It's how you interface with the patient, how you monitor and you also enthuse the patient or encourage the patient to comply and then having a backend process that provides decision outputs for the physicians that are valuable and are actionable, and it's the three-way system we think is critical to the future of this particular product of telemedicine in the ophthalmic space.

DR. BLUMENKRANZ: Thank you, Quinton. Pravin Dugel was supposed to be here and unfortunately, due to a family illness, he wasn't but he sent me a few slides. I'll just -- if you could turn to slide 22? I think the idea of processing all this information and having it be actionable is an interesting one. And it turns out that aside from using AI and DL and so forth, it's possible to use different ways of looking at data. For instance, in the office, we're used to looking at individual hand-written reports or typed reports or tabular data.

This is just actually a page from Epic here. And you can -- it's hard to read but that's --

they typically are hard to read even if you're not in the back of the room here. But you can see that data there. This is the same patient. This is tabular data. Can anyone pick out a pattern there? Is anyone quick enough to figure out what's going on? Maybe two, three, four minutes you'd be able to. How about that pattern there? Those are individual data points taken of a patient at home. Anybody starting to see anything going on?

What if you connect the dots, does it get more interesting? And what if you used a smoothing algorithm to interpolate between the points? Well, that's all home data. That's a real patient and Drug A is a drug that was given for treatment of exudative AMD, patient seemed to be doing poorly, switched to Drug B. I'm purposely hiding the names of the manufacturers so as to not be unfairly accused of favoring one over the other. And this -- that's the data, the patient was switched to Drug B and you can see immediately upon doing so, the visual acuity went up.

And this is the office data. This is the two points, they connect the dots and the smoothing and then finally, in the "light blue," you can see the actual office data. So it points out that there are lots of different ways, first of all, to acquire data, and there are also lots of different ways to look at data. And we may be still living in an era where we're used to looking at numbers, but I think if you think about the whole field of infographics and how to analyze large datasets, our minds are really based on pattern recognition and, I guess, Gestalt or however -- whatever the nontechnical terms are. And I think there is an opportunity for all of us to be able to use different ways of looking at the same data and acquiring more data but more of the same. Even simple numbers like vision and being able to acquire information. I'll stop at that point.

MR. OSWALD: Question, Mark.

DR. BLUMENKRANZ: Yeah.

MR. OSWALD What is the interval between the tests?

DR. BLUMENKRANZ: The interval, those are taken, on average, between three and five times a week.

MR. OSWALD: Okay.

DR. BLUMENKRANZ: And it's just a -- it's a visual acuity taken on a Smartphone.

MR. OSWALD: Okay.

DR. BLUMENKRANZ: Yeah.

DR. NISCHAL: Okay. So we're going to stop just for a few minutes for questions from the floor. Are there any questions for any of the panel? If you can just say who you are for the --

FEMALE SPEAKER: (Inaudible) from Columbia University. I have a quick question. We're generating all this data, offices are generating the data, hospitals are generating the the data, these data are required by imaging companies and AI companies to build these algorithms. Who owns the data? This is one. Second, in an era in the future, retinal images, iris images are going to become protected health information. What is going to happen then so?

DR. BLUMENKRANZ: Who wants to take that?

DR. NISCHAL: Not me.

DR. BLUMENKRANZ: Michael, do you want to --

DR. ABRAMOFF: I have strong opinions.

MALE SPEAKER: Yes, go ahead.

DR. BLUMENKRANZ: Why don't we have the two end Michaels talk about that. Between you, you have (inaudible) --

DR. ABRAMOFF: -- but I took (inaudible) --

DR. BLUMENKRANZ: -- perspectives.

DR. ABRAMOFF: -- okay, well, we'll see. So I think the patient owns the data or should own the data. I mean I would want to own my data. I do not want it to be sold by some hospital where I don't even see what they got for it. So -- but then I am a proponent of using data that is acquired for training algorithms, right, if we're testing algorithms, that is acquired fairly, you know, in a controlled fashion like for clinical trials and not just, you know, buy it from some hospital where patients don't even know that their data's being bought. So I would say I have a very strong opinion and competitors will think very differently, so I will leave it at that. So now you talk for the competitor.

DR. CHIANG: No, no. I -- it's (inaudible) that I am -- you know, we can talk about opinions and we can talk about sort of -- you know, sort of legalities of it and I think that, you know -- Lemma, I think my answer to that is that it's contextual. And, you know, if we see a patient sort of, you know, we own the data, from the hospital perspective, you know, the patient owns the data because it's their data, you know.

And if it's home-generated data, I'm not sure we have a clear precedent for, you know, for what happens with that. You know, presumably, that home data may be uploaded to the electronic health record in which case it may fall under, you know, the auspices of both of those. I know that's something that we've got to, you know, we've got to work out.

And, you know, as an aside to that, you know, I think as a medical community, I've personally seen a few situations where there's a little bit of ambiguity in that where patients get access to their own medical record and will say things like, you know, what are you talking about, I'm not a drug addict or I'm not an alcoholic; you know, can you change that from my medical record. And so I think there are things with oversight and patients, you know, sharing to this that I think are questioning some of the assumptions that we've had all along in terms of medicine. So I think it's an important question.

DR. NISCHAL: Can I just say I think that owning data and exposure to data are two different things, and I'll give you an example of what happened. And so the adolescent diabetics at our children's hospital were given monitors to monitor their blood pressure, and some of them were put on a beta blocker and some weren't. And the ones were put on a beta blocker, their traces at home actually were higher than the patients who were not a beta blocker. And it turned out that these children had access to what their blood pressure was. They could see it and that some of them were getting anxious about the blood pressure, and as they became anxious, the blood pressure went up. So it's really important that while the patient owns the data, it's not necessarily best for the patient's health to be exposed to that data, which comes back to the question of who analyzes that data.

DR. BLUMENKRANZ: So the Heisenberg uncertainty principle?

MALE SPEAKER: Right.

DR. BLUMENKRANZ: Okay, please.

DR. ORR: Hi. Susan Orr with Notal Vision and I have a comment about the amount of data as well. Going back to the physician, there was a slide at the beginning saying, I think, 100-plus apps have been approved by the FDA, which is an unprecedented amount of data that's inundating the physician who's trying to treat that patient. And in our experience, which Quinton has spoken to, the doctors are very limited in how much time they can spend looking at this data.

So I'm interested in a comment on the level of robustness and validation of the benefit of these apps in order to drive adoption across the physicians. Now just the example with home OCTs, since we've spent a lot of time interrogating it, doctors are not going to look at every scan on every OCT for every patient. So in order to extend the visits or have better outcomes, at some point, there has to be a reliance on that. And many of the apps don't have that level one evidence to support modifying the practice of medicine for a given indication.

DR. BLUMENKRANZ: Anybody? I can comment. I think you're absolutely correct. I think everything that's used in clinical practice needs to be very rigorously validated and I think efforts are now under way. And I think that's really part of this -- the whole idea behind this workshop is to both expose people to the potential benefits of this and also the pitfalls and the need for rigor and validation of anything that's going to be used. So I I certainly completely agree.

I think in speaking to the issue of data overload, that was kind of what I was alluding to before. I think that's really where automation-augmented intelligence and deep learning can really play a role. I think if it was left -- if we generate -- you know, if you look at the number of terabytes of data that are being produced every second in the world today and who's going to look at that, who's going to do something based on that, it's -- it would be impossible without using some sort of, you know, very augmented kind of computing power. I think -- and I think that's where it all fits together. That's worthy -- in my view at least, that's where the AI solves the problem of the data load and also the learning and making actual -- making real use of that data, not having it be just a botherance and then finally validation of that.

I'll just make one point because I've been -- everybody's been -- I think it was Paul Lee initially that talked about the issue of what's the gold standard. Fifteen years ago we published in the American Journal of Ophthalmology a study in which we were looking at whether or not a single mydriatic non -- a nonmydriatic monochromatic fundus image was as good as seven standard fields.

And we also got physicians at the Kaiser health system -- or it doesn't matter which one -- who were were -- who practiced in the art of ophthalmoscopy and diabetic retinopathy detection to grade those same patients at a separate sitting. And the first interesting part was that the digital nonmydriatic monochromatic images on -- in general were about 87 percent as sensitive as 7 standard fields. And we happened to be using that as the gold standard.

We then checked the ophthalmoscopy results and it was a 34 percent concurrence of the data. And so the interesting problem was that we had shown that digital was pretty good but that ophthalmoscopy, which was the gold standard in previous years, was no longer as good as either the new innovation or even the one that was existing.

And so it raises real questions as to what is -- you know, what are gold standards. I was interviewed by Ken Mills, who was the President of the American Academy in commentary on that, and he was not only bright but but wise and he said the problem with all of this is that when you introduce these new technologies in that case, those images were read not by physicians -- and we didn't have AI at that time -- they were read by graders at the Wisconsin Reading Center, so we know they were very good. And in fact, the nonphysicians graded retinopathy better than ophthalmologists.

Now it's were they better at really seeing it? No. I mean they had as many hours -- minutes or hours as they wanted to stare at a high resolution image on a screen whereas an ophthalmologist is seeing perhaps 30 patients in a half-day, the pupils not optimally dilated, no one's giving them the very best photo.

So in the real world, you know, situations are quite different than they are in clinical studies, and I think it's an important point that you raise, is how do you how do you get to the best data; what is the best data? I don't know that home data might not be better, worse or the same than clinical data obtained in the office, but that's what we have to do and that's the critical role that the FDA plays working hand-in-hand and collaborating with the people that are trying to develop this technology so everybody buys into whatever those results are. That's what they are and we know whether something's better, worse or the same than what we're currently doing. At least that's just a personal opinion

DR. NISCHAL: I'm going to have to move us along I'm afraid because we still have some really important questions to answer.

So, hopefully -- I'm sorry, Michael, we'll come back to you.

So we're going to move on to artificial intelligence which we've been discussing, and one of the first questions that we wanted to tackle was, how will I affect the use of our family digital tools in the future, which we've covered to an extent. And, Michael One, I wonder if you could -- slide 15, please.

DR. GOLDBAUM: If we can get back. Is -- which is One.

DR. NISCHAL: That's you.

DR. GOLDBAUM: Okay.

MALE SPEAKER: We figured it out.

DR. GOLDBAUM: I wanted -- I just wanted to make sure that -- so somewhere past 46, there's a slide that says "AI in medicine." But -- so the thing -- there are a number -- this has been studied. AI can break down into a number of different groups and it's something like 12 different subtopics. But the three that most interest us would be natural language, management of uncertainty, machine learning data mine and data mining, and image processing.

And the natural language, I guess, best would be for translation though it's also been used to to -- for other questions in medicine.

Management of uncertainty; in the past, we were doing things like expert systems, and that was labor intensive, and so it never got adopted.

With the deep learning, it learns from the data. You don't have to guide it and it does everything, and that really helps for us to be able to build these systems. And so we're using it for image processing and I think we'll continue to use it for image processing. We'll use it for image classification or interpretation and also for the component parts like image segmentation to find the various structures of importance in an image. And I think that will be -- it will be basically physician assistance in the beginning. Maybe eventually, we'll be able to learn from these systems but I think initially, it will be physician assistance in managing large amounts of data and learning, helping us to learn from the data.

DR. ABRAMOFF: Me? I need to see the slides (inaudible).

DR. NISCHAL: So thank you, Michael. We'll go on to the next part. Are there specific AI examples that help us negotiate these issues? Now, for example, interpretation of fundus photos for retinal disease screening and Michael Abramoff is going to tackle that question for us. While we're waiting for Michael --

DR. ABRAMOFF: I will just stand here and control my slides.

DR. NISCHAL: Okay. All right.

DR. ABRAMOFF: So two things --

MALE SPEAKER: Take the microphone --

DR. ABRAMOFF: This is good, this is good. So, Michael Abramoff. Shameless plug. I am briefing Congress, both the Senate and the House, on AI in medicine on Wednesday and I will be speaking about this meeting and that we had it and that FDA was involved, just so you know that we'll be speaking with Congress about this.

So now to AI. So, you know, Mike G., you're Mike One now, you know, did an excellent introduction. And so I just wanted to talk about algorithms for image analysis and specifically for retinal images where on the top you see sort of, you know, the way we do it where it's lesion-based so you have an image. You look at the image quality which is a big issue. We got many of the images coming from especially not so well-trained photographers will be insufficient and you need to know that so in real time you can tell them, hey, take it again. So that's an important aspect.

And then what we do, our algorithms do, is have specific deep learning modules that detect micro aneurysms and exudates or an abnormal disk, etcetera. And then that combines with anatomy, where the disc is, where the fovea is, etcetera, and that determines the outputs of the system.

And then Mike One and me, so Michael G. -- sorry, I get us confused all the time -- so we probably disagree about goal or role of black boxes, which is the bottom line, where essentially you have an image and you actually share it with an output and you don't really know what's going on. So instead of having an explicit exudate and explicit micro aneurysm, you say, well, I want you to associate this type of image with diabetic retinopathy or with glaucoma and this without.

So, you know, next slide. One thing I worry about is this, which is -- we showed this at (Inaudible) and hopefully, that publication will be accepted once. So --

(Whereupon, off comments/adjusting lighting.)

DR. ABRAMOFF: Oh, yeah, it's hard. You won't see it. So on the left is an image with diabetic retinopathy. Just believe from me that it's very obvious full of exudates, and Mark can probably confirm that it's DR, right, on the left. Yeah, there you go. And so there's exudates and hemorrhages and --

DR. BLUMENKRANZ: Probably.

DR. ABRAMOFF: Pardon me?

DR. BLUMENKRANZ: Probably. No, I'm just kidding.

DR. ABRAMOFF: Not probably.

(Laughter.)

DR. ABRAMOFF: Okay. It's the most obvious.

And so if you change a few pixels on the right, it's only minimally changed. And, you know. it still looks to me and you and Mark, hopefully, like DR. And then if you have algorithms that 99 so these are minimal changes and you have algorithms that are very sensitive to this and you don't know that, like black boxes, and we test this on a number of a black box, you know, CNNs, meaning convolution of neural networks, and they all started to see this image as normal. And so -- and experts would never do that.

So there's a sort of risk that it trains on things you don't really know about, and so I worry about black boxes in general. So I just wanted to bring that up because it's an interesting debate, and I'm sure you have something to say against it.

DR. GOLDBAUM: No. Actually I don't have anything to say against it, but what you can do is put your adversarial images in there, too, and label them correctly, and then it will learn how to get less than optimal images.

DR. ABRAMOFF: Yeah. So then -- but you don't know what the perturbation will be, right? So it can be compression or some noise and so you would have to train for all these different relatives.

DR. GOLDBAUM: You made this up but photographers will have various qualities of images. That's the real world. This one is not. So you can use -- you can train on the adversarial images created in the real world and the system will learn how to look beyond those adversarial elements

MALE SPEAKER: I think we're learning about adversarial communications here.

(Laughter.)

DR. NISCHAL: Can I --

MALE SPEAKER: -- (Inaudible).

DR. NISCHAL: -- can I just ask -- so, you know, with the question of poor image quality, I mean does anybody on the panel, anybody in the audience have experience with fractal analysis, because this seems to be one way of picking up retinal diseases looking at the the actual branching of the vessels? Does anybody have any experience, either on the panel or in the audience, of fractal analysis for analysis of these images?

MALE SPEAKER: Yeah, I do but --

MALE SPEAKER: Yeah.

MALE SPEAKER: -- yeah. You want to say something?

DR. CHIANG: We've done it --

MALE SPEAKER: Yeah, we've done it.

DR. CHIANG: -- and it works but it doesn't -- we have -- it doesn't work as well as the other things that we've done.

DR. ABRAMOFF: It doesn't add to the performance for DR or glaucoma --

DR. NISCHAL: Okay. We're going to keep moving. Could we go to slide 16, please? This is more about AI. This AI-enabled image analysis questions. So this is for you, Linda, because you've been very quiet and polite. And so are we ready for a fully automated interpretation?

DR. ZANGWILL: I think in some --

DR. BLUMENKRANZ: Slide 16, please, 1-6?

DR. ZANGWILL: -- I think there's good evidence in some cases. I think diabetic retinopathy is the closest to that. And as a non-clinician, I tend to defer to clinicians on this, but I really do think that the algorithms are close enough, and the -- it's compelling enough for diabetic retinopathy when the lack of access ophthalmic care, etcetera.

And I just want to also say that in terms of fully automated interpretation, I would take AI. We're talking about poor quality training at home. I think another avenue for AI would be to help train the people at home develop algorithms and training schemes to identify poor quality images or identify poor quality visual fields and bring that back to the patient to -- and improve the quality of those questionable data points.

DR. BLUMENKRANZ: Okay. Thank you. Michael A., this is for you. Does the AI DR algorithm give the patient or a doctor a diagnosis or a plan? We're going to go sequentially here just so you can see. Or does the patient's MD make the reading? Or does a third party doctor get to do it? So we're giving you the first crack at that. Does the -- does it go right to the patient or the doctor from the AI?

DR. ABRAMOFF: I think I've already spoken to them about this one.

DR. BLUMENKRANZ: Sort of covered that.

DR. ABRAMOFF: Yeah. So I think, you know, alignment with preferred practice patters really helps. I think again, it totally depends on the context. We're talking about normal eyecare professionals, primary care, it really needs to be (inaudible) and patients, you know, probably the same. I don't have experience with home monitoring but -- so, yeah, it should probably more -- be more of a diagnosis and a plan than, you know, probability of developing, you know, PDR two years from now. That is not something they can work with. We've thought about that.

DR. BLUMENKRANZ: And Michael C. -- or does the patient's MD make the reading enabled by IA? Is something that the personal physician should use and this is a tool available to them? Or does it go through a neutral vendor, if you will, or alternative source?

And then Michael G., we'll go to you next.

DR. CHIANG: Yeah. Mark, I think this is an opinion issue as much -- I mean is more so than a fact issue, and I guess I would say that my opinion is that machines are very good at -- machines can be very good at making diagnoses or by analyzing data. But I personally believe that doctors make plans; in other words, doctors make diagnoses and doctors make management plans.

And I guess what I mean by that is that I personally hope that we, as a society, will use these machines as decision aids the same way that I'll use my ophthalmoscope as a decision aide, or a cardiologist will use a stethoscope as a decision aid or an echocardiogram as a decision aid. In other words, they're all pieces of information that we use to piece together and make that diagnosis. And so I would think of these AI systems, you know, in the same way that it's another piece of information that I use that contributes to my overall clinical judgment and management of the patient.

And, you know. I think that one of the reasons, just just for the record, is that I think that as doctors, we do two things; one of them is that we diagnose and the second, we manage. In other words, you've got this diagnosis. What do you do now and how do you weigh the risk-benefit tradeoffs of one alternative versus another alternative.

And I think machines are, you know -- can be very good at diagnosing but I don't think they're very good at understanding patient preferences or understanding the context that we're going to apply those things in. And I think that all of that, you know, we have to consider in terms of developing and applying these systems and, you know, basically how to use them for patient care.

DR. GOLDBAUM: Okay. So if we can go to the slides 46 beyond --

MR. OSWALD: So just one comment --

DR. GOLDBAUM: -- who does the interpretation?

MR. OSWALD: Sorry, Michael.

MALE SPEAKER: Okay, go ahead.

MR. OSWALD: Yeah. Just one point

It's interesting in the last two months, we've had three different inquiries at Notal Vision about AI. We've had one from China. We read one from the UK and we've had one from a health system in the U.S.

DR. GOLDBAUM: Who does the interpretations?

MR. OSWALD: And I think the answer --

DR. GOLDBAUM: Who?

MR. OSWALD: -- to the question depends on what problem you're trying to solve for and I think will change by virtue of what you have available to you and what degree of trained personnel you have to deal with the issue. So rather than taking a U.S.-only context, I think there's a global context to this discussion.

DR. NISCHAL: Michael.

DR. GOLDBAUM: So there's one called, "who does the interpretation," but I'll read -- it's a single slide after that.

So if the machine does the interpretation, it's available 24/7; it's consistent; it doesn't get tired. It's a black box mostly; maybe we'll learn in the future how to get information out of it. And it should assist the physician at this point. And deep learning has allowed us to do a lot more with classifiers than in the past. The patient's regular doctor reads it. If the patient's regular doctor reads it, the data or the interpretation, that doctor has an interface between the -- that -- there's an interface between the physician and the patient and that's where the doctor still fits in.

That person is not available 24/7 and that person can be inconsistent, can be sleepy, can be wired, could be all sorts of things affecting him.

A third party doctor reads the results; no interface to the patient's radiologist, for example; no interface to the patient, but that person has the domain expertise that the regular doctor may not have; also not available 24/7 and also may be inconsistent. So those are the variables that fit with each of the three types of interpretation.

DR. ABRAMOFF: I want to go back to what Michael Chiang just said, which is I think it depends on the level. So, you know, we have been developing guidelines for autonomous devices for diabetic retinopathy with the American Telemedicine Association. So we go back and forth a lot with a group of authors, and one is this level. So that's a -- for the primary care physician, if you have a DR screening automated device, that's an assistive device; right? I mean --

MALE SPEAKER: Yeah.

DR. ABRAMOFF: -- they just hear, hey this patient is likely to have DR, manage this patient so maybe, you know, regulate better and also maybe refer. But it totally depends on the primary care physicians, so it's assistive. However, me, as a retinal specialist, I'm not having any influence of the results. So, for me now, as a retinal specialist, it's automated, so it's -- you know, it's terminology or semantics almost. So you have to be careful I think.

DR. NISCHAL: So we're going to move on to the last slide. Slide 18, please; 1-8. And I'd really like to give the whole panel an opportunity just to give a short answer to these two questions.

Firstly, how do -- and their safety of privacy concerns, you know -- h how do we address these concerns regarding the storage of information on personal devices in the era of common cloud backup for other data on personal phones and for technicians and patients? And how does monitoring of patient behavior and location relate to safety and efficacy concerns?

So if we'd like -- we're going to start with you, Michael Chiang, and then just work around and see what everybody has to say, and then we'll open the questions up to the floor.

DR. CHIANG: Yeah. Ken, I'm thinking about -- I think there was a really good discussion this morning about that that Natalie Afshari and Mike Trese did. And the one thing that I thought was -- that I would add to that discussion is that, you know, a couple of months ago, my 15-year-old daughter played in her first soccer game of the season. She came back cursing, you know, because she played 18 out of 80 minutes, and she felt undervalued as a player by the coach. And so, you know, I said, Erica (ph), you've just got to control what you can control, which is your attitude and your effort.

And I see an analogy with this, that we're sort of cursing about the hackers from China and India yet what we can control is the single most common security breach that, you know, I think is out there which is passwords that are either shared among people, or posted up on sticky notes, and -- or, you know, people use the same password for every system.

And, you know, I actually think that that's something that, you know, that, you know, to some extent that's sort of our low-hanging fruit in terms of these personal devices, sort of, you know people -- I think there are HIPAA rules are actually pretty good, you know, for protecting information. But the problem is that we don't apply them consistently and we're not very good -- so I think that if we could pay more attention to that, we'd go a long way toward solving, you know, this problem.

DR. NISCHAL: Michael A.

DR. ABRAMOFF: Well, I would just say that -- okay, the reason AI and deep learning is so popular right now is because of the enormous gains in computer power, and those are most achievable in the cloud or at least remote service because it's just more cost effective that way. And so there's a sort of push to do that because it saves you a lot of hardware and GPU costs that can be enormous. And at the same time, you know, because of doing that, you have traffic that otherwise you wouldn't have, because it would be processed locally. So it's sort adding a risk for a benefit, you know, making this AI technology even possible.

So there's a sort of -- you know, you need to find a balance there between security and even being able to do it. But it's -- you know, we're trying to solve it, all of us, but not fully solved.

DR. GOLDBAUM: Okay. So just move on to where it says "cloud." There. So I'm going to leave the cloud for now and just talk about security. So I think it's three slides beyond that. So if you just move three slides. Yeah.

So first of all, there is the -- in Europe, there's the European Union General Data Protection Regulation which is addressing a lot of these issues of patient data security. And I haven't found something comparable in the U.S. and there may be something comparable. If there isn't, it would be good for us to be looking at the same thing. And you can read about it at the website, eugdpr.org.

So there are ways to control -- ways of security. One is access; only authorized users;

you can have a password but even better would be a two-factor system where you put in this password and then it has to make a contact with your Smartphone and an app on the Smartphone says, "Yes, it's okay." So that's one of the methods that our institute is using right now.

Now transmission, there are various hypertext transfer protocol and various transfer methods that are more secure. And there's VPN which is just you and the direct communication to where wherever you are trying to communicate to.

The one thing that has not been addressed, and I don't know the answer to this one yet because -- it's the person going rogue. So the person who has access to the data and then decides that they're going to make it available to the entire world because of some feeling that they have. And so if anybody has an answer to that one, I'd like to hear it.

DR. NISCHAL: Linda.

DR. ZANGWILL: Yeah. I want to touch on, I think, the conversation this morning and my panelists. The human factor is really critical. I just want to remind everyone that Eric Snowden did not hack into the system. He took the data just like the person going rogue on the USB port, and that's something that's really challenging. And it could be -- that was obviously intentional -- it could be inadvertent that somebody wants to do more work at home and take something home, and then their laptop is, you know, lost, etcetera, etcetera. So I think the human factor in all these different systems and taking patients, monitoring -- home monitoring, etcetera is really going to be the challenge and the make or break of these systems really going forward.

MR. PATEL: So I'll just touch upon on a couple of points. I think one is, in my mind, is about trust. I think when we -- when FDA put out the guidance on cybersecurity, I think the fundamental principle in the guidance was about, you know, can the data be trusted and the person be trusted. So it's authorization and authentication about the data and the person accessing that data. So if you keep those principles, I think that concept needs to be sort of expanded in training and education as well as in use, and it can be one time and be done. I think it has to be -- or at a periodic basis to be reminded to people. So once we sort of think about those aspects, we get to a different spot and to maybe even address or identify or catch things that may be slipping away from us.

So I think thinking about not just cloud but just having data, sort of where the data resides; what it means; who do you trust it with and where do you get that information back, and who is accessing it is something that needs to be sort of -- that as an education level, should be up there and also awareness. So that's how I would think about it.

DR. NISCHAL: Thank you. So we're just going to wrap up with Zach and then Quinton.

DR. BODNAR: Sure. I think we had a good discussion in the earlier panel about the fact that it's very hard to secure these things from a technological point of view. If there are malicious actors out there, they're going to find a way to get in.

But this panel brought up something that I wanted to just continue on, which is that there is a human factors aspect to this as well and a lot of it is just not adherence to protocols. So to go back to like a classic example, the Enigma machine would have been an uncrackable device if the -- if they hadn't -- if they had used it correctly and that -- you know that's true to this day as well.

One of the ways that we could potentially mitigate this, but it's a little bit at odds with the principle of using that information to get as much from it as you can is, to compartmentalize it somewhat. So in the current practice of medicine, when you log into an EMR, you have access to every patient and everything about that patient. Should it really be that way? And when we go to a telemedicine-type environment where not everybody who has access to the system is even a physician, then I don't think that that's necessarily the right way to go. I think that you should have access to the information that's pertinent to you and what you need to do to use your job, and we have to do a better part -- a better job of compartmentalizing it that way

DR. NISCHAL: Quinton, last but not least.

MR. OSWALD: Two quick points. First of all, I think it's important, as a company, for us to maintain an external evaluation of our systems and processes. We do that with HIPAA on an annual basis.

The comment that the gentleman from the DoD made -- called me to ask a question to my CEO -- we use Amazon Cloud but we're not at the level that he indicated. And, you know, the question is what do we need to do to get there. So I think it's thinking about these elements is going to be important as we go forward.

DR. BLUMENKRANZ: Well, that brings us to a close. You -- some of you may have questions and

I would encourage you during the break that follows to sort of seek out the panel members.

I do want to thank all of you, all the Michaels and Linda and Bakul. And Quinton, you're thinking of changing your name, I know, to Michael and Zach. And so we will see -- we'll see you at, I guess --

DR. NISCHAL: 2:45.

DR. BLUMENKRANZ: -- 2:45. Thank you very much, everyone.

(Applause.)

(Whereupon, off the record at 2:27 p.m., and back on the record at 2:48 p.m.)

DR. HUMAYUN: (Off mic) a seat, would appreciate it. Thank you. So we'll get started with Panel 3 now and our panelists are listed up here. Lama will be going first followed by John, Nitin, David, and Eitan. If we go the next slide?

So the -- panel three was tasked to look at the effective safeguards and methods for mitigating the risks for an update on a digital health device and the assets threats and vulnerability to be considered and identified.

Mark Humayun, the the moderator and my co-moderator, Derek Sprunger.

If we go to the next slide? So we'll be addressing these items. What are the most effective methods of mitigating risk for ophthalmic digital health devices, safeguards built in the software and in the hardware, and methods to limit the intended users labeling for patient use training modules and tutorials?

The way we've organized this panel is we're going to have each panelist present a talk and try to address these questions during their talk, and then we'll open it up to the group.

Next question that we're going to answer is what are the assets, threats, and vulnerabilities that should be considered and identified as threat to the privacy of the patient for ophthalmic digital health device developers? Again, this is a topic that has been discussed previously, but I would like to ask the panelists to please focus in particular on how their device or how their technology has addressed some of these issues of transmission of information, storage of information, and monitoring patient behavior and location.

So with that, Derek, would you like to make a few comments.

DR. SPRUNGER: No. Just we're ready to go.

DR. HUMAYUN: Okay. So we're ready and we'll have Lama go first. So if you can go ahead and please make your presentation?

If you have any questions after the talk, please feel free to ask it at that time but again, we'll have a lot of discussion time to follow. So we have Lama's slide first?

DR. AL-ASWAD: So my name is Lama Al-Aswad. I'm the Director of the Tele-ophthalmology Initiative at Columbia University, and I started this effort because we launched a tele-ophthalmology project for identifying early disease in the community for diabetes, diabetic retinopathy, macular degeneration, glaucoma, and cataract. And this was based on a work that I did for seven years screening for glaucoma in the community, and we screened 8500 people.

But naively, when I started this project, I thought that I could set up this whole project within a year, launch it have it running. And I had timelines for every step of it, acquiring the system, acquiring the equipment, acquiring the -- you know, hiring people. And then IT security, I gave it for a months. And wrongfully thinking that IT security would take four months, it took a year. The server to be approved took three months at Columbia. The IT security to be approved took eight months and for multiple reasons. We were the first in a lot of them.

The electronic signature for consent was the first, so we had to tackle that. Having a mobile unit move around transmitting data to the institution, we had to tackle that. The question is can we mix it with the electronic medical record or not mix it with the electronic medical record; we had to tackle that.

But thankfully, it's launched and we've

been -- we've had a pilot and we've screened over 300 individuals with results but that's not the place for -- to talk about it. But in reality, this mobile unit goes into high-risk communities screening them for, as we said, for ophthalmic disease in addition to diabetes through hemoglobin A1C, blood pressure and BMI. And in this system, we created tunnels to maintain the data inside a closed system so there will be no leaks of the information that's being transmitted. It goes to its own server and it's protected in that server and there will be no leaks anywhere in the system.

And our system is as secure as the ambulances in New York or even more secure, some people told us, than that through the way we created the security in it.

But I was asked to answer some of the questions. The first one was what are the most effective methods for mitigating risk for ophthalmic digital health device. And from that question, I was asked the methods to tell -- to limit intended and users. So all our users have individual-issued IDs and passwords for the application, for the network, for the server, and they're not the same password FYI. And they're issued by the administrator. In addition, all the users have to change their password every 90 days. So we maintain that, we update that, and we continuously monitor that.

Labeling for information our individuals or participants in the study, they usually have to enter their information on an iPad. This is their regular information, protected health information in addition to answering a questionnaire about their health and their habits. So we -- those individuals don't require a password because their privileges are limited. They only have two screens. One is to enter their information, the second is to answer the health questionnaire. They cannot surf this iPad. They cannot look at anything else, and they cannot go back. And we have somebody assisting them during this process, so no alteration after they enter it.

But the challenging part which I learned, too, is training and tutorial modules. We developed a comprehensive system that requires PDF instruction guides to references; video recording tutorials; onsite training, scheduled or nonscheduled; screenshots that's everywhere for them to use the system without any identifiers; Retraining when we notice that they require retraining; and every now and then, we keep updating the system so we retrain and retrain, and we do do report cards.

And as Ken said from the prior panel, actually, I do audit the data that's being entered into the system. And I learned that after having the first month happen and I went back into the data, and I notice you do need to audit it every now and then. And according to my audit, I decide if that individual who was doing the reading, because this data is being transmitted real time to a reading center; there's a doctor, ophthalmologist, or optometrist there giving the instruction to the individual where to follow up. And based on those audits, we retrain the individual and based on that report card, we retrain the individual more to to better serve and either image or give instruction to the individuals or comments or recommendations for follow up. And we keep updating our system based on what we notice in that system to develop better tutorials for those individuals.

The other question I was asked was to assist threats and vulnerabilities that should be considered and identified as a threat to the privacy of a patient by a digital health device developers. So in our system, we transmit to a server and we have our own independent server that is not mixed with the electronic medical record of the institution. And that made everybody happy in the institution for IT security. The data-capturing system that we built actually is offline when it's not in use. The server is always online but the data-capturing system is offline, and that protects any vulnerability or anybody trying to open it or hack it.

The other tricky part is monitoring patient behavior and location. As a lot of you know, there are there are few states that have tele-ophthalmology licensure, like Maine has a tele-ophthalmology licensure but not all states. So as a physician practicing, let's say, in New York, I cannot -- if I don't have a license in New Jersey, I actually cannot practice telehealth in New Jersey.

So with our mobile unit, we go to areas where the reader has a license. So we have some people are licensed in New Jersey so when the mobile unit goes to New Jersey, the reader is licensed and can practice. But personally, I'm not licensed in New Jersey. I cannot be a reader when that happens why.

NYP, or New York Presbyterian Hospital, has telemedicine initiative and they have urgent care visits. They have virtual visits. And in those visits, they actually enter a contract with the patient, legal contract that gives them -- they sign that they are presiding in a state that the doctor that they are working with has license in and the legality behind that. But right now they're developing a geolocation into their app. So basically, if that patient, although resides in New York and the doctor has a license in New York, they go to, let's say, Wisconsin, the geolocator will notify the institution that this patient is not in New York. And if the doctor does not a license in Wisconsin, then the app is turned off and there's no virtual visit with that individual at all. So that's a different way of dealing with that.

Sorry, I forgot to do this.

So in general, these are things that we do to protect against the hacking, to protect IT security, and to train individuals for telemedicine. Thank you.

DR. SPRUNGER: Lama, thank you for presenting that, your experience. I think a lot of what we've discussed today is balancing safety yet convenience. And you're storing on a separate server. If that person then becomes a patient in your hospital, I would assume there's no crosstalk there. So do you have to start all over? And does that cause an inconvenience as opposed to being secure?

DR. AL-ASWAD: So a couple of things. With this initiative, we're not always working in the area the hospital is, and we've learned from our project before that you need to create systems where it's convenient for the patient to follow up. So we've contracted with safety net hospitals in the area that the mobile unit is, and those patients are sent to them and they, the patient, is given all their records and they can go with their records to that institution.

At Columbia, right now, we're working to merge -- create a different system that we can merge our information. Once the patient comes to our hospital, we merge it with our hospital so the data is available, but once they come. We can't guarantee every patient is going to come there.

DR. HUMAYUN: Okay, great. So I think next is John. If you could --

MR. REITES: Okay. I'm going to build on from what I was talking about earlier this morning just to give you some perspective of our project, what our company was doing. I mean my story is really quick. You heard blurbs of it, but I spent all this time in clinical research specifically and just realized that there were all these different stakeholders that needed to see patient-generated health data; right? Everybody needed to see it but they had a very different reason and purpose to see that data. The patients wanted to see feedback on the data so that they could feel engaged and know what was going on. A researcher or a provider wanted to see that data so they can make a decision maybe at the next telehealth visit or what have you. A sponsor of a study wanted to see that data at a macro view to make sure their investment was being triggered and that the patients were being enrolled as planned.

And so there's all these different stakeholders involved, and we kind of saw this ability to have this omnichannel experience as something that not just the patient needed but also the site, so the researcher needed but also the sponsor needed. And in doing that, one of the things that really came to fruition is the need to sort of make interoperability happen but not interoperability at sort of this high level that we talk about with maybe EMRs or other big assets. But if we were to come in and we were to collect data from a patient, remember that all-- there's a lot of different ways we can collect data from people, and they can be a medical device; they can be a consumer wearable; it could be, you know, scraping data off their phone; it could be authenticating them through KBA or some other technology. There are literally 37 ways of -- ways two ways you can collect data from a person through their phone.

And we realized that there were a few people that were nailing this piece or nailing that piece but really, we felt like the industry need to put all that together. So that's what we did. We put together a system that would help us to roll out and in one omnichannel patient experience, collect all the different data they need.

And the reason we did it is we did a ton of patient-focused insight work. So we went out and talked to patients and providers and actually got people's insights. And one of the things we heard over and over again is, especially in our world, that we had patients downloading three apps and two websites to do telehealth, provide an e-Pro, and connect a medical device.

And so it wasn't that the patient wasn't altruistic or wanted to contribute data or be involved. It was like they couldn't figure out all the tech. And so we're talking about like usability; you know, it's this button in the right place when really, we're not even -- we weren't even giving patients like the ease of just having everything in one app. And I know that sounds really simple and a lot of people that aren't tech, too, will say, oh, just put it all in one app. It's not simple. It took me like nine years to figure out and break and make a lot of successful mistakes in pilots and studies to figure out how to make this work. And so that's really the -- sort of the framework in which we see things.

And so I know it's hard to see on this visual but one of the questions that I'm tackling for the panel today is really, you know, when we're looking at these risks, how do we start to tackle training and helping people, helping our patients to actually do something we give them to do. And there's a lot of ways, there's a lot of tactics to that, but one way that I want to throw out to you, because we've really found some some really early progress and success with this method, is instituting what we call eDROs. These are electronic device reported outcomes. And essentially, what these are is another acronym because you know in our industry, we like acronyms so we just made one up. But the reality is is the acronym's important because what this thing does, what this eDROs is it takes an activity that a patient needs to do and it combines all those things together. So for instance -- let me give an example of what an eDROs is and it's a really simple one; actually, it workshop an app, an active task in Apple's research kit, but it will give you a framework for this to start.

So what this task looks like is we've been able to do instructional videos and tap training for a patient, and then before they -- so let me back up. Let me give you a for instance. So we've got a mobile spirometry and this mobile spirometry in the study requires the patient to do an e-pro, so they've got to do a survey. They've got to be trained on it. They have to make sure they do the reading exactly like they need to do at home. And then when they're done, we need to confirm that they completed that task correctly.

So think about all those different things they need to happen. And what we did is we combined all that into one activity. So patient gets on their phone, gets a notification or reminder and says, hey, it's now time for you to do your spirometry; they click button; button opens up activity; activity says, okay, John, let's walk you through the steps you need to do to do this.

And so it starts by training the patient, making sure they understand. You can put a quiz in there if you need to. And then it says, okay, now you have to do the activity, let's connect a Bluetooth device.

And so what it does is it takes something that could potentially be really complex and tries to make it as simple as possible so that any user can do it. And what I'll tell you is that -- what we found that's also exciting is that this doesn't have any limits in age and demographic in that we have patients of all different ages and different therapeutic areas using these app tasks with success. Doesn't mean they're all perfect but it does mean that we're seeing early success in the way that we're combining the effect. Does that make sense? So combining this is really a way that we're tackling the training.

And then the last piece is I wanted to also make a few statements about sort of these threats and vulnerabilities and data privacy, because obviously this is a really huge thing that we're -- that we have to be careful for. And there are a couple of resources that I direct you to. One is recently, with FDA and Duke-Margolis, we actually went through a process and released an in-health action plan. And in this health action plan, we didn't just describe the types of data you can collect on these devices. We didn't just give you a bunch of use cases, but we actually talked about some some practical things you can do to secure data privacy for patients. And so if you're interested in that after, we can give you that link to that information.

But one of the things that comes throughout sort of that plan and, frankly, in all the work we do every day is that we think about these different modes of dealing with patient data -- I want to start by saying the biggest sort of question people have is how do we data transfer; how do we use APIs; and we move data around. aren't we impacting patient's, you know, privacy; aren't we moving their data around? And what I would tell you is that one of the ways we've accomplished keeping that data private and secure is by doing tokenization.

And so if you're not aware of what tokenization is, tokenization is if I'm John Reites and I come into a study, when I come in that study and I enroll, my name is then turned into a hash, is turned into a really complex token, and then that token has data assessed with it and it separates my data from PIII to PHI to clinical data. And and it takes that data and parses it into completely different cloud servers. And so what you're doing is you're losing the ability to re-identify a patient, but you're really taking the most extreme stance on securing someone's privacy. And in a clinical trial, this is really what we've seen to be valuable.

And so when you go through that tokenization service, even though the patient and the app knows it's talking to me, John, in the data and everything else that we see, I'm just patient 00123 and all my data is completely separated. And so when you do that, your ability to do data transfers and API integrations from EMRs to other assets really opens up, because the data security and privacy of the data becoming public becomes a lot less of a risk.

Real quick I just want to touch on two other items. I know we've talked enough about local versus cloud storage. And I mean my two cents is that you should be using cloud. There's too many reasons to use cloud. And what I will tell you, even when I'm working with academic and healthcare centers, I would tell you two years ago, I definitely saw sort of this push for On-Prem. We're seeing huge advances in that in our own work. And what we're seeing is that the academic and healthcare institutions are learning more about other compliances for ISOs and SOC-2 to and other sort of data security and privacy things that you need in your cloud. And so if you're not aware what those are those, those -- there's a good educational component to know how cloud is actually providing, in a lot of sense, more secure storage than even your local Prem.

And then the last piece I want to touch on is this patient authentication. So I want to flip this discussion a little bit and throw out just one new piece, is we're talking a lot about how to how to keep a patient's data private and that's appropriate. But on the flip side, remember when we're working in today's digital health world, 99 percent of the data I'm getting in the studies is from a patient not in a clinic. They're at home. And so the question I would actually reverse is privacy aside, how do I make sure the person doing the data is the person I signed up in the study or is the person I'm actually treating. How do I know? You know, you've seen this old classic image of how you know the dog's not on the computer typing away or how you know the Fitbit didn't get put on a dog. Have you guys seen these things? There are a lot of different ways to actually authenticate a patient.

And so I would actually tell you that in this data privacy world, the other piece to keep in mind is how do we authenticate; how do we make sure people are who they say they are as they actively contribute and provide remote data. So lots more we'll talk about in the rest but that's it for now.

DR. HUMAYUN: That's very good. Again, the way we've structured this, each panelist will give a brief talk, and feel free to ask any questions during their or after their talk.

I had a question for you about tokenization. I mean I think that's good to take a name and turn it into this token, but as we heard earlier, you know, for us, a fundus image or iris image may be an identifier. Have you thought on it and, you know, have you guys thought about how to tokenize something that's very characteristic like an iris structure or a retinal structure when you're actually looking at findings in that structure so you do have to display it? Do you -- you know, do you somehow just decode the information, blur their -- I mean how do you -- how would how would you think about doing that? So John or Nitin.

DR. KARANDIKAR: Yeah, hi. So we actually have thought about this quite a bit. Interestingly, we actually asked DHS, the Department of Health and Human Services, if retinal images are by themselves considered PHI for HIPAA reasons. And frankly, the answer was kind of unknown. They didn't -- there's no real ruling on whether an image by itself, even a retinal image, is considered PHI purely. And this Kaggle competition, for example, has these hundreds of thousands of images, right. If that was PHI data, then you can imagine that's almost a HIPAA breach. But I don't think by itself it is.

But the challenge comes when you're -- and as we do, when you're combining the image with a patient's demographic information. Then it's clearly HIPAA information. And so what we are doing is that -- where tokenization comes in -- this is a great point John brought up -- if you separate the demographic information from the image storage and you're keeping the images in a secure location with -- essentially "hash it" identifying the image and you keep the hash back in with the patient demographics, you can still match those up for the purposes of analytics. But by themselves, then, you know, that makes it a lot more secure. So that's kind of how we are addressing it right now.

DR. SPRUNGER; So our next panelist will be Mike. No. We've had all our Mikes in the last session so we'll move on to Nitin Karandikar who is Vice President of Engineering for DigiSight Technologies. There he leads all software development activities and architects new functionality at the -- for the company's mobile cloud-based technologies. He's been doing this for 25 years.

DR. KARANDIKAR: Thank you, but you can call me "Mike."

(Laughter.)

DR. KARANDIKAR: All right. So like let's see here. Okay. Dr. Sprunger already talked a little bit about background. I've been doing this for a long time, been doing health technology from different aspects of it for many years as well across a variety of companies including security, HIPAA compliance, essentially enterprise integration, all of those things, within a variety of different solutions in health care, created or had teams build provider mobile apps, built in the HR for a while and then patient portal, home health, so a lot of texture and different things there.

One point about that about my background -- so my background and focus is on software really. I'm not a device guy so I was a little bit concerned about coming here, but it looks like the worlds are colliding, right, and digital health is going towards software increasingly.

A little bit about DigiSight Technologies; so we -- very easy to use technology solutions for ophthalmology providers at the point of care, so a lot of you here directly. It's composed of an iPhone-based app with a hardware imaging adapter that's a class 2 510 exempt device. And then on the backend, we have servers in the cloud, the ubiquitous cloud. We can certainly talk more about that. And we have integration, so we have HS7 and Diacom integrations for EHR impact systems.

What we do is we provide -- essentially streamline the workflow for providers to capture images and patient data, collaborate among the -- among providers and provider networks, and then document that information with the EHR in the back system. Obviously, we are HIPAA compliant. Our security is a core requirement, a core value for us. And then my role, as Dr. Sprunger said, is to lead the software development.

So mitigating risks; so if you were to design a new digital health software systems (sic) from scratch, what are the kinds of things you need to think about from a security perspective? So first of all, security is kind of a complex and evolving issue and frankly, you're never done. It's a process that you're continuously, you know, trying to improve security over time. It's a little bit like securing your house and, you know, you can lock the doors, lock the windows, but, you know, somebody could come through the walls. You continuously keep working on that.

Today's software systems are composed of multiple tiers. There are many different points of vulnerabilities and so you want to think holistically about the system security as a whole and basically build the security in layers so that an attacker gets, you know, progressive walls that they have to break through to get the data.

In terms of safeguards, there's a ton of stuff we do but I just want to talk about the top three things that I focus on certainly. One is encryption. You know, encryption, encryption, encryption, those plus three. But encryption, really, at every point where data is stored and during transmission, both over the internet and also within your network, that really helps you even if an attacker gets access to the system. If the data's encrypted, it's a lot harder for them to access it.

Second is employee training and comprehensive training for employees about policies and procedures. HIPAA actually mandates that so it's kind of part of HIPAA compliance. This really goes to -- you know, the previous panel was talking about internal, you know, people doing things inadvertently inside the organization. This also targets like social engineering where somebody compromises a valid user's activity. So all of those things, the more trained your users are and also if your (inaudible) on what each employee's role will be if there is a breach and getting ready for that, that really helps to put you in a good position, because with healthcare labor, it's really at some level a question of "if" -- I mean "when not if" there's going to be an attack. And so you want to be ready for that.

And finally, login and access control is pretty self-evident. Everybody, you know, you want to have the appropriate access at the web EPI level, at the -- for web apps, mobile apps, at different stages in the system.

So data storage is one of the questions we want to talk about. One approach we take is we try to get the data from the mobile app to the server as soon as possible and delete it from the mobile. So as soon as the app connects, essentially move the data to the server if -- so if a user wants to look at on the mobile, we re-download it and we interpret it, of course, so that at any point, there's less data accessible on the mobile and this also mitigates -- you know, device device loss is a real issue. People lose their phones and so you want to have the reader back on the server.

And one interesting thing we've seen repeatedly with customers, there's a lot of connectivity issues at the point of care in larger systems and practices. And this makes it extremely challenging to get the data to the server and it makes a difficult problem, you know. for for making sure that you are making the data not just secure but available and you have to make it reliable And so solutions we looked are caching and synchronization. We looked at like adding two-faced (inaudible) -- for the computers science folks among you. So there's different ways we can do to mitigate that but it's it's a serious issue for us.

Storing it in the cloud, I know there's this sense that the cloud is less secure and there was -- I think one of the panelists here talked about, you know, the Brink's truck versus the, you know, let's -- to pick on somebody -- 7-Eleven, you know, getting -- you know, there's a robbery there. And actually, the Brink's has a lot more security. And so in some ways, you could imagine that on the data cloud providers, if you go to a large reputable provider, they actually do a better job of securing the servers. And no offense to the ID things, right, but this is what these folks do like, Amazon Web Services or Google Cloud and they live or die by that. And so that is -- you know, as long as you have a BA with the provider and the provider is a well-funded, you know, reputable service, you might actually be in a better position to do that.

And we had some real challenges in the past when we were in with with a smaller cloud provider, and since we moved to a bigger, more serious kind of provider, life has become a lot easier.

And then there was a question about the data transmission. Seemed to have lost the slide there. Go up one further. Oh, yeah. And then in terms of data transmission through EHR and PACS systems, what you're trying to do is you're trying to get the data from your system to the remote system. You want to get it there securely. You want to get it complete and accurate, and you want to fit it within the provider's workflow. So you want to meet all of these criteria for it to work well.

So one of the key challenges with health system integrations, and I have suffered through this for many years -- one of the key challenges is matching patient records matching or matching MRNs or patient IDs across systems, and you can lose a patient demographic vector, and depending on what the partner has, to match those records.

And interestingly, what we are seeing is multipoint integration, so you are matching your health data with multiple systems within a given partner. So they might have an HER system for patient data, a back system for images, SSO system for single sign-on, and so you have to really orchestrate the order of the calls across all the systems at the partner site to make that workflow work. And that gets pretty hairy sometimes.

So that -- and you still have to do all of the other stuff like patient matching, you know, across all of the systems at the partner site. There are many different organizations or some organizations within the partner site and you're to make it all kind of work together.

And then there's the usual kind of IT things like transmission endpoint security; we can always go a lot more into that; completeness, accuracy, downtime. There's a lot of challenges to the system with system integrations. They're all solvable but it takes a lot of work and you have to kind of plan for -- around a lot of these.

So I could go on for a while but I know that we have time limit so I'll stop here. I'm looking forward to the discussion going forward.

DR. HUMAYUN: Any questions for Nitin? Thank you. So far we've covered a lot of the server and also software approaches, but now we're going to switch to David, and he's going to talk a little bit more of about hardware so please introduce yourself as well.

DR. MYUNG: So, hello, again. David Myung. I spoke briefly earlier and I'm Assistant Professor Byers Eye Institute and Co-direct the Ophthalmic Innovation Program with Mark Blumenkranz, but also, recently, Darius Moshfeghi passed the baton to me to lead the ophthalmology telemedicine -- the ophthalmology effort at Stanford and at the VA, so some interesting perspectives there, and a lot of learning today about that.

So my talk is, again, switching gears to hardware and Bakul mentioned earlier in this -- in the day today that now, you know, these medical device companies are, you know, out of 10 people only 1 is a hardware person and 9 are software people, and so I kind of feel like this one's for you, the 1 in 10. This -- you know, this is for the hardware -- hardcore hardware device engineer. It's also actually a lot of other shout-outs during this one because it's a bit of my own personal journey through this process of learning this process of mitigating risks in ophthalmic digital health devices through safeguarding in hardware, specifically on light hazards and the light hazards safety and electrical and EMC standards, EMC being electromagnetic compatibility.

So we're doing it through a kind of a case study and first of all, as a we have disclosure, I am a co-inventor on this ophthalmic camera system called Paxos that was actually licensed by DigiSight Technologies, which I'm now consulting and helped developing it and as a design consultant.

But this is really a kind of a story over of an aspiring and somewhat confused entrepreneur and inventor -- would be inventor and with an idea that many of the people actually had as well on using Smartphones. This is almost six years ago now. Some of these images are from about six years ago and a lot of people are looking at trying to take pictures of the eye with the iPhone which had just gotten to the resolution and camera quality to take pictures as an ophthalmic camera.

So we were -- I was coupling -- we were coupling and ophthalmoscopy lens to the iPhone through, first of all, some plastic parts that I ordered on Amazon. Then we started printing them in my friend's bedroom. He had a 3-D printer in his room. He had a bed and a 3-D printer and we would just -- said, "Make this" and he would make it, and then I'd attach the lens and we would take pictures like this.

DR. HUMAYUN: Did he print his bed, too?

(Laughter.)

DR. MYUNG: Yeah, maybe his friend. So, actually, actually that's him right there, Alexander, and he gave me permission to use his

PHI there. But it's -- you know, he helped drive this. But what I was was -- we were stuck. My co-developers told me we were stuck. They're like, "What do we now? We can take these pictures and -- but how do we get it to the next level. How do we get it in people's hands?" And all -- life all changed when I -- a mentor of mine many of you know, Dr. Emmett Cunningham, actually put me in touch with Dr. Eydelman who then introduced me to Brad Cunningham. I thought they were related but they're actually not.

He -- and (inaudible). He is actually not here. I think he's leading a relief team in Puerto Rico right now which is one of his many hats that he wears, amazing guy. But it was a conversation I had; I was in a parking lot of the county hospital and I was talking to him and I was telling him exactly what we were doing, and he said -- and just like, Can you help me; what's the next step?"

And he said, "Well, the next step for the FDA is we want to know, you know, what are you doing, the what." What is the -- what is -- and what is -- what are you're doing? You're -- -- this is an ophthalmic camera; you're putting light into the eye. You know, FDA cares about things that go into the patient's body whether it's a drug, a device or a device (inaudible), so we want to know what is that light source; what exactly -- what's the source of the light; is it an LED, is it halogen, is it a xenon light source, is it -- what's the intensity; what's the spectral of characteristics.

And for me, I mean literally, it was like the light bulb going off. Said, oh, okay, so that's taking it step-by-step. That's the next frontier to tackle. Characterizing that -- there's a (inaudible) another set of standards, very (inaudible) of standards that he turned me to.

And the other one was, you know, what is the, you know, electrical characteristics of it; you know, how is powered; is it plugged in; is it battery; is it plugging into the phone; does it use -- is it drawing power from the phone itself?

Those are really the two key things and so this is that story, a little bit of the process of getting eventually as a class two 510(k)-exempt device. And so we had a choice. At the time, we were using the light -- I was using the light source, (inaudible), of an iPhone light. But by the time I talked to Brad, two generation the iPhone had already passed. So I was like I don't even know which light source (inaudible). And in fact, there were many other Samsung devices and so on and so forth. So for me, it was -- well, for anyone developing these things, they're faced with a choice, like do you use the light drawing from the phone itself, or do you develop your own light source and you characterize that. And there's pros and cons to both.

The choice that we -- the pivot we made was to just develop our own, characterize it once and for all, and then let it work with other phones. Other people have taken other paths but that's the path we just had to take. We got some funding for the biodesign program at Stanford and then developed -- worked with a (Inaudible). You then do that kind of the nuts and bolts stuff of getting it sort of certified under these safeguards.

So this is the product code. Well, first of all, ophthalmic chemistry used to be regulated under product code HKI for almost all cameras and Ron Schuchard had mentioned this new code, PJZ. That was a huge turning point because actually, Jeff Shuren mentioned today, too, that there was an effort to sort of down regulate as much as possible these devices, because realizing there was such a huge volume of innovation coming through to really help, like this workshop's trying to, accelerate innovation. I think this was part of that.

So in April 2015, there was an announcement that new code had been announced and if you fall under their group one designation -- group one or group two, but if you get group one, then you follow under the PJZ and then you become -- you qualify as an exempt, 510(k) exempt device. So that was April of 2015.

And with that, talking about accelerating innovation, I'd been working on it for three years. Seven months later, we were on the market. We're registered as a 510(k) class 2 exempt device.

And so there is an algorithm for optical radiation safety, the ISO 15004 at the time, and now in 2016, it's actually the ANSI, the American National Standards Institute, Z8036 standard, very similar but there are differences. There's actually a flow chart that helps you navigate where you fall as a group one or group two. There's also electrical safety standards, the 60601 standard. I'll touch briefly on that in the next couple of slides.

But then also important are actually quality systems, having a 1345 ISO certification, and working with the group, we worked with a dive shop that did that. Basically, to me, that says don't print this device in your friend's dorm room, make it in their garage, work with someone who knows how to manufacture these devices and let it be safe for the public.

The second one is actually risk management. So, you know, what are the -- you know, we declare what the risks are. Does the device have sharp edges that can cut the user; is -- if it has to be unfolded, does it unfold, you know, properly without breaking down or wearing out; are there small pieces? Actually, I was looking through the 49-page document that we have on this and, you know, there was a part about are there small pieces that a child could swallow. These are all things that are important because these are potential hazards and what strategy that can mitigate those things. So they all kind of work together -- sorry -- so those four things.

And then just -- can you go to the next slide, please? If you comply with those thing as a general package, then you can fall under product code PJZ. The next slide, if we can advance it -- I'm trying to remember what the next slide was -- oh, the electrical safety standards so it's really --

MALE SPEAKER: (Inaudible).

DR. MYUNG: Yeah. The main thing is that for electrical safety, it comes down to two things. I like to boil things down. I like to boil things down. It's immunity and emissions. So one is is your device emitting some kind of energy or radiation and what are the implications of that. And number two, what -- is it immune to electrostatic discharge. So there's actually a test where you take a device and you give it electrostatic discharge in different places and you sort of record its performance. And then in terms of emissions, is it interfering with an antenna that you placed in a device. So it's a pretty well-subscribed set of performance criteria.

Just as a way of conclusion, so as we all know, no mobile device technologies have continued to evolve quickly. So since then, we're on iPhone X or 8 right now. The FDA has put in a place a set of straightforward guidelines for building safeguards in new devices anticipating all this change. Two of the main (inaudible) related safety issues are inherent to what you're doing so, so what is a camera. A camera needs light and the light needs energy. So you need light source characterization and hazard protection and then also electrical safety. But along those lines, two quality systems and risk assessments are critical. You need to have those in place.

And the other sort of comment that I'd like to make is that much like the theme of this workshop, I think what I learned from this sort of personal experience is just how approachable and accessible the FDA really is and talking to -- I've been talking to Brad Cunningham and then Michelle Tarver, Malvina Eydelman, and Ron Schuchard as well -- just how approachable they are, because they really do want to help us would be inventors, would be startups, companies accelerate their ideas into market just do -- so in a -- through a process and a safe fashion so thanks.

DR. SPRUNGER: David, this may be a very

simple question but for someone who is starting or someone wants to start something now using a phone, did you measure the actual light intensity or did you assume that the product specifications from the manufacturer were accurate?

DR. MYUNG: Oh, no -- yeah, you have to measure them. So there's -- with the 150040 and now the ANSI standard, there's a clear -- it's a, I don't know, 15-page document that goes through under different types of conditions. There's a -- first of all, I was going to show you -- there was actually a test set up where you put the device, shine the light. There's a radiometer and a bunch of other things, one in UV spectrum, one in the yellow light spectrum, and you measure the intensities at at certain wavelengths and under certain conditions at different working distances. And you have to record all those and you have to say in that test certificate whether you've met, you're below that threshold or not. And if you're not, then you bump up to the next level. So it's pretty -- it's very much a test and it's not something that's that easy to do as an individual. So we use a test house. So we use -- we outsource it to a group that -- a third party that can say, yes, you you've passed all these tests.

DR. HUMAYUN: So, yeah, Dave, I mean I'm always for not inventing anything that I don't have to --

DR. MYUNG: Yeah.

DR. HUMAYUN: -- or building anything that I don't have to, but I've learned with these bulbs that -- or light sources, the -- depends on how long you've used it or what --

DR. MYUNG: Yeah.

DR. HUMAYUN: -- period, there is a degradation.

DR. MYUNG: Yep.

DR. HUMAYUN: And, you know, I did some work and currently, you're doing some work in spectroscopy and there, it really does matter very much so. Can you comment on -- you know, we're using -- we're taking these devices and saying they'll have good light and, you know, the illumination will stay pretty steady. Is there any work done on the iPhone or the Droid, you know, how well those those light sources work and flashes, I mean after how many uses and so forth?

DR. MYUNG: Yeah. You have -- you do -- I think you just have to do the work, the characterization. I mean they -- it's not that they think one source is better than the other, but -- the agency, but they just know what you're using and how it works. So with ours, when we went to the external light source, we -- first we had to pick the battery, what kind of battery to use, a D battery or a little calculator battery. And it turns out not every battery is the same.

I just -- it was this whole new world for me. It turned out two of these CR2032 batteries sitting side-by-side and wired a certain way gave the longest life. And even then, it was a little bit of degradation over time but it -- compared to the other configurations where they pooped out in, you know, several hours, this one lasted some 18 hours.

But then, you know, prior to that, I was working with an iPhone 4S or 5 and first of all, if you're using that, the iPhone battery just drains very quickly so, you know, in two hours, you know, we were using it in a clinic and it would drain and the phone would get hot.

So that's when we were kind of like well, we want to use this in the developing world for instance and, you know, I don't know how that's going to work, so I might as well send them with a bunch of batteries, these are coin batteries. So I think every -- that's why every phone is so different. Some phones might have a brilliant light source that just lasts forever and ever but it's not touted as a major feature. Even if it was, you have to still do -- go do the work, because the moment you use it as an ophthalmic camera, it becomes a medical device so you have to -- you know, you as the developer, it's on you to demonstrate that it fits all the criteria

DR. HUMAYUN: Yeah. And please feel free to ask questions. I have one more for you.

DR. MYUNG: Yes.

DR. HUMAYUN: You know, clearly, you're looking at the ANSI light standards and electrostatic discharges. What about human factors? I mean I think we talked --

DR. MYUNG: Yes.

DR. HUMAYUN: -- a little bit about it. I could imagine somebody doing something at home, scratching their heads and so forth.

DR. MYUNG: Yeah.

DR. HUMAYUN: How do you deal with human factors issues, and how do you control for that somebody with a tremor in their hand, you know 75 year old lady who's trying to get a picture of her retina?

DR. MYUNG: Yeah.

DR. HUMAYUN: How do you address the human factors aspect of it? I mean, again, a lot of the devices I've built, I've spent a lot of time on the human factor. It always is the thing that I don't want to deal with but eventually forces me to deal with it. So any thoughts along those

CERTIFICATE OF TRANSCRIBER

I, LUCY T. TURNBULL, do hereby certify that this transcript was prepared from audio to the best of my ability.

I am neither counsel for, related to, nor employed by any of the parties to this action, nor financially or otherwise interested in the outcome of this action.

November 10, 2017 ____________________________

DATE LUCY T. TURNBULL