided some errors, or it will make it faster by faster pulling all the information together, or any kind of added value, and as soon as we have the data, then we have use cases and can introduce it, because it can be argued. We saw in quite a number of other fields, like computer-aided surgery. They have never provided those, added value, and then finally it did not make its way into routine treatment, that you cannot convince any payers to say, I need this expensive stuff, and don't show any benefit. It can be quite a range of benefits to multiple or very different parties, but this is what we have to demonstrate, and then we have a chance to integrate it in a useful way.

Audience member:

I'd just like to make a comment about you saying how this is, as part of being a doctor, and thinking about how we train doctors for the future, and in thinking about that, patients aren't homogenous. They expect different things from their physicians, and so the place of AI in terms of clarifying a diagnosis is just, as you said, one part of being a doctor. What does the patient expect in terms of how you will manage, how you will care, how you will communicate? Some of it has parallels to some of the folks who would say, actually, clinical examination is dead, and of course, that's not the case, and lots of patients don't think you're a good doctor unless you lay hands on them and examine them, and in my field of neurology, the clinical exam is actually quite important, and so AI has to sit within how on earth are you going to manage patient expectation of how the doctor's going to manage the situation with that patient? So how can AI help with expectations?

Professor Goldhahn:

I don't want to argue with you!—but interestingly, all we know is that the chatbot programmes that I showed you, there were some trials going on showing that those chatbots in chronic diseases, help patients' adherence, help managing patients. It is really interesting, because patients know that this is a chatbot, they know it's not the real doctor, and it still helps managing this relation, but of course, only in addition, not replacing. We are afraid of this kind of thing, where you just call a number and say, okay, for your chronic disease, press one; if you have a question on your insurance, press two—we don't want to have that theory. We want to have the personal contact, so it's not about replacing it, it's just supporting it. One good example are the so-called clinical decision support systems, so they give you a hint what to think about, what to consider, and even give you the likelihoods about certain diseases, but you, or we as doctors, have to make a decision based on this, so for me, again like autopilot, or let's say, supporting what we do, and then it may also have a function monitoring. In diabetes, we know there's a place where it can help monitoring, but it will not replace this personal interaction, and I would never argue about this.

Audience member:

Just to follow up on that, clinical decision support, the founder of Babylon just become (?? 0:46:08) provides decision support for general practitioners, (?? 0:46:16) used by trusts, and [he has] said that, like paraphrasing, it would become negligent in the future for doctors not to use artificial intelligence to help them in their practice, and he went further to say that, if that becomes the case, should doctors who do not choose to use it, be dealt with in a different way from doctors who do? I know that's a challenging thing. What are your thoughts on that?

Professor Goldhahn:

From a teaching perspective, I would see at first another challenge, if you like, too much on technology, the question is, how much essential knowledge do we have to have? That's the teaching problem, and the other one is then more for completeness. Do you consult a system just to make sure that you've considered everything relevant, and they're kind of two different things. Sometimes we make the wrong comparison. If you would compare a self-driving car with a very experienced taxi driver, in the best case you'd just reach the level of the taxi driver, not better, because the taxi driver knows your town inside out, but if you have a beginner on the driver's side, then your system could be very helpful, or would outperform the person, so I would say, within the career a young doctor would have, would take more advantage from the system helping him or her to consider things, and say, have you thought about this one, have you thought about this one?—just supporting it?—and I would not say again, for me, it's not the purpose to replace the specialist who has all the expert knowledge gained already. For me, that would defeat the purpose.

Audience member:

I'm just asking in terms of, this is going to be a stupid question, how do computers actually learn? I mean, obviously if you have a Duckietown and they drive round, and it probably senses the same, points out the lines, for example, if you have somebody who's in with crushing central chest pain, going down the left arm, and the computer says it sounds like migraine, how are you then, the next time, that a patient comes in, that you reduce the chance of you making an erroneous diagnosis?

Professor Goldhahn:

So typically, what you do typically is, you take a training set, a data set for training, that the computer can establish in certain relations, and as you said, a certain combination of symptoms will lead in a very high likelihood of this diagnosis, and then that it can do it unsupervised or supervised. Supervised means you label the data, you label all the diagnosis, the computer learns immediately whether the computer's right or wrong. Unsupervised would mean that the computer does not know what is right or wrong, and you do this afterwards, and then you fit basically the model of the computer, and you can come to a very specific fit, so that in this training set, the computer can get the sensitivity and specificity of nearly 100%, but the really interesting thing is then to apply it to new data, take this model and apply it to incoming data, and establish again sensitivity and specificity, and by doing this learning around just teaching the computer, when the computer's right or wrong. That's the way, how computers learn, and the interesting thing is, you can do it with very large data sets. Right now, we have a project where we want to feed in 15,000 patients, and want to establish a model about delirium and the prediction of delirium by combining different facts from the patients, the drugs he or she takes, and all other things, to calculate the likelihood of developing a delirium on the ward. The interesting thing will be then to apply to newcomers, and new incoming patients, and see how good can we predict it, that's basically the way we do this.

Audience member:

So one of the things which has hindered the implementation of diagnostics has really been bringing the regulators on to create an environment which is suitable. You could imagine if you have a self-learning algorithm in a diagnostic making therapeutic decisions, that there could be some regulatory implications? So are the regulators starting to think about this kind of thing?

Professor Goldhahn:

As far as I know, yes, they all have working groups, and there are quite some interactions between the groups working on the problem and the regulators, and in the whole digital medicine field is something where the regulators FDA and [CE Mark?] are, let's say, comparably fast. Usually it takes a while to catch up, but they realize, okay, they have to be much faster. If you go to the website of the regulators, you'll see that they have working groups trying to come up at least for position papers right now, and try to negotiate, so for instance, when Apple got their watch approved as a medical device for one channel ECG, there were a series of talks between the FDA and Apple, so it's kind of establishing the rules alongside.

Audience member:

I suppose my question is related to that. I'm involved, I'm also a radiologist involved in replacing the current imaging system, which has involved quite a lot of contact with companies, and the companies are falling over themselves to sell us various bits of AI which are presented as complete, ready and working, and I appreciated your description of it, it's a bit like the Wild West out there, that's a little bit what it seems like. I also agree that, in order to get it to work as you were saying, we need an ethical framework and a regulatory framework. I'm just wondering, where you see, which bodies that are already out there, which regulatory groups already out there, is that going to come from? Who's going to create these, or review and agree these ethical and regulatory points?

Professor Goldhahn:

In the new department, we created a professorship of ethics of big data, and the expert we could get there, she's now advising the European government in those questions. At the moment, we have to work on all levels to set up some frameworks, and I guess next year, or the year after, there will be a new professor of ethics in AI, and we want it covered, and again for me it's a very important point, because my friends, they are so heavily involved in the technology, and if we have a new solution, but the make or break is to develop the ethical frameworks and the regulatory and ethical frameworks. I think if we have the ethical ones, the regulatory we can derive from this, but the ethical ones right now are absolutely not clear. There are quite a number of working groups, a lot of discussion around, a lot of high level groups who try to put this in place from many different agencies, but other than this area we have a bit of Wild West, nobody knows right now what is finally the, let's say, the final guidance document basically. That's what we are missing right now—I'm not aware of one.

Professor S Elborn:

Eric Topol, who you mentioned a couple of times, has also written extensively about the inversion of the medical model, and that AI and machine learning will democratize not just knowledge, but democratize decision-making in medicine. Have you explored any of the issues that physicians encounter, in that their patients will know as much as they will know?—because the patient will become informed with the expert knowledge that will be then changing the conversations that they'll be having with their physicians?

Professor Goldhahn:

I mean, it goes a bit in the direction from your colleague. There are some patients who are extremely keen on knowing what they may have, what the symptoms mean, and there are other patients who want to use their right of not knowing, just to trust the doctor—you know best what's good for me. I think we will see a kind of diversification of patients. Some of the younger ones, those digital natives, may go in the direction you describe, and others, they still seek the care aspect and the empathy aspect, and don't want to know exactly like this, but what we clearly see are more informed patients, so this is, I think that's a fact. People come and they have already searched for their symptoms. They come with a certain type of mindset, and we have to deal with it. So I made an experiment with my students, that they have to give a presentation about a disease, and I give them an hour time for them, and say, you find out what's important, and it's so easy to get all the knowledge together. You'll find it all on the internet, but you have to process it, and this is limiting, or the bottleneck. If you're in the medical profession, you know how to process the data. If you're not, the question is, can you rely on the websites and all the fake news which is out there, and this is [?], but there are some websites now with quality labels, but it's also at the very beginning.

Dr S Hawkins:

So who do you think ultimately decide on the medical framework? Would it be the experts, would it be the academics, or would it be the politicians, or the big companies?

Professor Goldhahn:

I mean, if you look at the history …

Dr S Hawkins:

Who would you trust most?

Professor Goldhahn:

That's an issue! If you look at the history of regulation, final decisions, you can say, regulations come from politicians with advice from the experts, including academics, and I would favour this way. The question is, how fast are we in this process, and do we have all the facts?—generated by companies, as we see it right now, the companies simply generate facts. They generate the apps, they generate already the whole stories, and then the question is, do we act or do we react? My favourite way would be, yes, we have a political decision on this based on expert recommendation.

Professor McMullin:

So you could put it another way—will it help in any way, despite the ethics?

Professor Goldhahn:

I think so, it will happen. It will happen whether we like it or not. The only thing is, we should get engaged and shape and design it, because as we know from history or especially from industrial history, if things are more effective, if you can save money, it will happen, and we see this, and if it's not regulated here, you saw the CRISPR baby in China.

Professor McMullin:

That's what I was going to say.

Professor Goldhahn:

This is simply, people try out, what is, we have no legal frameworks and we try out.

Professor McMullin:

And I listened to George Daley at the American Society for Haematology, and the whole place was horrified about that, and the only answer, he said, "But it's happened! You can't put the genie back in the bottle. It's going to happen again."

Professor Goldhahn:

And this is what you see, for instance, with Uber. It was just an app, and it kind of turned round the whole field, and now different governments or cities try to regulate this, so they say hey, this has a detrimental effect on our taxi drivers, and we have to regulate it.

This is why you need regulators, to follow always what's happening.

Professor McMullin:

So the other bit is how much the other type of democratization, so we're the doctors, we're the experts, and we get paid more! We've seen nurse specialists moving into our clinics. My nurse specialist can manage my patients in that area probably as well as I can now. If you take out that expert bit, why should I get that? Is it because the rest of it, I did at medical school? There's a great democratization likely to come there, which we will not like, so we're not going to want to be paid less.

Audience member:

[I'm a cardiology registrar and PhD in science at the Ulster University?], and I think, this is the key, where basically medics need to get trained in data science and medics basically have to be more eligible in new technologies, in AI. We basically need to know what the future holds for us, then we basically need to protect our patients, because then we will be aware of what are the potential dangers, and obviously do no harm is the first, the most important aspect of our work, [?] for investing into medics, [?] the answer, how we can protect patients?

Professor Goldhahn:

Nothing to add!

Professor McMullin:

Anybody else? Okay, well I think we've had an excellent discussion, and an excellent talk. I'm going to send more emails out asking people to come, absolutely delighted, and I think that was a wonderful talk, so thank you very much, and I'd like to give you that as a little token of our appreciation, so thank you very much.