Can predictive AI reduce health disparities in orthopedic surgery? This UK researcher aims to find out

Artificial Intelligence is transforming health care. The promise of this technology is enormous and is already being realized to increase the accuracy of diagnoses, promote patient engagement, increase efficiency in health care and lower costs.

It's even being used to identify patients at risk of disease and predict patients who might be good candidates for medical procedures.

Done well, AI tools can help ensure patients with the greatest need for orthopedic surgery are prioritized for care, and help reduce health care disparities, says Luke Farrow, an orthopedic and trauma surgeon and clinical researcher at the University of Aberdeen in Scotland.

But without proper considerations, "you can ultimately end up with AI systems that worsen those health disparities, which is obviously the last thing we want. And there is definitely evidence out there to suggest that does happen if we're not careful."

Health Disparities podcast host Dr. Mary O'Connor spoke with Farrow about AI and health equity, and about his ongoing research on the use of AI to help general practitioners in the UK know when it is appropriate to refer patients to orthopedic surgeons for consideration of hip and knee replacement surgeries.

Never miss an episode – subscribe to The Health Disparities podcast from Movement Is Life on Apple Podcasts, YouTube, or wherever you get your podcasts.

The transcript from today's episode has been lightly edited for clarity.

Luke Farrow: This is definitely an area where I see AI potentially being a benefit. But I think, again, it comes back to that fact that you have to understand it, you know, within the kind of data that you're using. And I think that you have to encode and incorporate that into your models in order to allow that to happen. I think it is very easy in this situation, that if you're not aware of that, and you don't incorporate that into the way that these AI systems develop, that you can ultimately end up with AI systems that worsen those health disparities, which is obviously the last thing we want. And there is definitely evidence out there, you know, to suggest that does happen if we're not careful.

Mary O'Connor: You're listening to the Health Disparities podcast for Movement Is Life. I'm Dr. Mary O'Connor, Chair of the Board of Movement Is Life and co-founder and Chief Medical Officer at Vori Health. Artificial Intelligence is transforming healthcare. The promise of this technology is enormous, and is already being realized to increase the accuracy of diagnosis, promote patient engagement, increase efficiency in health care, and lower costs.

It's even being used to identify patients at risk of disease. And as we'll discuss more today, predict patients who might be good candidates for medical procedures. To learn more about one

AI approach with applications in orthopedic surgery, I'm joined by Mr. Luke Farrow, an orthopedic and trauma surgeon, and clinical researcher at the University of Aberdeen in Scotland, in the UK, where, by the way, surgeons are referred to as Mr. or Ms, or Mrs. and not doctor. Mr. Farrow, is leading a research project on the use of AI to help practitioners in the UK know when it's appropriate to refer patients to orthopedic surgeons for consideration of hip and knee replacement surgeries. Of course, this is a topic near and dear to us at Movement Is Life because of the disparities in utilization of these surgeries by patients of color and lower socioeconomic means. So we're very, very excited to have Luke with us today. Luke, welcome to the Health Disparities podcast. Thank you for being here.

Farrow: Thank you so much for having me on.

O'Connor: Okay, so let's start with the why. Why is there a need for such an AI tool?

Farrow: So I guess, one of the first things probably to explain is a little bit about how the kind of healthcare system kind of works in the UK. So in the UK, kind of anybody who has a sore hip or a sore knee, and might think that they have arthritis would typically be seen by what we term as a general practitioner, or a kind of PCP in the United States, who then basically see that patient and decide whether to send that person on to an orthopedic surgeon, you know, if that's appropriate or not. Now, the problem is that GPs, you know, they have this massive range of knowledge about many kinds of different kinds of conditions. But that means that they don't, you know, have the ability to have specialist knowledge about every condition, and therefore, ultimately, you know, sometimes struggle to understand whether a person might be ready or not to have a joint replacement.

This means that ultimately, we end up with quite a lot of people being referred for consideration of surgery, even if they're at a stage where maybe surgery isn't quite, you know, the right option at that time. And ultimately, this creates a big backlog of for people basically waiting to be seen by an orthopedic surgeon. And we know certainly those with severe arthritis that waiting for joint replacement surgery is not a benign thing, you know, people's health obviously deteriorates, typically, while they're waiting for surgery. And this understandably, has a bigger impact on people who are in a kind of poorer kind of health kind of state to begin with.

So our rationale for kind of developing this kind of tool was really that it could potentially help to really kind of decide who needed prioritized for orthopedic review, you know, without the kind of workload of having to go and sit down and trawl through all the kind of medical records, medical notes, and therefore hopefully make the system much more efficient, so the people that really need to be seen are the ones that are getting prioritized.

O'Connor: So that all makes perfect sense, especially if there's a long wait time, right? Because access to good health care for to be effective also has to be timely, right? So patients aren't waiting forever. Can you give our listeners like what's what's kind of the range of average wait times for a patient who sees their GP in the UK now to actually get in to see an orthopedic surgeon?

Farrow: Yeah, so I mean, unfortunately, in the UK, we've had some real kind of challenges with this recently, suddenly, the kind of COVID pandemic, which led to kind of, you know, a lot of elective kind of surgery, you know, hip knee replacement type surgeries, stopping for a period of time, within the UK, associated with kind of increasing demand has really seen our waiting lists kind of skyrocket. So certainly at the moment, within the UK, you could be waiting for probably up to a year from being referred from your kind of general practitioner PCP to see an orthopedic surgeon, and then maybe even another year after that, before you actually go ahead and get that surgery. So obviously, quite a significant amount of time.

O'Connor: Wow. Look, I can only say the people in the UK have a lot more patience than the patients in the United States, because I do not think that people would accept that here.

Farrow: Yeah. And I think, you know, it's really challenging. And obviously talking about kind of, you know, health disparities, certainly, obviously, one of the benefits in the UK has been having this kind of universal kind of health care system, where, you know, it's been kind of equitable kind of access to everybody. But now suddenly, what we're seeing is that, you know, we're seeing a bit more of a distinction between a kind of public kind of system and a kind of private system, whereby the people who are able to afford to pay to go and see somebody, you know, in a more timely fashion are able to do that. And unfortunately, it's the people from the lowest socioeconomic backgrounds, obviously, can't afford to do that. And the ones that end unfortunately, end up kind of waiting at the bottom of the queue.

O'Connor: Right, because you have the National Health Service, which is the universal health care, and then you have patients that can access your private system in the UK? Where they would pay out of pocket or have some kind of additional insurance coverage for the private system.

Farrow: Yeah, that's right. Yeah, absolutely. And certainly, you know, kind of historically, you know, the vast, I think, obviously, unlike in the US, the vast, vast, vast majority of our healthcare system functions through that kind of public kind of state. So, you know, probably more than 95%, you know, but again, certainly, we're seeing that proportion kind of start to change, you know, with these kinds of pressures that are on that kind of poke public kind of healthcare system at the moment.

O'Connor: Alright, so, we've got the why, right, you need the tool, because it will help prioritize the patients in terms of who has access to the surgeon in a more timely manner, given these crazy, you know, up to one year wait times to see a surgeon. So, walk us through, then your approach for developing the tool, as a machine learning module, to really try and, for the tool, to understand who is the right patient to see the surgeon in consultation next.

Farrow: Yeah, so what what we're trying to do is really trying to kind of replicate the information that a surgeon will have their kind of fingertips when they're kind of making those kind of decisions. And we've collected that data from within our kind of our kind of hospital systems. So

this will be all types of kind of patient kind of records. So you know, things like a person's kind of X -rays, there'll be kind of text information, you know, that's been stored in there, there'll be information about their kind of demographics kind of who they are. There'll be information about their background kind of health, what their kind of current symptoms are at the moment what medications they're on. And then basically, all this kind of information is is kind of stored within a very kind of safe, secure kind of research environment, which is specifically basically designed for this kind of AI kind of development. And they'd be we basically kind of build the models from there, so to start off with that information.

Typically, what we're doing is a process called fine tuning a model. So we're taking a model that's been often kind of developed kind of elsewhere and then basically adjusting it to fit a specific task. So, you know, if you talk, for example, about X-ray kind of interpretation, we might take a take a kind of machine learning model that has been initially kind of trained on just recognizing kind of everyday kind of objects. And you're then kind of adding to that and trying to see if we can teach it, you know, to look at the X-rays and to try and define some of those kind of key characteristics within the X-rays that a surgeon will use to make those decisions about whether a patient is suitable for surgery or not.

O'Connor: And so my understanding is that you've looked at about 2,000 patients that you selected to help you and your team developed this tool, where you're looking at the decisions that the surgeons made, and how the information that they use to make those decisions. So did I get that right?

Farrow: Yeah, that's correct. So, I mean, basically everyone was included, that were people within that kind of health kind of health care system that basically been through this kind of pathway. So had been referred kind of by their kind of general practitioner, through to the kind of orthopedic kind of service. And basically we have that information, so from about 2015, to 2022. And that was predominantly actually, just because that was the kind of start of when they kind of electronic healthcare records really started to kind of come in, within our region. So obviously, the, you know, the ability to kind of have and utilize that data is very much kind of dependent on that. So, you know, unfortunately, we aren't really able to go back kind of too much further than that, which which limits the data that we have a little bit.

O'Connor: And so how, aside from obviously, like X-ray findings, right, one could easily understand that the X-ray findings would have to show at least moderate to severe arthritis for for this tool to then, you know, say this could be an appropriate patient, right? What would some other factors be that would go into the tool?

Farrow: Yeah, so I mean, you know, predominantly some of those things that we talked about. So you know, and these are obviously very much important things to be considering when, you know, we're talking about health disparities, because, you know, I'm sure we'll come on to, you know, I think surgeons are not perfect in their kind of decision making or their ability, you know, to be completely kind of unbiased in that decision making. So, you know, it's things like a person's age, as I said, it might be things like a person's kind of background, medical health,

obviously, some people, you know, might not necessarily be fit enough to have an operation. So looking at that, obviously, a person symptoms, you know, generally, hip and knee replacement is, you know, primarily a pain relieving operation. So obviously, if somebody doesn't have significant enough symptoms, and again, there's probably not much point in going ahead and doing the operation at that stage. So as I talked about earlier, really just trying to replicate that information that a surgeon is using, to try and essentially kind of, you know, see if we can basically copy that kind of decision making process, but obviously, you know, without them without the surgeon necessarily being kind of directly involved with that.

O'Connor: So, you know, here in the States, we have a, we obviously have some health disparities. And certainly, when it comes to hip and knee replacement, we know that there's under utilization of these surgeries, typically by lower income, and I'll say, under insured patients. Now, that's not so much an issue in the UK, because the vast majority of the population is under the National Health System, right? I'm making an assumption I'm asking you, right, what kind of disparities in selection or utilization of these surgeries have you historically seen? Is there, you know, disparities still related to some kind of socioeconomic status, concern about weight and BMI? What are some of the factors that would be of concern right now?

Farrow: Yeah, so I think certainly those, you know, three that you mentioned, would were definitely very, you know, high up on the list. You know, certainly there's good evidence that, you know, within the UK kind of health system that people, you know, despite the fact that it's a public kind of health care system, you know, people with a lower kind of socioeconomic status do probably struggle still to access kind of health care in the same way, you know, that perhaps other people might do. And that's very much true, you know, within the orthopedic kind of surgery kind of space as well, we do see that, you know, I think, even work, particularly, you know, with the way that the kind of waiting times are at the moment, you know, you do tend to find that, you know, sometimes it can be the people that shout the loudest, you know, are the ones that end up kind of getting prioritized.

And obviously, again, that doesn't tend to be those people, you know, from from a poorer kind of socioeconomic kind of background, again, in terms of trying to... one of the ways that they've tried to obviously deal with our current kind of waiting time system is to kind of set up isolated surgical sites where there are a bit of a way from some of the other kind of healthcare pressures, but the problem is these sites that they often don't have kind of higher levels of kind of perioperative care. So, again, they aren't necessarily able to look after the patients that are, you know, are less fit or have more complex kind of comorbidities. So, again, you get that almost, that double effect, where you end up with the kind of fitter kind of population being siphoned off to kind of, you know, have procedures in those centers, but actually the ones with a poorer socioeconomic status, and actually, probably the people that really need it the most are the ones that end up kind of getting left behind with that kind of system. So, you know, we do, in exactly the same way have to be very, you still have to be very careful about how we, how we manage it and tackle those disparities.

O'Connor: Right, so how do you build a tool that doesn't incorporate the current bias, right? in the decision, either the decision making process or whoever yells the loudest to move up in the queue? Right?

Farrow: Yeah, so, you know, it is a difficult challenge, because ultimately, a lot of the time with these models, you know, the outputs that you get a very much related to the data, the information that you put in. But I think one of the key things is, is knowledge and understanding that and understanding the disparities that are present, and how we can potentially impact those. And it may be that you have to, and you can build these into these algorithms where, you know, it can adjust for those disparities.

So for example, it can adjust for the fact that, you know, somebody who has a lower socioeconomic kind of status, it can almost prioritize those people within that system. And suddenly, as part of our research we've done quite a lot of work really trying to look and understand both from a patient and surgeon point of view, what people really, what really matters to people in terms of, you know, who should be getting kind of priority? You know, as I said, should it be based on socioeconomic status? Should it be based on how bad their symptoms are, you know, the other factors involved, if somebody's, you know, a working kind of parent and has carer responsibilities or other things, you know, how did they impact on how you kind of developed these models? So I think that is certainly the first step is that that kind of understanding and appreciation of the fact that these things exist, and then really trying to kind of uncouple, that kind of unconscious kind of bias that we sometimes kind of see kind of seeping in.

One of the other key things I think to appreciate is that bias is different in different places. And again, with developing these models, you have to be very careful when you're taking them out of a system that they've been developed in and applying to a new system, because it's very easy if you're applying it to a different population, that the models actually don't end up doing what you're trying to achieve, because they've not been appropriately tested and trained within that specific population. So that's one of the things that we really look at in terms of, obviously, yes, we've developed this kind of model primarily locally, but if we're going to use that wider field within the UK, or you know, even bring it across to the states, you know, again, you would, that model might well look very different for whatever kind of area you are applying it to.

O'Connor: Those were great responses. I mean, you know, there's no tool that's going to be perfect. I think everybody understands that. But if the tool can reasonably mitigate against the unconscious bias that the surgeons currently have, right, to try to minimize that, and still, in effect, do a greater good, because it's allowing and supporting those patients that really need the priority access to get seen, then, actually, it has the potential to promote health equity. Is that how, ... do you think that's a fair statement?

Farrow: Yeah, absolutely. And I think that's, you know, that this is definitely an area where I see ai ai potentially being a benefit. But I think, again, it comes back to that fact that you have to understand it within the kind of data that you're using, and I think that you have to encode and

incorporate that into your models in order to allow that to happen, I think is very easy in this situation, that if you're not aware of that and you don't incorporate that into the way that these AI systems develop, that you can ultimately ended up with, you know, AI systems that worsen those health disparities, which is obviously the last thing we want. And there is definitely evidence out there, you know, to suggest that that does happen if we're not careful.

O'Connor: Right. So Luke, what's the status of the AI tool now? How close are you to, you know, having it completed? And how do you test it before kind of more widespread utilization to see how well it's performing, and if it's helping to promote health equity and not, you know, make things worse?

Farrow: Yeah, absolutely. So we're currently kind of still in the kind of research phase at the moment. So we've started kind of developing and building the models, and as I said, primarily basing that off the kind of X-ray, because we know that that's probably one of the kind of key things as you mentioned, if somebody doesn't have kind of arthritis severe enough on their X-ray, you know, then the likelihood is that they're not getting a joint replacement, and then building in all these kinds of other things kind of on top of on top of that, as well. I think, again, one of the kind of key things with developing these AI systems is, again, it has to be done in a very kind of structured and kind of evidence-based kind of way, just in the same way that if you were developing a new drug, or if you were developing a you know, and a new orthopedic implant that there has to be a very rigorous kind of testing kind of system that we need to go through in order to basically make that happen. So, usually starting off with essentially what we're doing now, by using a kind of historical kind of information, you're then taking that information out kind of forward. So testing the model in a bit more of a, a kind of day to day kind of purpose.

But ultimately, from my perspective, I think to ensure that these models are safe and effective, I think just in the same way that we have randomized control trials for the drugs isn't said for other things, I think we do need to see randomized control trials have similar kind of AI tools, in order to demonstrate that they are effective, and they are doing what we achieve. And I think that's the only way to really test them. And then finally, I think we need to monitor them appropriately as well, because the other thing that can happen with these models is that over time, we get what's called drift. So the model can drift or the data that's coming in drift, say for example, the threshold that people have for operating on people, you know, with hip or knee replacement changes over time, then obviously, that potentially kind of impacts on your model. So it's very aware that we have that kind of longitudinal kind of process as well, to continually monitor and make sure, again, that our models are doing what we're trying trying to achieve, and not ending up, you know, achieving something different.

O'Connor: How do you think patients will react to the tool being used? And I just had this thought popped into my head, if they don't like the outcome of the tool, would you envision some kind of, you know, review process or appeal process? Like, I went to see my GP, and my knee is killing me, but you're, but the tool says I'm not bad enough to go see the orthopedic surgeon. And I say that's malarkey. And I want I want someone to review this.

Farrow: Yeah, no, absolutely. I think again, it's a really important part of the process, particularly, you know, when we're dealing with AI in a situation like healthcare, where it's a really kind of high stakes environment, you know, it's people's lives that we're dealing with, isn't it? So I think there's a number of different ways that you can essentially approach that kind of challenge. So, you can obviously have, you know, somebody's checking through the results. And certainly, if you look at, for example, in radiology, that's a lot of one of the things that they'll do. So, they'll have perhaps an AI system read a chest X ray, but at the same time, they'll have a radiologist take a look as well, just to say, I agree with this, or, you know, I don't agree with this.

But I think having some sort of safety netting kind of process is really important because just in this, you know, these models again, in the same way that humans aren't perfect, the models are never going to be perfect, and they are going to make mistakes. So it's so obviously, it's about mitigating those mistakes as much as possible. And making sure that you know, when those mistakes do happen, that it's not just a kind of 'goodbye, see you later, you know, we're never gonna see you again' kind of thing that, that there is a kind of process in place that people can still kind of feed back into into that kind of loop.

O'Connor: Yeah, that's great. That's great. When do you, in your crystal ball, if everything is going well, when do you think you'll be testing your model or kind of bringing it online?

Farrow: Yeah. So I mean, we're certainly hoping within the next couple of years, that we'll have certainly so you know, we have essentially a kind of prototype at the moment that, as I said, we're just kind of testing within this kind of secure environment. The next step is to then bring it outside of that environment and kind of test it elsewhere. There's obviously, you know, with all these kinds of things are quite stringent kind of processes to go through in terms of, you know, most of this kind of AI kind of software is classified as a kind of medical device. So, again, it has to go through quite kind of, you know, stringent kind of testing from that side of things, which, again, is a good thing, but, you know, it takes time for that kind of process to process to happen. So that's certainly the hope within the next couple of years that we'd have something that could be up and running and usable.

O'Connor: Well, I think your work is very exciting. And I'm really looking forward to seeing how the tool works at the end of the day. And, and appreciate and truly appreciate your, your commitment to working to make sure the tool is as unbiased as possible. Because the point would be to for it to be less biased than us humans, right, because that would still help advance health equity.

Farrow: Yes, absolutely.

O'Connor: Yeah. Any other closing thoughts that you have for our viewers and listeners?

Farrow: Thanks. I mean, I think the only thing I think may be left say is that I think promote, you know, I think I, you know, does have a lot of potential to help, you know, with the health disparities, I think, not only from what we talked about in terms of using it from a, you know, as a clinician or a surgeon or, you know, a doctor, but also, I think from the kind of patient kind of facing side as well. So, you know, I think helping, obviously, people to understand, complex medical, their own complex medical notes, the ability to summarize that information and put it in a way that, you know, people can understand the ability to have, you know, kind of individualized kind of treatment plans or chat bots that can help people through kind of healthcare problems that again, the ways there that can hopefully it kind of helped to open up kind of access to healthcare in a way that we haven't necessarily seen before. Yes, obviously, there's, you know, some technology barriers that kind of come in with that as well that we need to be appreciative of, but I think there are definitely a lot of potential avenues with which I can really kind of help in this space.

O'Connor: Well, thank you, Luke. So to our viewers and listeners, you can learn more about Mr. Luke Farrow's research at the link in our show notes. And I'd like to thank our guest, Mr. Luke Farrow, for joining us today. This brings us to the end of another episode of the Health Disparities podcast from Movement Is Life. If you like what you hear, be sure to subscribe wherever you get your podcasts, and take a moment to leave us a rating, it really makes a difference.

Also mark your calendar for our upcoming Annual Summit. This year we will be in Atlanta, Georgia, an Thursday, November 14th and Friday, November 15 at the Whitley hotel in the Buckhead area of Atlanta. You can go to our website and register for more information. And of course, Luke, you are warmly invited to join us in Atlanta, Georgia, if you're on this side of the pond at that time. I'm Dr. Mary O'Connor. Until next time, be safe and be well.