NEJM AI Grand Rounds Podcast TRANSCRIPT

Episode 1: Dr. Euan Ashley on AI, Genomics, and Cardiology

[00:00:00] Welcome to NEJM AI Grand Rounds, a new podcast from the NEJM Group that explores the deep issues at the intersection of artificial intelligence and medicine. I'm Andy Beam, co-host of NEJM AI Grand Rounds, and I'm delighted to welcome our listeners and to bring you our first conversation. I'm Raj Kumar your co-host of NEJM AI Grand Rounds for our first episode.

[00:00:21] Our guest is Professor Euan Ashley of Stanford University. Dr. Ashley is a pioneer. In 2010, he led the team that conducted the first clinical interpretation of a human genome, and he holds the record for the world's fastest genomic diagnosis. He even has a Guinness World Record to prove it. We'll have a wide ranging discussion where Dr.

[00:00:41] Ashley will share the stories behind these feats, his experiences applying artificial intelligence to genomics and to cardiology, and its views on whether and how AI will change. Dr. Ashley will also describe how AI seems able to extract information that human experts may not see. For example, [00:01:00] cardiovascular risk factors from retinal, fundus photographs or laboratory values from echocardiograms.

[00:01:05] We ask Dr. Ashley to speculate on the prospects of a pan-diagnostic non-invasive test enabled by AI. We also asked Dr. Ashley for his most controversial opinion, what evidence is needed to translate AI into practice and how physicians should think about AI. Without further ado, here's our conversation with Euan Ashley.

[00:01:27] So Euan welcome to NEJM AI Grand Rounds. Well, thank you for having me. It's wonderful to be here. A question we like to ask all of our guests is, could you tell us about the training procedure for your own neural network? How did you get to be where you are and what were the, the training data points that led you along your path of gradient descent?

[00:01:45] Well, I mean, I think like any kid, I was exposed to the science fiction of potential future of artificial intelligence from an early age, whether it was talking cars or spaceships or really anything that would push [00:02:00] us towards a bright new future where computers were helping us out. So I think that was the starting point in the formative years.

[00:02:06] I was always a computer nerd. I didn't think with my first computer that had little rubber keys that I was programming in basic that I would ever really be running neural networks or building artificial intelligence. Nor did I ever particularly try, although I built little computer games as a kid and tried to do

[00:02:22] somewhat smart tax software for my dad when I was a young teenager, but none of that could quite be called artificial intelligence. I was however running deep learning models, or not very deep, actually pretty shallow learning models, but neural nets in the 1990s, which I think dates me quite a lot. Actually some of the earliest large-ish scale medical data that we had available of database of patients and electrocardiograms and exercise data.

[00:02:49] So I was running some fairly shallow neural networks in the 1990s and I think that that interest in data and interest in statistics in general then came out through my [00:03:00] professional life and a PhD in, in Oxford. The latter half of it really became much more about data and statistics as I found that there were data I was generating in the wet lab that I really didn't have a good model for and didn't really have a good mechanism for analyzing.

[00:03:15] And then I jumped a across the pond to to Stanford where I've been ever since. And of course here surrounded by the history of artificial intelligence and the burgeoning tech of Silicon Valley that both predated and post-dated me and then this third coming I guess of, of AI that is the one really impacting medicine that we're all so excited about.

[00:03:34] So those are kind of the steps along the way for me. I've always been a nerd. That much is very clear. So, Euan, could you tell us how you got interested in medicine? Yeah, that one's actually pretty easy. My earliest memory is telling people who asked that I wanted to be a doctor, so I came from a medical family, and I th I don't think there's any doubt that my dad, who was a GP and my mom, who's a midwife, influenced me in a certain direction.

[00:03:59] Although when they were [00:04:00] actively influencing me, they were trying to push me away from medicine. I think it was just seemed like the thing I was put on this planet to do. I was always the guy who my friends came to when they scraped their knees and when there was a first aid course, I was first in line and generally I would go home and get the answers from my dad and bring them back.

[00:04:17] And so I was, I was just always in a doctor mode and it just seemed like the only thing that I was ever. It's certainly the only thing I ever considered

and one of the bonuses of life is that I, I found I could combine the interests I had in my spare time, which was computing and data, uh, with my main profession, which was medicine, which initially I did not plan or did not think would even be possible.

[00:04:42] That one's a bit more longstanding and, and it is somewhere, somewhere buried deep in my soul. is the idea of, I think trying to help people in, in whatever way I can, and that, that's probably what led me to.medicine So Euan, we want to focus on some of your research. You've been a pioneer in the clinical translation of the genome.[00:05:00]

[00:05:00] You wrote a book, the Genome Odyssey, and you've really helped steer the field. Going back to 2010, uh, if I can take you back that far, I read one of your papers, uh, while I was in grad school, clinical assessment, incorporating a personal genome in Lancet, uh, and it had a big impact on me. So I'd like to start there.

[00:05:18] Could you tell us about that? Yeah, well that's, that's very kind of you. It's always nice to hear that. Well, it's always nice to hear anybody who's read your paper, but to hear it had more of an influence is an even, even bigger deal. So, um, yeah, I mean it was an adventure and for us, really where the genome adventure or our own, if you like, genome Odyssey began, it was late.

[00:05:37] 2000, 2008, 2009. At that time there had been maybe four or five individuals on the planet who had had their genome sequenced. And one of them was my colleague Steve Quake, who was chair of the bio engineering department here at Stanford at the time, uh, with Russ Altman. And I was meeting with him really just for.

[00:05:57] Uh, standard, like let's set up a genetic [00:06:00] seminar, I think because our new chair, Mike Snyder was coming and we went to talk about that, and then in the end ended up talking about his own genome. I mean, he had it on the screen in front of him, and that to me was just a mind blowing moment where he started showing me the ASTs, GS and Cs of his genome on a screen in front of him.

[00:06:18] But in particular, he pointed out certain variants and genes that I recognized that actually were associated with inherited heart disease. And so quite quickly that meeting turned into almost like a clinical discussion where we started talking about his family history of heart disease. And he mentioned, oh yeah, you know, my family has been asking me to go see a cardiologist anyway.

[00:06:36] And we realized in fact, that his cousin's son had died suddenly age 19, a really tragic event and uh, one that lo was left unexplained by the postmortem and and medical establishment after. And really that whole thing came together to the point where we essentially invited him to our clinic to be assessed, to be checked out for this potential family condition.

[00:06:58] And at that moment, I think we both [00:07:00] realized that he was about to be the first patient, really to walk into a doctor's office with his genome on a hard drive and somewhat scarily to me,that doctor's office was mine. And we had to think pretty quickly about what we might do when we had a patient who had an entire genome in the clinic room with them.

[00:07:18] And, you know, it's something that we'd been thinking about for a while. Anyway, the, the cost of genome sequencing was coming down from the 3 billion of the Human Genome project, you know, a hundred million for Craig Venter and all the way down from there. And we were around $50,000 at the time. So certainly not something that you were gonna immediately jump out and pay for much like a haircut, but it was very clear

[00:07:38] where this was going and it was heading rapidly towards the thousand dollars genome, and we felt that somewhere around that mark a genome would finally become useful for medicine. And, and those of us who are interested in data and taking on data challenges should really start to think about what that might mean for medicine.

[00:07:53] Because in, in my world, in in cardiology, . You know, we're used to a cholesterol panel, like five numbers, and that's pretty [00:08:00] easy to deal with in a short general practice visit of 12 minutes. But when you have 6 billion data points of a genome, well that's a little bit different. And so also, how do you bring to bear

[00:08:11] the entirety of the genetic literature and really everything that is known about associations between genes and disease and variants and disease. Uh, how do you bring that into, uh, an algorithm that can be essentially deployed in the context of one patient in a even primary care practice, but certainly in, in a clinic?

[00:08:28] So that was the question we set ourselves and, and, uh, fortunately I had some good friends who joined in. Steve and his team joined. Of course, Ross Altman and his team joined Atul Butte was here at Stanford at the time. He joined in. They all really, basically as soon as I called them, just said, that sounds like fun.

[00:08:43] Tell me where to be . And so all these people showed up with grad students and postdocs and suddenly we had a team of 25 people and really started to work just for six months, flat out on doing nothing but analyzing Steve Quakes genome, for his risk of every disease we could think about. And yeah, to [00:09:00] cut a long story short, in the end, the paper you mentioned was what came out and it really became a pivotal moment for, for many of us thinking this

[00:09:07] could be the future. If genomes really are available, then the sorts of approaches we were trying to put together would be the ones that we would need to be able to deploy this in medicine. So yeah, it was, it was a really exciting moment and, and a great example of teamwork. I mean, very much like a very, and there was a group from Harvard who joined as well from the Personal Genome Project.

[00:09:25] So really by coastal effort, uh, really focused at, at trying to gain insight from one person's genome. So do you see the genome interpretation challenges today as fundamentally different than those that you solved back then, or as refined or maybe even harder in some way than when you did this for a single genome back in 2010?

[00:09:46] Yeah, yeah. No, I think they're pretty parallel, actually. I mean, we, we have a lot of advantages now, of course, compute power, uh, significantly greater sequencing. technology is, is better and faster and cheaper. Our [00:10:00] ability to understand the darker elements of the genome, if you like, the harder elements, definitely significantly improved in that time with better chemistry and better technology.

[00:10:10] But essentially, it's a very similar thing if we think about the, the three tracks that we thought about back then. Basically thinking about risk for rare disease, so looking for rare variants. Thinking about common complex diseases and polygenic risk scores, and then thinking about pharmacogenomics,

[00:10:24] that's approximately the framework that we would adopt today and that we actually are adopting today on the front lines of medical genomics. . And so that part hasn't really changed much, but we have so much more power. And one of the major elements of that reduction in costs is manifest not actually an individual level, but a level of the entire population.

[00:10:43] Because back then, I mean, we absolutely were looking for control populations in the single digits. You know, we, there were 10 people whose genomes we sequenced that we could look up in the clinic at the time. Although

our, the, the sequencing we were doing there was for specific focused [00:11:00] genes for maybe three genes for long QT syndrome or seven for cardiomyopathy.

[00:11:05] When we got those reports back and they tried to tell us, well, this is a new variant, this might be disease causing, then that was based on a comparison with a control group of about a hundred. Anonymous blood donors. That's where we were back then. 100 anonymous blood donors. We knew nothing about their ethnicity or their ancestry.

[00:11:24] There was only a hundred of them, and we literally knew nothing about them. And if we didn't see it in those a hundred, then we kind of concluded that it was probably novel and might be more like to be disease-causing. We were so naive, , you know, and I think that's a lesson really in, in, in every field.

[00:11:38] But genomics particular, just when you think you've got somewhere, Need to be really humble in, in the face of data and, and now we're in a great situation. We have hundreds of thousands of individuals from varying populations that it's not diverse enough yet. I'm sure we'll talk about that, but it's much more diverse than it was.

[00:11:54] And we can now with a new variant that comes up in a patient in clinic, make reference to hundreds of thousands of [00:12:00] individuals who have no known disease or do at least do not share the, the disease that is under question for the patient in front of us and use that as a reference to decide how likely the variant we've just found is likely to be disease causing.

[00:12:12] So from that perspective, we've really moved into another dimension, which is very exciting, but in other ways it's very parallel to what we were doing back then. . So it strikes me that I think we have a lot more resources now and a much richer understanding of reference variation across populations, as you pointed out.

[00:12:30] And this allows us, I think, as you're saying, to filter out some of these signals that when previous era might have been labeled as disease-causing or pathogenic. So for some of our listeners, Who've been involved with ordering genetic tests or working with genetic testing laboratories. I think as with 2010, in 2022, they'll still see a pathogenic label on the report.

[00:12:52] Right? So this the spectrum, the way in which we categorize genetic variation, as with 2010 and in [00:13:00] earlier eras, is still one centered

around this concept of pathogenic. I'm curious, you brought up polygenic risk scores. Do you see us moving in medicine away from the pathogenicity label towards a more quantitative understanding of risk?

[00:13:15] Something like penetrance or conditional probability of disease given a variant? Is this even feasible or is this just a dream right, that we have? That of course makes sense, but it is very hard to realize in practice and will this really help patients? Yeah, I mean, these are great points and it's, it's really a great question.

[00:13:34] One of the things I'm fascinated by is the extent to which the framework we lay on the world with history and, and culture of humanity, uh, really reflects what we might otherwise find from the underlying data. So obviously we like to go from a. Data science perspective towards unsupervised approaches first and everything that we do to try not to bias what comes.

[00:13:55] We try to, to have the data speak to us in terms of patterns, but when we live in the [00:14:00] world, the world outside of our very controlled scientific environment, let's say moving into a medical world by a medical world, then we really are saddled with. The history of each of these technologies, and I think the structure to which you refer in particular is our clinical framework for classifying variants, where we have a wide range in the middle where we would say this is a variant of unknown significance, and that we cover the vast majority of variants, or certainly the vast amount of our confidence or in variance.

[00:14:29] And then there'd be a a 10 or maybe 20% at either end where we'd be able to classify a variance as likely pathogenic, likely, Ben. Or pathogenic and, and benign and. Kind of five category, uh, classification system is one that has grown up really as genetics has grown up in, genetic testing has grown up.

[00:14:50] And so it was, it was very interesting to then, and you brought up polygenic risk scores. So the history of polygenic risk is completely different. So that comes from genomewide association, which [00:15:00] is really something. That statistical genetics, uh, really invented. It was really an invention of science, which was run.

[00:15:07] Uh, and really, I think also with the crisis of replication in that prior historical period, the kind of over a correction that represents, that Bonferroni represents, was really the standard. Uh, and so there was this incredibly cautious but highly quantitative approach from statistical genetics and science that is then combining in time and place with this

[00:15:29] five category classification of clinical genetics, and I am absolutely with you that we need to merge those and move forward kind of together to a point where we add more quantitative rigor to the history of the clinical world. But more application than real world application to the, this sort of statistical genetics history with its history in scientific association studies.

[00:15:55] And I think that will take us to a really interesting place. And so one of the things I've been [00:16:00] really excited about in this last year at Stanford, we have now replaced our, uh, panel testing- traditionally for patients who might have inherited cardiovascular disease. We would send panels of, and these days it's more like 30, 50, 70 genes.

[00:16:13] But now every single one of those is actually a genome in the background. So the backbone for those tests is a genome, and we read out the specific genes that are known to have the relationship with the disease. But that opens the door to this polygenic architecture for every single patient because clearly these two things don't exist in a vacuum.

[00:16:33] And in individuals you have an entire genome. And across the site frequency spectrum, you have variants. So you have variants of large effect, your variants of small effect. The small effect ones will add up at some level. Potentially to impact the larger ones. And they're, they're all essentially on a spectrum.

[00:16:49] And I think we absolutely need to think about quantification, uh, all the way into the rare disease part of that spectrum, which is where your question began. And I think it will be possible to do that. I think [00:17:00] some technologies will help us and it won't happen immediately. We obviously are already far along because of these population cohorts.

[00:17:06] That's tremendous. I think we're about to hit a. Where Variant effect maps, cellular effect maps, where we can use gene editing technology. Crispr, we use CRISPR X in our lab, but there's a few other approaches to that where you essentially do saturation, mutagenesis, create every single variant in a cell, uh, or multiple cells, and then phenotype them would allow you to, to have a pre probability in.

[00:17:31] For every single variant in the genome. And I think that is a technology along with moving to the millions and tens of millions of people in population cohorts that would allow us to get to the point of a very reasonable pre probability that would be ready for us when a patient comes with a new variant.

[00:17:48] Switching gears just a little bit. So you published a paper in 2022 . In the New England Journal of Medicine and it was titled Ultra Rapid Nanopore Genome Sequencing in a Critical Care Setting. [00:18:00] So I think ultra rapid there is well deserved. It is the key word. I heard you set some type of Guinness World record with this paper.

[00:18:08] So I have two questions. First, do you still hold that world record? And second, could you tell us what the major bottleneck was that you solved and what you think the next frontier is? Yeah, of course. And, and thanks for the question. I, I do believe, to our knowledge, we still hold that record, but, uh, I'm happy to be in a race with anyone.

[00:18:26] I think pushing technology forward is, is what we do. And we were inspired towards this by the amazing work of Stephen Kingsmore who. I learned recently. He is also from Scotland originally, although he was brought up in Ireland. So he, he has an Irish accent, unlike mine, but born in in Glasgow, same place as me.

[00:18:42] And he really ha for a decade, has focused on how fast can we bring genomic diagnosis to the bedside of critical care patients. In his case in particular. Neonates. And I think that is just a, an incredible contribution and a really important group of patients. We're, we're focused on, on many different patients.

[00:18:59] Of [00:19:00] course, those with rare or undiagnosed disease have been particularly the recipient of the benefits of genomic technology. But those who are critically ill with rare or genomic disease, it's, it's hard to think of a more deserving population of our time and, and efforts back in the early 2012, around that time, uh, Stephen had set, uh, a record around 50 hours.

[00:19:19] I think we were moving around that time somewhere in the 40 to 50 hour range, trying to bring the computational component of that, squeeze that down. He then halfed that down to 26 hours and, and he was, uh, awarded a, a Guinness World Record for the fastest genomic diagnosis, which he then broke down to 19 and a half hours.

[00:19:36] And that was kind of the starting point for an idea that we had, which was to apply Oxford Nanopore sequencing to this idea of rapid genomic diagnosis. Lumina sequencing, of course, because the clusters have to grow and because of the basis of, uh, sequencing by synthesis that there's a certain time over which that has to happen.

[00:19:55] And in general, it's, it's around 18 hours and so it's very hard to [00:20:00] sequence faster than around 24 hours because you need the majority of that time for the sequencing element. And obviously then you get the rest squeezed down as much as you can. With nano port. It's a completely different technology. Of course, these are small protein pores through which, uh, DNA falls or is, is kind of breaked as it falls through.

[00:20:18] And Oxford Nanopore, the company based in the UK, had built a machine that the chief technology officer had informally called a tank , which, uh, basically has 500,000 of these pores that can be deployed at once. So you can basically, with this machine, generate a one x genome. So cover the genome one time in 90 seconds.

[00:20:39] I mean, 90 seconds is just mind blowing, . So to generate enough data to do a kind of clinical grade genome, it would take about an hour, but that's significantly faster than we could achieve with, with our aluminum machines. So we were, we were excited by that. And the technology, the Nanopore technology is as well as the Pacific Biosciences long-read [00:21:00] technologies have.

[00:21:00] Coming on leaps and bounds in the last few years from the perspective of cost and affordability, but also improvement in accuracy. And so we really felt the ti the time had come to, to try this out. And so yeah, we, we put a team together. What we realized almost immediately, and this speaks to your.

[00:21:16] Bottleneck question is that while we or the machine could generate the data very quickly, the the big box that makes a lot of the noise of a 7 47 aircraft right next to the compute box was not capable of keeping up with that rate of data. And so in fact, if we ran the standard algorithms on the box that came with the machine, it took 18 hours.

[00:21:38] Interesting. It was 18, but took 18 hours to do the basic compute- to do the calling. And so clearly we could sequence in an hour, but 18 hours of computation, just to get to the point of thinking about filtering the variant call file was, was not gonna work. So a very talented, uh, grad student, you know, Goinca uh, in from double E here at Stanford, to, on the task of, of thinking about how to [00:22:00] deploy the individual algorithms were faster obviously by breaking up the packets, sending them up to the cloud, making sure we were running individual containers for the individual algorithms that were optimized

[00:22:11] with their compute architecture according to the algorithm. So for example, the variant calling is mostly neural net based, and so that's optimized for GPUs. But the alignment, uh, is, is a compression algorithm that is actually optimized for cpu. So if you want this to happen kind of in real time, you need to optimize those different processes and the cloud obviously is.

[00:22:31] Is really helpful for that. So she built an approach to doing that and we tested it out, realized that it could actually, in the end do almost real-time alignment and then variant calling, uh, came shortly after with the alignment phase instead of having to wait 18 hours. It was just 30 minutes after sequencing.

[00:22:49] So really incredible job in in solving that one on her part. And then we passed this over. This was also a collaboration with the Google Deep Variant team and the folks Benedict Patton and his team at UCSC. [00:23:00] And we used their long-read Collar for Oxford Nanopore, which is a combination of three different things.

[00:23:05] Pepper margin deep variant goes by, but essentially mostly neural network based. There's an HMM in there as well, but mostly neural network based models, so again, optimized on GPU and it would appear every week we had a meeting and they would come and they'd speeded it up a bit and prove the F1 statistic.

[00:23:21] And so it would really was an exciting moment where every few weeks it seemed like this was really getting faster and better. And finally we were kind of ready to, to start and we opened the doors and we just were getting calls every day from across the intensive care units at Stanford. And it was great to see this was something people were really interested in.

[00:23:38] And when we were able to start returning diagnoses in under. 12 hours. I mean, I, I think our first one was 19 and then they dropped mostly to around 12. Halfway through we pulled the bar coding step, which we realized we didn't need. And then the, the last five genomes were all blood to diagnosis in around eight hours.

[00:23:57] So, so we're very happy with [00:24:00] that. And we indeed were contacted by the, the Guinness World Record Group, and they awarded us the record for fastest genome sequencing, so that that was a. Moment and, and very different group to talk to than our, our usual, uh, communications with journals around scientific findings.

[00:24:17] Uh, very different. So is the remaining gap between where the world record now is and like a 10 minute sort of standard lab diagnosis, just waiting for Moore's Law to do its thing? Or is there another bottleneck that you think we have to solve before we can get there.? I know that we can go faster. Uh, first of all, and I think there are others, I hope, and I'm sure there are around the world who are already trying to beat the record there.

[00:24:41] I think that clearly there's always a balance between accuracy and time. So, you know, you, you could certainly do a much less accurate genome, maybe one that still makes a diagnosis. But overall is much less accurate. Maybe it's lower coverage. You know, each one of the steps has kind of a high accuracy, slightly more intensive approach [00:25:00] or a slightly lower accuracy, uh, faster approach.

[00:25:03] So you could certainly make some balanced decisions there. Blood to answer in 10 minutes might be a stretch, but I , I, I think we could probably take another few hours off this and where we plans are underway to do that. And, and actually that's right across, cuz even the wet lab element, there was a rapid kit that we, in the end didn't use because we wanted to prioritize accuracy.

[00:25:23] Now the kit is, is at seems to be at a point that really is much more accurate, so we should be able to use that rapid kit. So even from the wet lab portion, we should be able to, to strip off a few hours here and there I think, and, and at least get down to, to three, four hours. I suppose there's a question eventually as to how fast do you need a genetic genomic diagnosis?

[00:25:42] And we've already had some people say, well, you know, often two or three days is enough for me. And I think what I'd say to that is, that's great then, then that's what we, we should have a system where you can dial the diagnosis to the time that you need it. But there are clearly cases and, and Stephen Kingmore's more's work has shown this, where the faster you can get a diagnosis,

[00:25:59] the better [00:26:00] it is. There. There are real changes that can be made, and we've told stories of patients with cardiac conditions who were suffering multiple cardiac arrests, little babies in the early point of their life, and we had no idea what we were treating. And if you could say you can have a genetic diagnosis whenever you want it.

[00:26:15] When do you want it? Well, I'll say I want it right now. You know, like, I have a baby in cardiac arrest. I, I want the diagnosis now I'm reminded of the Jeff Bezos quote that he said, I don't know what people are gonna want in

the future, but I know they're gonna want it faster and cheaper. Yeah. , I think that's exactly right.

[00:26:30] Yeah. Uh, this fits. Perfectly with, with where we are in, in genomics. So yeah, I don't, I don't hear anyone in the NICU telling me that I should just, uh, slow down and that we shouldn't bother trying to speed this up more. Right. So I, I, I would like to shift gears just a little bit, and I was struck by, uh, your responses to Raj's questions about sort of what it was like to be in cardiology in 2010.

[00:26:52] Mm-hmm. and I, I'm, I'm curious if there's a similar sense of excitement now. With the things that are happening in the rest of artificial intelligence and cardiology, do [00:27:00] you think we're in a similar moment in time, were there lessons that you learned during this sort of new technology of genomics in 2010 that you're carrying forward with you now?

[00:27:09] Could you gimme the lay of the land? Yeah, I think there is a, a similar excitement to that. Uh, I do think that. And this is not unique to cardiology, but I think specialists, doctors who go into cardiology do tend to be quite tech focused and, and quite tech interested. It's, it's a subspecialty with technology that is prominent and has been in the history of the profession.

[00:27:30] And so I think that we are all very excited about what's possible. And then I think there's also low hanging fruit for, for medical AI in cardiology. And we've seen that we have. In major hospitals, you know, a million electrocardiograms sitting around already at some level over read by humans, uh, with labels that can be good training data.

[00:27:52] We've seen examples of course in imaging as well. Similarly in some of the room work in that area is an example. So I do, when talking to my [00:28:00] colleagues, see a similar kind of excitement about what is new and what could be possible that isn't now. I mean, and, and how we can augment what we do today with this new technology.

[00:28:10] And it's exciting cuz I think your mind, IM immediately goes to the far future. And I think that that, that, looking back you do get quite far. Over the course of 10 years. I mean, that's Bill Gates's quote, isn't it? I mean, he's like, most people overestimate what they can do in a year and underestimate what they can do in 10.

[00:28:27] And I think that that's a, that certainly was my experience with the genomics kind of revolution in, in that sort of time period. 2010, I think we

were able to do Steve Quake's genome, and I think we thought, especially with all the calls we were getting, like within just a few years, like everyone will have their genome sequence, of course.

[00:28:43] Didn't happen. Um, but if I look back 10 years later, I mean, we have come a very long way. I mean, talking about the potential for a hundred dollars genome now and talking about really a situation where we can talk all the way across the site frequency spectrum as we just did. So if we, if we put that into [00:29:00] the, the mold, we are now thinking about medical AI,

[00:29:02] we're at a very similar moment as you observe where we've got these initial proofs of concept. Uh, but we haven't yet really got to the point where we've had major implementation and it may be possible to go faster with ai. There's, there's a few reasons that we could discuss, uh, for that many of the.

[00:29:19] Challenges though are going to be similar. And, and as always, I think they're somewhat historical and cultural rather than technology related. Yeah. So, so that's great. I, I do want to explore that sort of open space that you mentioned. Uh, but first for the benefit of the audience, I'd like to anchor on a specific project that you were involved with so that we can sort of flesh out some of these details.

[00:29:37] Yeah. So the, the paper in question is Video-Based AI for Beat to Beat aAssessment of Cardiac Function that, um, you and some fellow Stanford researchers published in 2020 in Nature. Um, I know I have trouble recalling papers that I wrote even last year, but I'm gonna put you on the spot and ask you to give us a TLDR from that paper that you were involved with two years ago if you don't.

[00:29:55] Yeah, of course. Uh, and this is really, uh, exciting work and a, and a great [00:30:00] collaboration with James Zou, who's a colleague in computer science and biomedical data science here at Stanford, and led by David Ouyang, who was one of our cardiology fellows now faculty member at Cedars. And the principal was that we, we'd seen some really interesting applications of AI for static images.

[00:30:16] So it's the first place you would go in medicine, the chest x-ray, think about a CT scan of the chest, and we'd seen some really nice applications of AI approaches to those. But of course we're obsessed with the heart. And the heart is, is marked by the fact that it's in continuous motion. Essentially, you know, 3 billion times in your lifetime your heart is gonna beat.

[00:30:35] And so the technologies that we have for looking at the heart really have to account for its motion. And while in the tech world, AI for computer vision had really moved into the video world, and, and even if you think about some of the work that was being done to analyze. Soccer games or football games or basketball games had clearly moved into the, the vid, the realm of video.

[00:30:56] Uh, medically I was still very much focused on [00:31:00] static images, and so that was, was a lot of the inspiration that, that David took. But, but we didn't have to look far because just as I mentioned a moment ago, we had a million electrocardiograms sitting around. We had several hundred thousand images using ultrasound of moving hearts.

[00:31:14] And so the issue at stake was how do you scope this down to something that may be achievable? Uh, because these are, uh, many of them, and most of them in fact, have human labels already. So the basic premise is easily set up. You have great training data, you have good videos. So then it's all about the practicalities of the file formats and getting them in the right place, and how do you divide the data set and then also what other information do you have on those, those individuals?

[00:31:40] And we do use, uh, very commonly on almost every patient this metric called the ejection fraction. It's simply the, the fraction of blood that is ejected per beatta out to the heart. And you measure it by essentially using the scroll wheel on your mouse to move through the video. Stop at the frame where the heart is maximally full.

[00:31:59] Measure the [00:32:00] size and then, uh, continue to the point where the heart is maximum contracted and then measure the size and kind of project that into three dimensions and turn that into a fraction of blood that you believe is being ejected per beat. We thought that was really an ideal place for, for an AI to, to intervene and potentially go faster and do.

[00:32:16] Better from an accuracy perspective than, than a human. So that's, uh, the work that David presented and, and perhaps not, not surprisingly, given the history now of, of, uh, biomedical imaging ai, we, we were indeed able to build that train that. and demonstrate that it, it worked better than than humans do a little bit.

[00:32:36] Uh, it also had certain other advantages. It takes so long to do those tracings that whereas the AI can operate, you know, almost instantaneously, it can look at multiple frames and multiple heartbeats. And was that the example

in that paper of. Uh, that we demonstrate for patients who have an arrhythmia called atrial fibrillation, where each beat is actually a slightly different shape and size, and the manual approach is just to pick [00:33:00] one.

[00:33:00] But of course, the AI can look at every one of them instantaneously and give you the answer for each beat of the heart. So there's clearly a few areas where we can really play to the strength of the AI and that was the essence of what we presented in the paper. Awesome. Thanks. I'd like to ask, um, another important contribution of this paper was the public release of a lot of echo data, if I'm not mistaken.

[00:33:20] So, right. I have kind of an institutional question is how do you get Stanford to sign off on that? Is this east coast versus West coast sensibilities , because I fought this battle many times, but. So far unsuccessfully. Yeah. Well, I think that it's, and I, I, um, would not take credit for this personally. I think that, uh, it, it is a testament to the foresight of the leadership here at Stanford, but also to the foresight of certain individuals.

[00:33:47] And, and David who led the paper, led the release of the data, which we were as PIs. We were absolutely behind, uh, some of the, the credit needs also to go to our AI and medical imaging group here at Stanford, who really worked [00:34:00] with the IT group from the hospital and worked through many of the issues in terms of, um, de-identifying imaging in a way that the institution was, was comfortable with.

[00:34:11] So that a lot of that groundwork had been laid by the folks in radiology, uh, and we were able to kind of jump on top of the work that they had done, uh, and get the institutional approval to release it. And to your point, I think it's. So important when we see what has been possible in computer vision and, and the tech side, the leaps and and bounds that they have made forward have only been because they've had access to, to imaging and, and the genomics world that we talked about earlier as well, is another one where they've really got the idea of data sharing early.

[00:34:40] and that's been to the benefit of the entire community. So I think it's still in pockets, uh, happening in for medical AI and imaging. Um, I would love to help make some connections with folks here who've worked it through to see if we can help on, on, on your side. On the East Coast, we have, I think one example from a Boston Hospital in Mimic, [00:35:00] but that's a great example, uh, that that's all compared to the data released by Stanford.

[00:35:05] I think a drop in the bucket. So I'd like to ask you a question about a follow up paper to this one. And it was Deep Warning, Evaluation of Biomarkers from Echocardiogram Videos. So I like this paper because it was. Obvious but not trivial to me that diagnostic tasks were gonna be capable of being done by algorithms given sufficient training data.

[00:35:25] I think what surprised me and a lot of other people was the capability of the same models to extract things that may not be in principle extractable by humans. And I think this paper is a good example of that, of reading lab values directly from the echo. So could you tell us a little bit about that paper and whether or not you found that also surprising?

[00:35:44] Yeah. So this is work from a student also that I share with James Zou, uh, Weston Hughes, who was actually presenting this morning at our, our, uh, computational biology meeting. And, uh, with this, this paper actually came up, uh, briefly in our discussion this morning. [00:36:00] Great. And interesting work for just, I think e exactly the, the reason you mentioned, uh, we spent so much of our time, I think, in the medical AI world.

[00:36:08] Essentially using human labels. So it's, it's back to this question of the, the framework. We as humans have kind of already put on the world and in some other work we've generally tried to look beyond that. If we think about the general area of clinical studies, we think about. Hard endpoints first, and that starts with life or death,

[00:36:26] So in general, if we think about clinical trials or work in the clinical arena, the thing that is really undeniable is, is someone alive or dead? And so that's the hardest endpoint of all. And so whether we're testing a drug or we're testing a diagnostic or a screening tool, anything short of that is going to be

[00:36:43] somewhat of a surrogate, but of course, that's the world that we live in. Hopefully it's a meaningful surrogate, uh, or a gold standard diagnostic. So in the work with the electrocardiograms, there's two approaches. You could train a model on the final diagnosis, if you like, [00:37:00] on when, how long the patient lives or dies

[00:37:03] but much more commonly and much more easily, you can train it on the human labels that are there. And so that's generally been the approach and certainly the, the first approach. Even if eventually you, you're sure that you want to take the model beyond, uh, training the, the kind of proximate surrogate marker, you're probably gonna start their enemy cuz that's what you have.

[00:37:21] And so, so much of our work is based there, but where I think it becomes really interesting is to go beyond anything that has been created by, if you like, by a human as a human label, and to see other signals in here that we just don't. Perceive that can predict really interesting things from the medical record and just having an infrastructure or framework where we can have all the electrocardiograms and the echocardiograms and the basic, uh, elec electronic health record data in one place allows us to ask questions like Weston asked in this paper.

[00:37:54] Basically deploying the model, it was possible to look at some things that were, first of [00:38:00] all, reasonably predictable. So there are cardiac markers. Troponin is a, a marker of cardiac damage, BNP. The, even though it's brain natriuretic peptide is really a cardiac marker. So not too surprisingly he was able to predict those from the echocardiogram.

[00:38:14] But I think as you pointed out in the question, much more interestingly, he was also able to predict hemoglobin, which would be to diagnose anemia essentially. And while most doctors could start to think in their brain of ways of connecting the heart and low hemoglobin, it's not immediately obvious to anyone exactly what you know, which element of the echocardiogram is helping this algorithm predict the hemoglobin.

[00:38:38] And, and I think that's, that's what's so exciting about the application of these kind of technologies to medicine is I think where it can go beyond. Human label and historical human label to predict things that that are buried, signals that are buried in there that we just can't even perceive. I'll follow up and we'll wrap up this segment by asking you to speculate wildly.

[00:38:58] Um, so given sort of your work [00:39:00] on echoes and some of the work that others have shown about the information, content of things like retinal fundus photographs and the ability to protect cardiovascular risk factors from those, do you think that there's some, at some point in the future we could have a pan-diagnostic non-invasive modality where blood draws are kind of like bloodletting from medieval times and you have this

[00:39:21] the tricorder from Star Trek where you can get a high resolution read out of someone's physiology without having to poke them. I would love, I spent my life hoping for a tricorder, you know, I think they even have XPRIZE competitions, don't they? For people to, you've got Scotty down, so, right. Well, yeah.

[00:39:36] You know, I have one advantage there, . My mine's a slightly more authentic accent. , although it's probably not this, my Scottish friends would probably tell me I sound quite American these days. But, uh, uh, I, you know, the opposite is also true. I think that's a, that's a dream, isn't it? I and I, I do think. I am not certain yet

[00:39:56] if we can get that far. Because the [00:40:00] depth of molecular data that there is, but the depth of, of molecules, if you like, the number of molecules that are circulating in blood is such that I think we're pretty long way from being able to replace everything. You know, in another, with another hat on, we're talking about omics and moving beyond genomics to measure all the proteins and all the metabolites and you know, to think that we might have a non-invasive approach to measuring every single one of those when we're still

[00:40:25] thinking about how we do it invasively seems quite a leap, but is there a a really vast range in which we still have to, to operate where we can take non-invasive data and find signals buried in them for things that are currently invasive? Yeah, I, I totally. Agree with that and I'm really excited by it.

[00:40:45] I think also your phrase is sort of pan- diagnostic. I think one of the things that, the dreams that we have, if you like, as futurism people, we love to, to think about, uh, the future is that any of the data you have and when, anytime you [00:41:00] have a test, there should be a computer brain in the background.

[00:41:04] Essentially with knowledge of all of whatever it was every e ECG ever done and knowledge of every ECG you've ever had, and that knowledge should be brought to bear on the test that, that you just had to, to maximize the ability for preventive care, for intervening before you get disease rather than waiting till you have it.

[00:41:23] Uh, and I think that, Future is not so far away. I think that is, is quite achievable, and especially if we start to build infrastructure for it now and start to think as we move forward that that's going to be possible and that's how we should arrange our data, then it's, I don't think that's too much of a leap at all.

[00:41:39] Great, thanks. So we will now transition to the lightning ground. I, I hope you're ready. So the rules for the lightning round are, you can, uh, respond in no more than two sentences, so two sentences or fewer for each question that you get asked. Are you ready, Euan? I am ready. Should AI be explainable? AI does not have to be explainable, [00:42:00] but if it is, that probably helps. Euan,

[00:42:04] what's your favorite novel? Um, my favorite novel, um, the, well, I'm currently re reading an autobiography, a biography of Sylvia Plath. That's not a novel. Uh, so the Bell Jar , I guess. Wow. Dark horse . Um, uh, do you think Twitter is a good medium for meaningful scientific discourse? I think it could be, but your use of the word discourse would lead me to say strictly no.

[00:42:39] Do clinicians need to understand machine learning to contribute to machine learning projects? I think they should understand the basic principles of machine learning, but they do not need to know the details. Will machine learning or AI reduce costs in healthcare? I think they definitely could reduce [00:43:00] costs in healthcare, absolutely.

[00:43:01] Through early diagnosis and intervention. Are preprints, Annette Scientific good? Definitely. If you could have dinner with one person, dead or alive, who would it be? Barack Obama. All right, and now to the conclusion. Euan, so we've talked about a lot of things. You've led a nascent technology sequencing into the clinic over the past few decades.

[00:43:25] We've spoken about challenges of AI and cardiology, and challenges ahead of instrumenting and implementing AI in the clinic. How should clinicians think about the impact of AI on medicine? Well, I think the impact is only beginning to be told. Uh, but I think there are lessons, and we've discussed many of them just in the last little time together here, that can be learned from the application of new technologies just in the last decade or so.

[00:43:52] that would really help us move things forward quickly. I think we're at a really interesting moment where we've demonstrated there is power and [00:44:00] clearly we can demonstrate that computers can predict things. Both that, that humans can and do it faster and probably better, and predict things that humans can't.

[00:44:09] So I think the cases made that we should, as a community, move towards. The best way to integrate AI into medical practice. And so this is the moment we're at, like, how do we do that? And let's make sure we fund that properly. Cuz that's gonna be more expensive than those initial studies where you just gotta release the data and get the data trained and, and have people beat up on it, you know?

[00:44:31] But the next stage is all about clinical trials essentially, and following patients over time, preferably prognostically, so that you know, and that's much more expensive. A nice example was the work from Suchi Saria and

her group, uh, recently published that I know you're both familiar with, but moving, uh, a sepsis screening diagnostic AI into the real world.

[00:44:54] And then not only demonstrating that if you deliver that data back to [00:45:00] doctors, the patients who are found to have sepsis do better. So you're showing a medical outcome, but also integrating the feedback from the individuals who received that data into how it worked and, and that's exactly the framework that we are thinking about in our cardiology AI work.

[00:45:16] We're, we're looking to get it into the workflow and then work out about the, essentially define the cycle of human intelligence to artificial intelligence. So the training data is almost always. coming from, from humans. So that's the AI. Then we've trained the AI, we've put that into medical practice, and now we're in a loop that that is inseparable because now you're augmenting the human.

[00:45:37] So the next round of training data is actually the augmented data, and so you never again, really will have a situation where you have completely clean data. So I think we need to study that cycle by sort of asking questions of the human intelligence, which is going and talking to people and asking questions of the artificial intelligence, which is explainble.

[00:45:56] So just as a follow on and deliberately provocative, uh, will [00:46:00] AI replace doctors? . Yeah. Uh, I mean, I feel like saying yes, just to be provocative, but, uh, obviously I, I put myself out of a job. No, I think that AI will allow doctors to focus on the reasons they went into medicine, which are the human ones. I think that AI can allow actually AI somewhat paradoxically, and, and maybe this is surprising to some, can, can take the computer out of the consulting room in a way and allow the two humans

[00:46:29] that are part of this interaction to make a closer human bond, and I think medicine will be better for that. AI restoring the humanity in medicine. I love it. Um, so I have a companion question to Raj's. So you are someone who has had a quantitative background, a medical background, um, who has navigated the introduction of genomics.

[00:46:49] Do you have advice to early career clinicians who are maybe still in medical school or residency who are perhaps slightly more pluripotent than some of our senior colleagues are? What should they be [00:47:00] thinking about? What courses should they be taking to prepare for AI in. medicine? Yeah, I, I mean, I think it's, it's an exciting time to be someone in medical school on a early stage of training cuz so much is, is gonna change.

[00:47:12] And I think this future really is gonna be in the hands of people who have that dual training and at a minimum of dual understanding. I think your question earlier was, was a really good one. It's like, how much do doctors need to know? Most doctors don't need to know much, but we're going to really benefit if a small number, uh, know a great deal and really straddle the two worlds where they're thinking methodologically, they're thinking about developing new methods.

[00:47:37] They read non-medical AI papers. I mean, one of, I, I love to read autonomous driving AI papers on, on archives just for fun, my spare time. It's, it's, you know, I think that there's so much to learn and, and some of. Augmentation methods are, are inspired by the work in physics. So it's you, we get into these little boxes.

[00:47:57] Uh, and I think so, so I think [00:48:00] back to the, to the question, you know, I, I would absolutely encourage. As much exposure to the methods element as as possible for those who have that inclination and to realize that the power they have, because actually you're not writing these tools from scratch anymore.

[00:48:15] You're really walking out how to, to use those tools and then understanding what the parameters are, you know how to tune the hyper parameters is, is really a and, and that's a task. Better done by someone who has a deep domain knowledge. So, so I do think it goes both ways for, for the folks who are predominantly on the method side, uh, the, the, the PhD scientists, the computer scientists, and the data scientists, for them to understand as much as possible about.

[00:48:40] The medical application will also be really helpful. And so certainly we have a variety of approaches here. Uh, our bio design, uh, course for example, at Stanford is one where we're mixing people from the, the business school, from engineering, from from the school of medicine together so that they can learn from each other.

[00:48:58] And I think that that will be a really [00:49:00] important part of this going forward. Will machine learning exacerbate healthcare disparities? , that's a really important question, and I think I would say if we are not careful, uh, it will, uh, but there is a potential, first of all to minimize that. And there is the potential to realize a disparity com compressing element.

[00:49:21] Uh, which is to say that bringing the technology to everyone could improve the medical care of, of everyone. So, so, To deal with that in two parts.

Obviously the first part is clear and we're all very familiar with if, if the training data we have has bias in it, then we're gonna be training algorithms that have bias.

[00:49:40] And so we have to be very, very careful about that. And I think there's an increasing understanding of that, which is step one. I was very encouraged by some of our recent searches. For faculty members here at Stanford where a large number had prioritized this element in their career. The, the idea of really understanding [00:50:00] and eliminating bias in, in artificial intelligence as we start to apply it to medicine.

[00:50:04] So, so I'm, uh, encouraged compared to where we were a few years ago, but I do think we need to be very careful and we probably need to do even more than we're currently doing. . For me, the exciting point is actually that technology can actually level the playing field a bit. I mean, we are in a situation now where, you know, there's six or 7 billion phones, most of which we would really call smartphones.

[00:50:26] You know, around the world there's a significant compute device in the pocket of most of the adult citizens of the planet, which is an exciting moment when you think. What that device can do, and obviously deploying diagnostics that could help eliminate disparities globally is one of them. So I think this intersection of digital and AI is a very interesting

[00:50:48] place to look for, for really bringing technologies. And, and, and you know, an obvious one is if you can get a, a Harvard level, uh, radiology read, uh, from a a, a village in Africa on your chest [00:51:00] x-ray, you have eliminated a disparity in the world. And so I think there's opportunity there that we should take account of whenever we can.

[00:51:08] Thanks. So another opportunity for controversy here, Euan . Um, what is your most controversial opinion, or said differently? What do you think the orthodoxy of our field is most likely wrong about? Yeah, that's interesting. I mean, I have a lot of controversial opinions. Pick your favorite in general. Um, I'm, what I'm wondering is, um, do, how many of them are specific to, to AI?

[00:51:34] Um, . Oh, I have a lot of controversial opinions by soccer and , you know, about medicine. Um, even the ejection fraction we mentioned earlier, which is this measure, uh, that we used in the, in the, in the nature paper earlier was, you know, I, I, I am not a fan of that in general in terms of, of ai. I think one of the things I just mentioned of ogo, I, I, I sort of try to keep an eye on AI in other, in other spheres, particularly tech.

[00:51:58] Of course, [00:52:00] I, I look at the investment. Just to take two examples that are local to us here, but you know, but that Google and Tesla have in, in AI and thinking about real, real world AI on the Tesla side, and obviously search and human understanding on the Google side, and I look around. And I look at what we do in medical AI and I just see that they just exist in completely different dimensions.

[00:52:25] Dimensions of investment, uh, dimensions are just the sheer number of engineers, dimensions of the sophistication of those groups. And I, I really wonder why. We as a world are spending all that money on, uh, you know, why are the best engineering teams, the ones that are thinking about how to make more money through advertising clicks, uh, and not how to make our patients in intensive care units better.

[00:52:52] So I suppose that is marginally. Controversial. You know, I think that the challenge is like, what do you do about that? I mean, we live in a [00:53:00] world where people are rewarded according to, to what people are willing to pay for, and because people are willing to. Pay a thousand bucks for a smartphone, then there's a lot of money in the coffers of the tech companies to, to hire those engineers and, and make them even better.

[00:53:14] The way our healthcare system works, it, it doesn't really allow anyone anywhere. I mean, there's no healthcare system anywhere in the world with like hundreds of engineers devoted to medical AI . But if you took a step back, or if, let's say another way you land as an alien on the planet, you might wonder where the priorities of a population are,

[00:53:32] you know, when we have hundreds of engineers focused on clicks and advertising, and we have only a handful in our healthcare systems focused on how to better integrate these exciting new technologies in into patient care. So may maybe, maybe that's probably the, the best answer for now, . So, final question and we'll end on a, on a upbeat note, what are you most excited about in the next five years for medical AI?

[00:53:58] Well, I think it's this idea of [00:54:00] implementation, isn't it? I mean, we've come to this really exciting moment where we can, we know the computers can make useful predictions, uh, and, and it's time to deploy it. And so that's what I think is, is exciting. But I think it will take much of the next, you know, three to five years to really build momentum in that arena.

[00:54:16] Because medicine historically changes quite slowly and it changes through rigorous trials that are blinded and that show outcome changes. And

then those trials are reviewed by large groups of wise world people who then write guidelines. Uh, realizing I'm now in some of those groups. Um, and those guidelines are then produced by professional societies and then slowly,

[00:54:42] opinion changes and then payers change and that's the rate of change. And then the regulatory bodies obviously are there as well. And so I, for that reason, I think things change much more slowly. You look at the, the, the speed at which our tech companies change the hardware every year. They have, they have new things and, [00:55:00] and, and then look at how fast our hardware, let's say, have a

[00:55:02] defibrillator has changed it. It doesn't change at that kind of speed, but there's, there's actually no reason why it wouldn't, other than the fact that medicine has to move more solely because of regulation and because of rigorous trials are required and they take time and money. So I think that five years probably a good

[00:55:17] time period in which to think about the future really changing. And I think over that period we'll start to see true implementation. We'll start to see patients diagnosed early because of the AI, and intervened upon and lives saved. , and I think that's really exciting. You think about what starts as a technology that we deploy on our local computer and then on a server, and then for a population in, in a hospital, and that eventually lives are being saved because of an algorithm we develop.

[00:55:44] I think going back to the next generation and thinking about potential inspiration for a field, uh, I think that that's it right there. All right. Well, I think we'll have to stop there. So I wanted to thank you. I know you're very busy. I found the conversation to be very thought provoking and [00:56:00] engaging, and I'd like to thank Euan for being on NEJM AI Grand Rounds.

[00:56:03] Oh, thank you so much. It's been an absolute pleasure. Thanks. We're grateful to Euan Ashley for joining us. And we're grateful to you for listening to NEJM AI Grand Rounds. We hope you found Euan's comments as insightful as we did. We hope you'll join us next time and please follow us on Apple, Spotify, Google, or wherever you find your podcasts.