Informatics in Disease Management: What Will the Future Hold?

Mike Hogarth, MD

February 2019

Speaker 1:	00:00	It's my pleasure to introduce our last speaker, Dr Mike Hogarth is a clinical professor of biomedical informatics at UCSD and he'll be talking to us about informatics and disease management.
Speaker 2:	00:14	Thank you for having me. and thank you for the invitation. I'm not a nephrologist, I'm a general internist, so I'm generally a danger on the wards I tend to induce a lot of the injury that you have to treat, so prevention is key, right? So, will talk a little bit about that. So I'm here to talk to you about informatics and disease management and what will the future bring. And so let's see here. I'm just going to cut right to the chase, there is actually a paper in Critical Care Medicine, I think last March that shows that they predict AKI onset with up to 48 hours before it actually happened with an area under the curve, of 0.9. so that's a piece of software you should have running today,
Speaker 2:	01:08	then you'd know how many people you'd have to round on in two days , at least that's a benefit, so I'm going to talk a little bit about how that's happening and why that's happening now. So machine learning and artificial intelligence and give you a little background on that. So actually machine learning has been used for a long time, in a linear regression is a very simple maneuver and it's used for predictive analytics. A lot of people wouldn't consider that machine learning today, but it really is the same kind of point. But we'll talk about what machine learning techniques are being used today and their impacts. So what's new? So most AI and machine learning, particularly machine learning algorithms have been around for awhile. And so why is machine learning and deep learning in particular new revolution in the use of AI?
Speaker 2:	01:57	And you see there the transition from this sort of idea of artificial intelligence to machine learning, which you hear about a lot today or now. And then this deep learning. So we'll talk a little bit about all of those. I felt that it's necessary to really, make sure that Ted Shortliffe was given some kudos. He was really the pioneer that started the whole idea of whether computers could, actually, help decision making in medicine. What's interesting is the graphic and grades correct, you'll notice they actually do, reference Ted, and that he had a predictive algorithm, to classify I think, blood disorder what

		really was with the name E-Mycin, you probably can imagine it was actually, to provide advice on, prescribing, antibiotics. And we don't use it today because all those things have changed and the bugs have changed as well.
Speaker 2:	[02:55]()	But the basics are there. That was a rule based system. today the machine learning approach is to use thousands and thousands of records and many variables. And that's possibly because computing is so much better today. so key informatics trends now in 2019 are really what are causing this to really come to the bedside. So HITECH and EHR adoption, as much as we love that, right, we all love it. No, we don't, has made massive amounts of data. I call it dirta because it really it isn't clean data like you're used to in clinical studies. It's really what we call real world data. So I call it dirta but it is useful. some of it, and it's really fuel for machine learning algorithms. The more they have, the better they are at being able to allow you to do either classification of images or patients or predicting what might happen at a particular person.
Speaker 2:	[03:53]()	At the same time, we have high performance computing that is highly commoditized. In the past everybody's saying, Oh, it's looks like the mainframes of the past, not really the mainframes or the past where millions of dollars and, only mainframe ninjas had access to them, nowadays you can get your access on a virtual data center, with a credit card. And it's that simple. And, about six months ago, I took all the CMS data and all I wanted and did very simple, computation as to what were the top procedures and what was the most expensive hospital in the country. terms of charges. The two top. And, I did that with 11 machines.I just let them run for six hours, but the actual computation took less than a minute.
Speaker 2:	[04:45]()	So I could run the computation every minute for six hours, and it cost me $3 and 96 cents on AWS to do that. And it would've been just as cheap in any of the other cloud providers, you just had to know what to do, and the biggest thing was just getting the data, you know, the 9 million rows of data. Anybody want to know the answer to the question? Stanford, you know who the second hospital was, you'll never believe it. UC Davis Medical Center where I spent 25 years of my career arguing that we weren't that expensive iI turns out the data doesn't lie. we are that expensive. so I encourage you to do the same thing. You can do it for $4, so you can do the computation. you can't do it in Excel. You probably can't do it on the laptop. it's a lot of data but, you can do it in the cloud easily. at the same time we see machine learning engines and algorithms that also became commoditized. So the tooling around machine learning is very

simple. Now, programmers don't have to know, they don't have to be PhDs in computer science to do this. it's in a variety of languages and most of it is as a service as we call it. And I'll show you a little bit of that. so you see the trend there about 90% adoption of EHR and the massive growth of data, it's pretty amazing. The growth is about 50% increase in healthcare data per year annually. And we're headed for, probably around a thousand exabytes and we are going to 2300 exabytes by 2020. That's just enormous amount of data that can be used.

Speaker 2: 06:16

A lot of it can be used by machines to help us make decisions. So that has garnered the attention of a lot of tech companies. So the three big ones, Amazon, Google and Microsoft that provide sort of cloud services all have machine learning platforms that you can use tonight if you would like. They are very easy to on board. You just have to have a credit card, you create an account a few shekels and they have some experiments you can run tonight that are very easy to just get introduced to using machine learning. So very easy to do. They see the value of AI in healthcare in particular you see the $150 billion estimate of a market share for this kind of technology, and in addition, 5 billion of that being sort of service oriented machine learning as a service, so that will come to your EHR soon. If it hasn't already, why is that such an attractive thing? Well, you know, this is the usual pain point, right? That is 2015 data out of a Health affairs article. That's real data. That's national health expenditure at over $3 trillion. and that's 17.8% of GDP. So, we're flat on the percent GDP. You would think, that's good news, no because the GDP is increasing. So actually the total cost healthcare is continuing to rise. So I think right now it's probably headed towards, within the next 10 years, something like, $5 trillion. and even though it's around 16, 15, 17% GDP, just the GP is going up, you can see they're $250 billion of wasted, money on just processing the transactions to get paid, which is kind of weird. We have some medical errors that happen.there's been a lot of controversy on what constitutes a medical error, I run the death registry for California. There's actually, no cause of death called medical error, which was a highlight of a article a couple of years ago out of Hopkins that there really should be a physician as a cause of death or the system or something. So we can start classifying these. Everybody is using bad data for this and that's all. We don't have good numbers, but the estimate is around $20 billion. and of course 2,500 deaths a year that we can prevent perhaps, and then a lot of dollars at waste. So, tech companies see this and they see the expenditure here and they see a big opportunity to reduce cost and take a little bit of that home as a consequence.

Speaker 2:	[09:00]()	So you see that, you see Microsoft, IBM Watson, Google and DeepMind. Of course they focused on kidney disease and AKI. The latest articles on that though were about how Google somehow got the NHS to just hand over, identified data, to do that work, which wasn't really Okay. And so they retracted that. But nevertheless, the idea I think was a good one just, the execution was very poor. Microsoft taking on blindness with machine learning and then IBM Watson, which we heard a lot about, so you'll see a lot of research into ML/AI models for various tasks. It's a very hot area, you know, the VA wants to use DeepMind to prevent kidney disease, just like NHS. I doubt they'll be handing over personalized data like the NHS did, but nevertheless, there's a big opportunity there.
Speaker 2:	[09:51]()	Notice some of those articles, there's one that just came out last week. So this is a very hot area in a variety of journals. So, and yes, there's a magical thinking in this space as always anytime HITECH gets involved, this is the new shiny object that has a lot of money attached to it. There's a lot of magical thinking. And of course, IBM Watson is the classic. well, I won't call it vaporware, but nothing's really happened with IBM Watson, when it comes to healthcare very much. And I have some theories about that, that I think they just pick the wrong thing to do. But quite frankly, anytime that you've ever tried to put a computer against an expert in healthcare, they can be equivalent, but they're not better. Well, you really want to do is have, the computer help people like me, a general internist, be as good as somebody like you when it comes to the diagnosis, treatment, you know, that kind of thing.
Speaker 2:	[10:43]()	And that has been demonstrated time and time again. In fact, that was exactly what Shortliffe demonstrated. The E-MYCIN system with the general internists was as good at choosing antibiotics as an infectious disease expert. So we can be experts as generalist if we have this kind of help. What they did there was they tried to, help oncologists who are already experts in what they do. So just a bad choice, I think. of course, IBM, you know, didn't feel the pain as much as MD Anderson. they canceled the project almost $40 million paid IBM by MD Anderson. It was originally, priced at $2.4 million. So again, very unrealistic, magical thinking, et cetera. But let's not throw out the baby with the bath water. This is a great little paper, from a couple of years ago. looking at, machine learning, and trying to classify bone marrows the flow cytometry, wise, and AML and what they trained it with 200 samples from disease and non deceased patients.

Speaker 2:	11:54	And this thing was, able to correctly classify, predict, 90% of the cases that would become relapsing AML. And it was 100% correct and distinguishing normal from abnormal AML bone marrow. That was the machine by itself. So that clearly is helpful. So how does it really work? You know, again, it's not magic. You have training data usually, or you have data that it has to ingest. There's a training algorithm that uses a variety of techniques and we can talk about those. I'll present some of them. it creates a model that basically after which you have the model, you just present new data and it either is classifying that new patient or that new image or what have you, or it's providing a prediction, and that's basically how it works. there are really three types.
Speaker 2:	12:46	There's supervised learning, unsupervised learning and reinforcement learning. When it comes to machine learning approaches. The first supervised learning is the most common. You have labeled data, there's direct feedback provided usually in terms of points, and the system is almost Bayesian in its approach of trying to become refined and, providing a better outcome every time. Unsupervised learning is essentially, you give it the tub of data and it clusters it and you tell it the variables you're interested in and it'll create clusters and it'll put particularly patients, for example, in different clusters. And so the new patient comes in and gets put in a category, with that, I say it's kind like GWAS you know, association is not causation. So you might have some weird associations, but, if it's predictive and it works,, I'm not sure you don't want to use it.
Speaker 2:	13:37	But nevertheless, that's the one the most black boxcich of all of them. And then there's reinforcement learning, which is the most sophisticated, and that's, where you really have a piece of software acting as a person that is being rewarded or punished based on an action it has taken. And so I'll show you it's called the hit and train kind of maneuver, hopefully you don't do that with children, but, with a computer it's Okey. So here's just the supervised learning training data set and you train it. once it's trained, you put new data in and it gives you suggestions or predictions. Here's the unsupervised learning. Again, no person in the middle and, so much less effort. You can take massive amounts of data. So a lot of people like this and thereafter, a lot of the data that you're producing in your hospital for this kind of thing.
Speaker 2:	14:26	And then reinforcement learning, which has a lot harder to do, a little bit more AI like, where the computer is almost becoming the decision maker. So I wanna just mention artificial neural networks because people talk about those all the time and it's

really the sort of the step before we can talk about deep learning. So it's an AI system loosely based on biological neuron networks and brains. It's not really an algorithm. There are algorithms that are used in that setting. The setting is really one of a framework, kind of a computational framework that has this notion of nodes that are acting like neurons that have inputs. And once you meet a certain threshold and the inputs and different weighted inputs, it fires off itself and it can fire off into another, neuron. it can be supervised or unsupervised.

Speaker 2:	15:18	These in terms of it's use for machine learning. Why is it called artificial? I love history, so I always try to figure out why is it called something weird like that? Well, it's because, in the beginning they actually created the first neural networks with circuitry and that's actually a picture of the Mark 1 perceptron as it was called. I think they quickly realize it's much easier to do this with software, right? Then wires that are plugged in and sometimes get loose. And if anybody has an electrical engineer here, you know what I'm talking about, a cold solder is, the biggest pain you could ever have so, rather than the machinery, which is very fast by the way, as computers became more prominent programming languages were easier to use to implement this.
Speaker 2:	16:03	This became artificial neural networks, which is what we do today in a computer. So that's why they're called artificial neural networks. the idea all came out of a neurophysiologist, the mathematicians sitting around trying to figure out modeling how neurons work. And that's become this artificial neural network, which is very useful tool. Actually. I'm not sure that's how the brain actually works. I think that's still yet to be sort of laid out. But, this is a model that seems to work for computers to really emulate what, people can do sometimes in terms of decision making. Deep learning is just multilayered artificial neural network where the first layer captures the raw data, as an input layer. And then there's multiple layers underneath that have a function that modifies the data. It's not just a simple weight and those functions can be changed.
Speaker 2:	16:53	So deep neural networks can be different, and you're transforming those features all across the network and propagating that out. So a lot of times it's just simplifying the data and, getting it to be, more, categorizing of usually images. So this is usually applied to images. And so that's where we see a lot of deep learning being applied today very successfully, and then there's something called convolutional neural networks, which is applied to imaging, and it actually came out of the idea of trying to figure out a hand written notes. So this was out of

AT&T bell apps. And you see, they're actually a photograph of the system that they originally developed trying to figure out three, eight, four as they moved the camera around a real time and you'll see it picked up eight and four.

Speaker 2:	17:40	And as you move it around it eventually picks up three, eight four. and it was inspired actually by how the brain and the visual cortex of the brain actually works and how it's laid out in layers. So that's why they tried to emulate that. Yann LuCun did that.There's a significant resurgence of interest in CNNs because, we have a high availability now digital image data sets that are annotated. So they're large training sets that you can use, in a supervised learning mode. And you have high performance computing with something called GPU. These are graphical processing units, the ones that all the kids want, and their computers for their Nvidia, gaming and what not. they're very useful for parallelized computing when it comes to this kind of thing. So now we have, AWS, Azure, you can get a GPU cluster instead of a CPU cluster and it'll perform much better and it can do a CNN much faster.
Speaker 2:	18:37	So then it's just a matter of how sophisticated how dense the images are and how much computation power you have, as to be able to create the model. So this is why your phone can recognize your face. This is what it's using. So this is how it works. Input images, feature extraction, feature vectors, they become numbers, and the machine learning algorithm will categorize those. And then new images basically get matched to something in the archive, by vectors, you know, so it's not really doing images is transforming those images into numbers and feature vectors as they're called. So here's a conventional deep neural network in dermatology that, actually, did pretty well, they used one point, almost 1.3 million images from 18 clinician curated repositories of dermatologic, photographs that were curated.
Speaker 2:	19:35	So these are annotated by clinicians, they tested a CNN based classifier on 376 new cases, head to head with 21 board certified dermatologists on two scenarios. So it was scenarios and these were real biopsy proven lesions by the way. So differentiating benign S K from basic carcinoma and, there were 135 of those. And then differentiate malignant melanoma either with a simple photograph or, with a little higher resolution camera differentiate melanoma from a nevus and the performance was very good. You see the red dots there in those curves are actually the dermatologists. So the system was as good as the dermatologists were. And again, I'm not saying that this should replace every dermatologist, but have we seen

dermatologist sometimes misdiagnoses or miss things? Yeah, so I think this could assist, physicians and dermatologists at diagnosing. I don't see that replacing, there's a lot of things we do that a machine can't.

Speaker 2: 20:38 And obviously this is just two very specific use cases or scenarios where there were actually 2000 different conditions that were represented in those other images that it ingested. So, here's a CNN, for, echocardiography, all they were trying to do here was not make a diagnosis. Obviously that's, a little harder. they were just trying to actually image, classify, the views correctly, you know, so parasternal long axis right ventricular inflow, short access to much level, these kind of views that are traditionally done by the echocardiographer and so they did this with 200,000 images and you can see there the video views, it was almost 98% accurate and still images 92. The four board certified echocardiographers had 79.4% accuracy. So it was actually better than they were, which is interesting. Maybe they need more coffee or something. I don't know.

Speaker 2: 21:40 Maybe they're almost too dark or the fellow was asking too many questions. I don't know. It was surprising to me that it actually outperformed, the folks that do this every day, Here's deep learning with retinal images, trying to detect diabetic retinopathy, the same paper describes a glaucoma, and changes related to that. But when it came to this almost half a million retinal hemorrhage images, you can see the area under the curve 0.95 and 0.936 for diabetic retinopathy that is, vision threatening diabetic retinopathy or referral early on diabetic retinopathy. Very good at doing this, which is why if you go to the, optometrist today maybe go to ophthalmologist, you'll get a funduscopic photograph and there's a piece of software. It's actually doing this, commercially now, and helping the optometrist find llithum and things like this.

Speaker 2: 22:34 So, anyway, this is where that's headed this is reinforcement learning as I mentioned, that's a little harder. and this is called the AI clinician. So this is actually a paper, I think from last year, in Nature describing the use of a data set that's available that anybody can download. That's real data that's anonymized, that came out of Boston Hospital by 12 years worth of ICU data. All of it. I mean, the notes, every alarm, all the labs, all the reports, there's narrative texts, reports and this system, use that, and then they went to a new set of data, an they had, information about what had happened, in the clinical course. and they then applied AI clinician to it.

Speaker 2:	23:32	And whenever AI clinician said, to do a particular thing and there was a matched action by the clinician in that data set that was a much lower mortality than other cases. So in essence, they did pit a clinician prospectively. But if it were to assist somebody in an ICU setting, it could potentially reduce mortality to some extent. So and all it was doing was providing advice on fluid resuscitation and presser administration. This is a gentleman who's coming to UC San Diego shortly from Emory. His name is Shamim Nemati, they trained a validated on 31,000 Emory ICU admission, 60 features. So a massive amount of data. And they constantly, are getting data from the ICU Philips monitoring system., and they created a sepsis prediction algorithm with an AUC of 0.85 . And, it tends to recognize sepsis 4 to 8 hours before, the current mechanisms do, and they've actually implemented that and integrated that into their care, at Emory.
Speaker 2:	24:47	So there's actually a room where this is being shown to a set of nurses who then if they see something happening and the system alerts them, they'll contact the nurse at the bedside to start the bundles, the sepsis bundles. So, with that said, regulation, is the FDA ready for this? It turns out they kind of are, they came out with 2017 guidance on decision support software, which includes AI, but AI is a little different. It's more black boxich, but basically they view it as a device, and, so anything that's intended to diagnose, cure, mitigate prevention, treatment of condition falls under FDA oversight. Anything that's low risk, weight management, mindfulness to, et cetera. And not under FDA oversight. They distinguish between clinic decision support and patient decision support. so you see there, that's the guidance, these of the recent decisions for AI based systems, which they've made.
Speaker 2:	25:40	Some of them are de novo, under the regulatory pathway for de novo. So they don't have a preexisting, algorithm out there. Many of them will be under that kind of, regulatory pathway. And the legal profession has also caught on. So this is the American Bar Association and I'm particularly, to see who done it, you know, civil and criminal liability for poor outcomes when you're using AI. And you just got to love lawyers, look can't and to what degree may physicians delegate diagnosis, tasks, AI systems without fear of exposure. And to what degree can we not use AI, and not be liable. So, you know, you might be liable either way. All right, so this is a nice article. Talks about the future of disease management, in the world of artificial intelligence, machine learning, this is a great little paper out of Stanford.

Speaker 2:	26:29	What's life going to be like in 2030 and this is a comment out of that paper. It could be great, using abase application, here's, it could also be go horribly wrong. After reviewing how much your treatment is costing the hospital, I'm going to recommend assisted suicide. Or you see there the machines are walking around handing food and money to us so what's the moral and ethical issues and I think we're going to be struggling with this in the next 10 to 20 years. What does it really mean to have a job when a job is, replaced by a computer or an AI machine, that kind of thing. Everybody should buy a robot probably and then let the robot work and then we all play golf or whatever and still get paid.
Speaker 2:	27:13	One thing I know for sure AI has come as part of society now it's here to stay. I actually ordered a pizza from Domino's in San Marcus here and it was delivered with that car. It was fascinating. I went out to the curb to see it coming. It was kind of funny, I'm sure me and my son will probably think that's really odd. You know, dad's watching the car coming, of course it's going to come, it's going to deliver a pizza. yea, it's like this is how the world works now. But it really did. It delivered the pizza. It opened the back window. You slid your credit card in there and the pizza was nice and warm. And then we took the pizza and then it just started driving away. So if you want to order Domino's pizza here in San Diego, go to San Marcus dominoes and wait for the car.
Speaker 2:	27:59	It's fascinating. All right. when are we going to be replaced? That's often a kind of a question. What am I? smart machines. Sorry, I'm a little over here. But, surgeons, right? I noticed this thing. This is a paper from AI experts, it's a viewpoint and that's where that graph comes from or that low display and says surgeon. And I thought that was kind of funny, because you know what we think of surgeon. So then I put nephrologists in there and I said, surgeons will be replaced within 50 years, but the nephrologists, the quintessential internist, It'll take 300 years. You don't have to worry, but chances are you will have apps to help you. Okay. But you won't be replaced for a heck of a long time. But surgeons, I think their, time is coming, so they better switch their specialty. All right. Thank you very much.