FOOD AND DRUG ADMINISTRATION

OPHTHALMIC DIGITAL HEALTH WORKSHOP

October 23, 2017

7:55 AM - 4:12 PM

2 Montgomery Village Avenue

Gaithersburg, Maryland

Reported by: KeVon Congo,

Capital Reporting Company

C O N T E N T S

C O N T E N T S

(Continued)

C O N T E N T S

(Continued)

C O N T E N T S

(Continued)

MODERATORS: Mark S. Humayun, MD, PhD;

C O N T E N T S

(Continued)

P R O C E E D I N G S

DR. EYDELMAN: I m the Director of FDA's Division of Ophthalmic and ENT Devices. I am honored to welcome all of you to today's first Ophthalmic Digital Health Workshop.

Digital technology has been revolutionizing all of healthcare. Ophthalmology, with its ease of obtaining anatomical images digitally, has been inundated with opportunities to improve patient care via digital health technology. My staff and I are driven by our vision of bringing U.S. patients safe and effective medical devices in a timely fashion. Today's workshop, by fostering innovation in ophthalmic digital health will help us bring our vision to reality.

Before we start, I want to take this opportunity to thank all of our six co-sponsoring organizations for their hard work during a whole year to make today's event possible. And now I'm truly honored to introduce an individual who has reinvented FDA's oversight of digital health technologies. Dr. Jeffrey, our Center Director of

Devices and Radiological Health will now share his vision about ophthalmic -- about digital health innovation plan. Thank you.

(Applause.)

DR. SHUREN: Thank you, Malvina. It's a pleasure to welcome everyone to today's conference. I also have to apologize and send regrets on behalf of Dr. Gottlieb, our Commissioner. So I was supposed to open up the conference and he was supposed to end it. Unfortunately, he is tied up in hearing prep all day. I think he is testifying twice this week before Congress starting tomorrow so, again, sends his regrets, but this is an area of deep importance to him and to myself.

As you heard from Malvina, there are tremendous innovation that's going on and opportunity for greater innovation due to digital health technologies, I mean all the way from enhancing existing functionalities, like the opportunity to provide more precise placement of ophthalmic implants to entirely new

functionalities with learning systems and decision support to greater connectivity, connecting technologies so they can share information but also impact each other's function, and then to provide care remotely through telemedicine.

Now the north star for the Center for Devices and Radiological Health at the FDA is our vision, that patients in the U.S. have access to high-quality, safe and effective medical devices of public health importance first in the world. And it's not about a competition between countries. It's a recognition that we want medical devices to provide benefit to patients but is of limited value to patients unless they have timely access. And first in the world is simply a good metric for that.

Now we at the FDA face some challenges in achieving that vision when it comes to digital health technologies because the regulatory paradigm, while risk-based, is also very product-focused, and it was designed around hardware technologies. And even as they evolved to have

software, it was more as a component because digital health technologies are different. And when we talk about these technologies in the international arena, the term we use for "medical devices" is "software as a medical device" where the technology truly is the software. So software as a medical device, SaMD; and when it's in the device as a component, it's Sims. So we think about very different.

So hardware technologies; well, they have rapid innovation but it's more around the order of months to sometimes years and in very competitive spaces, we'll see next generation technology about every 18 months. You can learn a lot about hardware technologies by looking at them, taking them apart, kicking the tires, if you will. And their impact on patients tends to be very direct and observable. You can see changes to the structure or the function or measurable biological or physiological parameters. And the knowledge that's generated about one device is often transferrable to other devices within that

category.

But when we deal with software, it's very different. The innovation cycles are much faster. There are new challenges, like cybersecurity. You can't just look at software and have a good understanding of what it's going to do. And when you go ahead and test it, the impact may not be so direct on patient health. These may be impacts on cognitive and behavioral aspects of the patient or the clinician. And what you understand for one software program isn't necessarily transferable for other technologies even when they have similar functionality.

So it's a very, very different kind of beast and we started to revisit our approach on these technologies around 2010 when we started to receive inquiries from software developers, many of them not in the healthcare space but looking to enter it or just having entered it and wanting clarification regarding the FDA approach. And so we started to revisit how we think about these technologies. And at the time, we had already

cleared over 100 mobile applications over a period of 15 years but often for very traditional kinds of functionalities.

So based upon this deeper dive, we put out draft policy and guidance back in 2011 that we finalized in 2013 on mobile medical applications. And there were three key principles that came out of it; first, a recognition that we should not regulate unless it's truly value-added. So we were seeing lots of technologies being developed, very low-risk, functionalities we'd seen before in hardware but not lots of innovation going on. and we said we might better serve patient care if we backed away from it. So we engaged in the largest deregulatory effort we had as a sector in over a generation.

The second was recognizing that what really matters to look at is the functionality, that we would be platform agnostic. So if you made a software application and you put it on a ubiquitous platform like an iPad, we don't regulate the iPad. We'll regulate that software

and then the software developer is responsible for the whole system.

But the third thing is we recognized we need a new paradigm, we need a new way of thinking about these kinds of products, the old model that had been put in place simply didn't work. And then we started to further expand on that idea of got to be value-added, so you saw policies come out on what we call medical device data systems, essentially the technologies to receive and send, store and display information from medical devices and then applied that same approach on general wellness principles and general wellness claims.

Now you're going to hear from Bakul Patel in just a few minutes about our current thinking regarding our approach to digital technologies. You'll hear about our digital health innovation action plan, our pilot on pre-certification that we just launched in July. You'll hear about our efforts in interoperability as well as our efforts to expand this approach in the international arena and drive international harmonization. And you'll

hear in particular about a different way of thinking, not focused so much on the product but much more about the firms and the idea that if firms can demonstrate that they conform with excellence principles, let's say they're very good at software and testing, then we may be able to rely on that in lieu or some of all the kinds of evidence we might see pre-market, particularly for lower risk claims, allow those products out there. Then we'll gather that information in the post-market setting, feed that back on a levels of evidence approach that as we learn more, the applications for those technologies can expand.

Let me leave you with one thought -- that not only do we need to think about a different paradigm for digital health technologies, but we need a different way to approach it. The traditional model of government acting in a command and control fashion does not work well here. We have to do this collaboratively and today's conference is just a great example. We need to establish a forum or multiple forums where

we can bring together the interested stakeholders to work collaboratively and proactively to address common challenges and even unique challenges of the various stakeholder groups through a collective responsibility approach, what we call a "collaborative community," something that we are in the nascent stages of setting up in a variety of areas. But if there's any place where there's truly a need to problem solve in this collaborative community approach, it's here in the digital health technology space.

So with that, let me turn it over to Mike Repka to speak on behalf of the American Academy of Ophthalmology. Thank you.

(Applause.)

DR. REPKA: Thank you, Dr. Shuren, for coming out this morning and everybody else for traveling either far or wide for this platform. It's my pleasure to welcome you as one of the first of the stakeholder organizations that assisted the FDA in sponsoring this, the American Academy of Ophthalmology, who also recognizes the importance,

the mission, if you will, of innovation to improve patient care. So our 24,000 members congratulate the FDA on their willingness to be open.

We have -- the other sponsoring organizations will come up and if you guys can see who's next, I think it's Derek and then Ken. Thank you, Mike.

DR. SPRUNGER: Thank you, Mike. I'm from -- President of AAPOS, which is the American Association for Pediatric Ophthalmology and Strabismus. We feel this is a very important meeting so children can be represented as well in all this. It's -- it's pretty timely for us as we have a lot of interest in screening for ROP via telemedicine vision screening, so we have a lot of interest. Like to thank the FDA for allowing this to happen and also for the people, the organizing people. It's been a great group to work with. We look forward to a great meeting. Thanks for being here.

DR. REPKA: Thanks, Derek. Ken?

DR. NISCHAL: Thank you very much for allowing a representative of the American Academy of

Pediatrics. I'm Ken Nischal. I'm on the Section of Ophthalmology. The American Academy of Pediatrics has 66,000 members. One of the things that the Section of Ophthalmology's been very keen on doing is working on the interface between the primary care physician or pediatrician and the specialist. And we've been very interested in some of the new digital health applications for screening for amblyopia, which is one of the commonest causes of visual loss in children under the age of eight. So we think that the importance of digital health in getting to these children can't be underestimated. And again, thank you very much for arranging this.

DR. REPKA: Thank you, Ken. Dr. Afshari.

DR. AFSHARI: Natalie Afshari representing American Society of Cataract and Refractive Surgery and a warm welcome to you all. The mission of ASCRS is to promote deliver of cutting edge surgeries as well as promoting delivery of care by working with patients, medical agencies, medical communities as well as government. So

that delivery of care will be much more possible with the explosion of digital health. So a special thanks goes to Dr. Malvina Eydelman as well as the FDA for spearheading this effort. We look forward to a great meeting and a warm welcome to all of you. Thank you.

DR. REPKA: Thanks, Natalie. Mark?

DR. HUMAYUN: Okay. Good morning and it's great to be here. My name is Mark Humayun. I'm a Professor of ophthalmology at University of Southern California. I've worked for a long time with the DA medical Device Division, developed many products and perhaps you know me best for working with the Second Sight Argus II retinal implant, which is the only FDA-improved implant to restore sight to the blind.

But I'm here today as the President of the American Society of Retina Specialists, which is the largest retina society in the U.S. The mission of the American Society of Retina Specialists is to provide a collegial open forum for education to advance the understanding and

treatment of vitreoretinal diseases and to enhance the ability of our members to provide the highest quality of patient care. So this is a very important conference for us and a workshop. I look forward to working on the panel and sharing that. I'm looking forward to learning a lot from this and thank you Malvina, very timely to have this workshop.

DR. REPKA: Thank you, Mark. And finally, from Stanford, Dr. David Myung.

DR. MYUNG: Good morning, everyone. Thank you so much for having us here. We're honored as representing here -- representing Stanford here and the Byers Eye Institute at this really wonderful and important event. My name is David Myung. I'm a member of the faculty at the Byers Eye Institute. Also, Co-Director of the new Ophthalmic Innovation Program.

So Stanford has a number of collaborative educational and research programs in place with the FDA's Center for Device and Radiological Health. One of them is this educational one-year

year-long fellowship in ophthalmic innovation directly collaborating with CDRH. It's a project-based didactic, hands-on fellowship that teaches fellows the -- in a step-wise some often sequential or sometimes parallel stages in development needed for successful commercialization of new medical technologies. The fellows get to collaborate with members of our Department, other Stanford Departments, like Department of engineering, other Silicon Valley innovators and colleagues at the FDA. It's affiliated and, in many ways, inspired by the Stanford Byers Center for Biodesign, which teaches courses like biodesign for mobile health and biodesign innovation. So the fellows get to work with faculty members in our Department who've led the way in a number of new technologies, taking it from bench to bedside.

And digital health has been an important recent focus of our -- members of our Department. We've been trying to pioneer devices such as, you know, Smartphone-based visual acuity testing,

ophthalmic cameras, new ways of doing visual field testing, new applications in virtual reality headsets, machine learning and artificial intelligence and novel ways of doing ophthalmology telemedicine.

One of the things I'd like to do here is actually introduce our next speaker. Our first speaker of the speakers we have today is actually our inaugural innovation fellow, Dr. Zack Bodnar. Zack completed an innovation fellowship in June of this year, is currently one of the surgical vitreoretinal fellows at Stanford.

He has a very unique background and perfectly suited for this workshop. He graduated with a bachelor's and a master's degree from MIT in computer science. Then he worked in the tech industry for a number of years before going off to medical school at Dartmouth, then doing his ophthalmology residency at St. Louis University.

As the first ophthalmic fellow, he accomplished quite a bit; for instance, successfully drafting a mock pre-submission

package for an app that he developed as a resident that measures the degree of ptosis in affected patients using a Smartphone. He then -- he also coauthored an editorial in *JAMA Ophthalmology* on the very subject of this workshop along with Doctors Malvina Eydelman and Dr. Michelle Tarver from the FDA, so in my opinion, a true expert now in the emerging field of digital health in ophthalmology.

Please welcome Dr. Zach Bodnar to the stage to tell us about accelerating innovation to encourage new frontiers in ophthalmic digital health.

DR. BODNAR: Thank you.

DR. REPKA: One thing, Zach, before you start, just for minor -- you can come on up -- for the audience. We are on an extremely aggressive timeline today. There are many, many talks and I'm going to ask the speakers to adhere as carefully as you can to times, because we do have these introductory lectures, talks this morning followed by perhaps the most important part of the session which are the panels, the dialogue in

which CDRH, the Agency, is looking for guidance from the community on how best to handle the very difficult issues that are going to be presented. So Zach, go ahead.

DR. BODNAR: Thanks. Do I have a clicker or --

UNIDENTIFIED MALE: Well, you could --

DR. REPKA: -- which means I'm going to have to cut speakers off, I guess.

DR. BODNAR: Okay. Well, I'll try not to be the first. It's a pleasure to be here, everybody. I have one financial disclosure to make which is that I've done consulting work for DigiSight Technologies over the past year. And it's my pleasure to talk to you about mobile medical devices and digital health.

So there really have only been three technologies invented in the modern era that human beings are willing to carry on their person at all times, and that is the wristwatch, the credit card, and the mobile phone. Since the iPhone was introduced in 2007, the number -- the percentage

of Americans who have Smartphones has increased up to 77 percent, so they're pretty much ubiquitous devices now. And they're great platforms for the development of mobile medical devices and just medical devices in general. They have graphics processers, which are excellent for hardware acceleration of graphics, and they have high resolution cameras for the capture of photo and video. They have biometric sensors as well and they have very flexible dynamic user interfaces. The touch screen is very intuitive and can be customized in a myriad of ways that make it possible for medical devices to be customized. In addition, developing for a Smartphone or a tablet is a rapid process, which increases the rate of innovation and deployment of these devices.

Now, of course, digital health doesn't just encompass mobile medical devices and Smartphones. It's things that we're also familiar with that are becoming integrated into part of the internet of things as computational power is added to things that are traditionally what you could consider as

analog devices. So now we have digital phoropters and all kinds of things that are integrated with Bluetooth and wireless and other technologies that enable additional computational power.

Now that raises a question of what exactly is a medical device in this context now. So the FDA, based on the Federal Food and Drug Cosmetic Act, defines a medical device as anything that's not a drug but is intended for diagnosis, management or prevention of disease. And as was mentioned earlier, software in and of itself can meet that definition. In addition, a consumer medical device like a Smartphone can be transformed into a medical device either by adding software or by adding hardware extensions that enable those capabilities.

So there's great potential for this. We're already mentioned telemedicine as one application and you can see from the right side of the slide, there is already a pretty rich ecosystem of applications and devices that have been developed on these platforms. These allow patients to

personalize their health data collection. They enable home health care, disease monitoring, and will create innovations for screening, diagnosis and new management of ophthalmology diseases.

But the create challenges for device developers; namely, as was mentioned earlier, many of the companies and organizations that are developing these devices are not traditional medical device developers. They're software developers that are entering the space for the first time, and so they may not understand the rules and regulations regarding regulation of medical devices and whether or not their application even meets these criteria. Things like unmodified hardware change their risk profile when they become a medical device so while the torch on a Smartphone camera is perfectly safe in the hands of somebody taking their family photos, it becomes a different risk hazard and different profile when that same torch is used to take camera -- take pictures of the back of the eye. In addition, this all depends on the setting in

which the devices are used. For example, within the operating room, these devices have to be able to coexist wirelessly with all of the other technology that's available. And it depends on the intended use and the user that is trying to use the device. A patient is going to use a device much differently and understand its operation much differently than a physician.

Of course, all engineers -- experienced engineers know that small changes in these complex systems can have a cascade of consequences. They have robust quality assurance frameworks to ensure that the catch errors early but they need to think about these things in terms of safety and also the usage profile for patients. Small changes in the user interface, like the size of a button or the label or its color can profoundly change the way in which a patient may use it, and that can potentially impact its safety.

These are data-driven devices, which means that HIPAA considerations are important as well. It goes without saying that personal private

health information should be encrypted on these devices but it's also important to recognize that not everything about this is completely within the control of the developer. For example, the operating system and the hardware itself is developed by a third party. That means that when operating system updates come out, which are often done in order to address security flaws, there -- it's the responsibility of the developer to notify the end user who is then responsible for installing the device -- the update. But the ultimate responsibility for safety and security rests on the developer itself.

There are questions related to telemedicine as well. Because these devices are mobile, it puts them in the hands of non-physician users in different settings which raises the questions of when is physician oversight necessary, is there a need for realtime synchronous communication.

And when patients are the users, that changes the risk profile as well because there, the errors and patterns of misuse that patient has are going

to be different than those of an expert user.

As I said, most technology enterprises have a robust system, development cycle beginning with design, development, quality assurance including integration and unit testing that is pretty well-established and they can leverage that as long as they recognize that they are also testing for safety and effectiveness of devices. One particular issue is that because these are consumer devices, it's not possible to test every possible software and hardware configuration that might be deployed.

So one way to mitigate that risk is to limit the possibility of installing your application onto only tested configurations. Of course, as I mentioned, it's important to recognize human factors in the testing of these devices which involves robust testing of the user interface and documentation and making errors very clear to the user, to the patient so that they understand White House en there are safety issues.

But with all that, there's great promise for

this technology. It has many advantages. It brings the technology to the point of care. That can be the Third World, it can be the E.D., it could be a school. It improves the efficiency and automation of many routine tasks that physicians do. It streamlines communication between patients and their providers.

The other thing is because these are mobile devices, you have the opportunity to gain insight into disease states outside of the clinical in the interval in between visits, so for example, tracking intraocular pressure while a patient is at home or at different times of day. The network connectivity of these devices provides information into their actual operation while they're deployed in the real world which means that safety signals can rapidly be recognized and quickly acted upon.

The FDA has developed some guidance as was mentioned. The pre-submission program which I had the opportunity to be involved in is a great way to -- for device developers to submit a plan submission to the FDA and learn about how they can

improve their development plan and their testing for safety and efficacy. Initially, there's a digital health mailbox that device developers are willing to -- are able to send emails to you.

I just want to thank Mark Blumenkranz, Stephen Young, Malvina Eydelman, Michelle Tarver, and Ron Shuchard, who were my mentors over the past year. They did a great job in helping me to develop some of these ideas and really enriching my educational experience. And with that, I pass it on.

DR. REPKA: Thanks, Zach. We do have a question period at the end of this session, assuming we have any time at the end of this session. So our next speaker will be Mr. Bakul Patel, who is coming to speak about the regulation of digital health. Mr. Patel is Associate Center Director for Digital Health at FDA and has a longtime interest in sort of this area in business development improvement.

MR. PATEL: Thank you -- thank you so much. Thank you, everybody, and this cannot be even an under -- it cannot be said enough but this

collaboration and the need for collaboration in digital health is exemplar here in this room. And the more peoples trying to connect the dots is actually even more important as time goes n.

What I wanted to do today is just give you a perspective, sort of how we got here, things we have done, how we are sort of getting to a place where we can now all sort of starting to rely -- start to rely on these technologies that are becoming part of our lives.

And Zach, as he rightfully mentioned, there -- there -- enough convergence is happening in the space that's taking, really, healthcare -- or healthcare to all walks of life.

And this is really a slide that talks about how digitization and sensors and software is now moving to every part of life and it's becoming something so ubiquitous that computing power sensors, connectivity and software has become the leverage -- the data that we are generating and then converting that into -- for -- converting that into -- for health purposes.

So two years or three years ago, I think the conversation was just about variables. The conversation was just about actively tracking but now, I mean today we are talking about how can we use that technology in true healthcare purposes, and how can we use that connectivity that exists and the ability for people to have this technology at a much lower cost and how can that be sort of brought together and make it meaningful.

As Jeff mentioned this morning, we truly care about how to get these products to be patient-centered; how can care be delivered in a patient-centered way. At the end of the day, we all look for these technologies to be high-quality and safe and effective but more importantly, for us, as FDA, as -- for us as a community, how can we partner together to sort of create an ecosystem lean forward, you know, 23:23:45 that can get us ready for the digital health future. And that's really what's required at this point in time. I think we'll stop talking about being ready in the next five years but today, I think we need to be

ready. And how does FDA do that? So the talk I'm going to give today is about how we got here.

When we first looked at what kinds of technology and software exists in the world of, you know, medical devices and we chunked it into these three buckets; software that is simply a real medical device but on its own; software that it's inside a medical device; and then software that is used to make medical device, and they're all becoming important as we move forward. I have an anecdote. You know, a few years back, or actually, I would say 10 years back, there used to be, you know, one software engineer and 10 hardware engineers, and I think today if you ask anybody, any organization, it's the opposite. It's nine software engineers and one hardware engineer that is employed in any of those companies. And the reason is you can do many things very easily. The hardware has become so malleable that you can change it and change things that it's intended for by just changing software. And that has been sort of the trend as we move

forward.

How do you sort of take that and soft of bring the two to this place of healthcare and what does that mean? So you take those three things that I talked about, software that can show up as, and it raises a lot of issues and a lot of sort of topics for us to connect, to think about. And it's all driving by connectivity, driven by technology, driven by sensors so it brings up, you know, can it be used for diagnostics, too; can it be used for simple telemedicine. And in a way, that may actually touch -- cut across care giving to care managing across the spectrum.

And then you have sort of advances in technology in terms of artificial intelligence, machine learning, algorithms are continuously sort of evolving over time and what does that mean in terms of regulations. And as we talk -- and I'm hoping today we'll have -- get some ideas about how we sort of move forward in those spaces.

But one thing I would say, that connectivity also has raised many, many, many different issues;

cybersecurity and interoperability. There's a need for interoperability and then the need to protect that interoperability solution through -- and being careful about cybersecurity.

We do have policies that we started off in 2013 with mobile medical apps but just scratched the surface that really talks about, you know, how do we sort of move forward and focus on things that are really important and that can be value-added. We have, over the last five years, we have published many, many documents and it's really about how do we focus on the high risk functionality and really allow the low risk functionality to flourish and give a pathway for people that can create technologies in the high risk phase. And that's really how we've sort of been approaching this area altogether.

What does that mean? It does not mean that we can only do this locally here. Digital health, by definition, is global. Software, by definition, is global. That means that connectivity and operability of those products can be -- does not

know the boundaries of the geography, right. So we work together with our international partners and regulators to figure out what should that starting principal look like and how should we even think about going forward.

Over the last -- in the last five years, we've been -- or four years we've been sort looking at how do we sort of take the concept of software as a medical device and have a framework that can actually help people understand and get off on the same page. As Jeff mentioned, we can't touch and feel software when you look at it and the concept of, you know, not knowing what it looks like, what impact does it have starts off with a very fundamental thing called definition of what the product does. And it seems pretty commonsense but what we found over time is that description was not standardized so we ended up defining -- if you define it in a certain way, you will actually know where the software sort of falls into and what markets.

Now again, this is a principle level that

crosses the globe, that regulators across the globe have agreed upon and we just published, just a few weeks back, the final document in the series of our documents on software as a medical device that lays out a starting point for us to start considering in each jurisdiction what -- what's -- how should we approach it, how should the community approach it. And this was meant to be a guidance, a technical document for regulators as well as industry as we start moving in this area. So that's the second stage. So we have a bunch of documents that we have released that define -- that showed how you would focus on and where we would focus on in terms of regulation.

Internationally, we worked on coming up with the framework that led us to, most recently or last year, the 21st Century Cures Act, which took our policies that we had so far on where core focus should be and codified it. There are a few things in that 21st Century Cures Act which talks about certain types of functionality would not necessarily be regulated by FDA and gave that

clarity based on some of the things that we said we are not focused in our guidances, for example, mobile medical apps, general wellness, medical device data systems, that have become so ubiquitous and become such -- at a low risk that it did not warrant FDA to oversee it actively. Congress took that and codified that into law.

So moving forward, we will provide -- we'll be providing more clarity on how those policies are affected in the coming months and that's part of the innovation action plan we published a couple months back.

And then let me bring back to like the challenges Zack was talking about of why software is unique. We talked about -- and Jeff mentioned this morning -- is the pre-market timeline from a regulatory Perspective were best-suited for hardware. On the flip side, software development timelines can be continuous. Some of you may have heard about DevOps which basically happens all the time where things develop, created, tested and delivered to the customer on a realtime basis.

That is an extreme case of how software can be delivered to patients or users at the end of the day.

So just an example of how we have been thinking about the challenges and unique opportunities -- so challenges about developing software and delivering fast is one aspect. The other aspect is things change in software so quickly and we -- our approach to those changes with software, if it doesn't change, I think it may actually cause risk to patients because change is absolutely necessary. So how do you sort of take those unique aspects and turn them into public health benefit is something that we've been looking at.

The other point is how does -- how do we take that -- take the unique opportunities that software avails us because you can connect to the patients and the user directly and have that information be collected? How do you leverage that real world experience a user has with software and turn that into something that's

useful for patient benefits and for public health? So we've been thinking about those and that led us to this concept that we are going forward, is bringing those four things together and going to a paradigm that really talks about focusing on high risk products or high risk functionality that's aligned software development timelines and aligned with industry best practices but -- and on top of that, being consistent and aligned with the global regulators. That will -- that, we think, will yield better products to market that are safe and have high quality to all patients.

So this leads to, you know, what we are going -- what we are doing and what we are going -- how we are going forward with this. The concept emerged of how will you change the paradigm to not focus on product-by-product but rather than product-by-product, take it to a company level, we then trust the companies making these products, we -- based on their organization's excellence and culture of quality, we could imagine a streamlined regulatory pathway.

And what would that look like? The pathway looks like if you're -- if you know what kind of products you're making in terms of risk, with the framework that we use in the IMDRF, and you know what excellence looks like, you could afford a straight to market pathway for certain low risk products or changes that could happen over time.

Or the flip side is if you're not as excellent, are you building products that are at a high risk. We would come up with a paradigm that's different than what we have today in terms of how we review it. Imagining a software review done by a paper submission is just simply not cognitively, you know, connecting for me or for many others. So how do you change how we review, how we change what we review is something also we'll be looking at this -- during this development. It's all predicated on how we leverage the unique abilities of software that can collect post-market information, real world information, and then feed it back to the product itself.

We don't stop there. We don't want to stop there with the development of this program. What we want to do is we want to learn, as a regulator, that this program does not become static. We want to keep this program self-learning as well. So how do you take the evidence that's been created, how do you take the learnings that have been learned by the companies, and how do you aggregate that and bring it back to us to see whether we can make the right -- whether we have made the right choices, can we tweak and can we learn and grow and be scalable for -- as products and technologies grow, because we know this will change over time as we have seen already.

One other concept in the IMDRF document which we are adopting in our thinking for the new paradigm is how do you sort of take the continuous learning that can happen, allow for the right clinical evaluation to sort of take place so that there is a wide path for people to start with the small claims that they can make or small functionality claims they can make and learn with

product being existence in the marketplace and being used by users and prospectively collect that information that is -- that can supplement your next level of substantiation of claims. So how do you sort of do that more (inaudible)? So in essence, we are creating -- we are taking a chapter out of the book "*Agile*," that developers use today and applying it to regulatory paradigm and saying how can we be more agile, how can we be learning, how can we sort of allow -- how can the regulatory system allow for products to be safe, at the same time effective, not changing the bar on that safety and effectiveness but being more nimble and more iterative as far as products can get in the market. We feel that will help us get there.

How do you sort of get there? So many unanswered questions in this paradigm. This is exactly where we have started when we talked about what an excellence looks like. So I'm going to touch upon this really briefly. We will probably not have time to go over all of these, but this is

really what we thought the starting point for what excellence should look like. When we think about patient safety, we all agree that there are many things packed in that word "patient safety," which means people need to make choices for patients all along and throughout the entire product life cycle. What's -- commitment to product quality; commitment to clinical -- be clinically responsible, matching claims, understanding and doing the right evaluation, etcetera; what does that look like? Cybersecurity is one of the top topics right now and at the end of the day, we still want people to be proactive.

How do you have these five fundamental principles that we all care about embedded in every organization and if they achieve that, how do we give that credit back to those companies that can then be afforded a path to market in a more trusted way. So this, we believe, is the starting point. This will evolve over time.

I know next year at this time, I will be thinking about how these things refined over the

last year but really, what -- in the pilot, what we're doing is we're developing -- taking this concept and taking it to the next level and developing a program that will sort of refine this further, get people to include these principles in their own balance scorecard so they're monitoring, and we are able to just sort of observe and not have inefficiencies because somebody has to start reporting something to us or creating something special for us We are trying to leverage what exists in the -- by existing in the practices today in these organizations. And that's why we announced the nine participants to take what they do best, learn from them and create a program that is best suited for the community.

We want other input along the way so I would encourage every one of you to be -- stay engaged, pay attention to what we are putting out. We are going to do the nine participants but we'll share what we're doing with them publicly as well. So any input from that perspective we very much appreciate it.

So the concept here is taking the five principles, looking at it from five different -- four different lenses. If you allow people the process flexibility and we have allowed people for measurement flexibility but still anchoring on those five principles, can we determine trust; can we prove trust, not just by up front but also after the product's been in the marketplace.

So I'm packing a lot of this information that I spent an hour-and-a-half talking to the pilot participants, but this is where we are going and I'm happy to talk a little bit afterwards as well. But as we move through the program, we are iterating as well, we are learning as well. We are also taking what we're hearing from folks, building it in the program. But the way we are approaching to build those three big components is we are focusing the first few months on the excellence principles and how do we identify organization excellence.

The next phase we'll focus on -- and not to be very serial on this but we will focus on the what

should a review look like and we'll practice what a post-market observation looks like, what -- how do you get access to that data, perhaps change the word from "reporting" to "access" or supplementing the word "reporting" to "access," how do you do that. So we are going to look at different ways of one, informing the agency; two, informing public health; and three, sort of how -- keep us sort of current and learning along the way.

I think that was the end of my slides but I want to close with saying if we're -- it's high time we need to look at a paradigm that is best-suited for high risk technologies that are emerging, are collaborating. As you heard, players at med tech companies are not -- they're not previously considered to be part of digital health are now partnering with people that you have not imagined. I mean you -- if you think about the folks like Google and Omada Health and others who have not really been in this space and purely in general wellness are now moving quickly into medical device space, and how do you sort of

provide a path for them so they can actually deliver products, bring that innovation to healthcare that we all see for. Thank you very much.

(Applause.)

DR. REPKA: Thank you, Mr. Patel. So our next speaker is Dr. Ronal Schuchard who is a lead reviewer for Ophthalmology Division, Medical Devices at FDA. His talk will be on FDA Perspectives on Ophthalmic Mobile Medical Applications and Telemedicine.

DR. SCHUCHARD: Good morning. Thank you very much for attending and giving me the opportunity to share with you. From the point of view of the ophthalmology review teams that are looking at this area, the ophthalmic digital health area is a broad area.

There is a variety of topics that must be looked at when a device comes in for evaluation. You've heard many of these already from previous speakers, Bakul and others, that -- so today we'll look at primarily, because there's not enough time

to cover all of these, is software as a medical device, interoperability, and mobile medical applications and last, artificial intelligence or deep learning. Some of you, I think, are aware that, for example, the commercial applications of deep learning is driving this in a big way but obviously, it is overlapping into the digital health. There's a 100 percent investment, Tesla, et al, with autonomous college, if nothing else, in this area.

So the types, as already been talked about a couple of times, but just to reiterate; software in a medical device that we see; you're going to hear in a few minutes how perimeters have software that evaluates the results coming out. This is a prime example of software in a medical device where abnormal or normal results are part of the perimetry system. And also, we're seeing more and more and more OCT images where color-coding is an aspect of what is reported or what is put out by the device to enable better reading of these devices.

But software as a medical device is where the innovation and where the prime areas that are challenging us in terms of digital health are. And so the top right shows you one example of a cleared device that is using a Smartphone to do at-home testing of vision and, of course, you'll hear many times this morning about how cameras and other things within Smartphones are being utilized.

So it is a rapidly-evolving landscape. Many ophthalmic devices we already have and will continue to expand in terms of the digital technology. Software diagnostics, CADX, computer assisted diagnosis, and advanced analytics, which is the computer assisted detection are rapidly emerging. The greater connectivity and the interoperability is going to introduce new but it has greater aspects of the things that we must consider.

And so all of these areas are rapidly -- and I put a little slide there that shows *Moore's* law -- and yes, this is an expanding field that is only

going to increase.

So telemedicine in ophthalmology is here already. We are seeing that telemedicine is practiced in ophthalmology but to note that it is not regulated, the practice of telemedicine is not regulated by the FDA. There's many devices within the telemedicine or the tele-ophthalmology -- the -- there are many devices within the telemedicine world that are regulated but the practice. You'll find that there's a lot of what we refer to as medical device data systems, that is systems that transfer, store, display the medical device data but they don't control or alter the function of the device. And these devices in themselves are also as part -- as Bakul has mentioned, the 21st Century Cures has changed how we do things and functionality is no long part of the definition of a medical device.

So telemedicine systems, the devices within the telemedicine that you'll typically see are like ophthalmic cameras. Many of the cameras that produce the images, they are regulated and at-home

vision testers can be regulated although they can be class I devices; therefore, they would be 510(k) exempt, so we are seeing visual acuity, an Amsler grid and a variety of visual function tests. So the distinction, again, is that class one or class II 510(k) exempt devices within this world allows one to have perimeters with databases so that their class I. They can be part of a telemedicine program and there's not a distinction between class I and class II with group one light sources for ophthalmic cameras and group two light sources, and you'll hear a little bit more about this in the panel discussion.

In all of these things, people will come to the digital health world thinking that they can compare their devices with what's ongoing in telemedicine. But as I've already shared with you, since telemedicine is not regulated by the FDA. Oftentimes we have difficulty with the claims that are made by the telemedicine systems in terms -- especially in terms of the sensitivity and the specificity of the systems. And so,

therefore, a comparison to an unregulated or to a practice of medicine is difficult and one should be careful about comparing your digital health devices and especially in the ophthalmology world in terms of these sensitivities and specificities let alone the application to a particular patient which we would be looking at the positive predictive value or the negative predictive value of these systems or devices.

I'm having a devil of a time with -- so the categories of health IT is another category that is found within ophthalmic realm. There is administrative functionality. There are the health IT which it talks about admissions, billings, and a variety of those kinds of things, and there is the health management functionality, those aspects of things with managing a patient. But it's not until you get into medical device functionality that we see primarily that the FDA would start providing oversight in terms of those types of things.

then in terms of clinical decision support

software, which is also an area that is part of or allied with health IT, we see that, again, the health management functionality in terms of clinical health records and drug dosing or reminders of preventive care and those types of things are distinct from a medical device functionality where we start seeing computer-aided detection or computer-aided diagnosis and refraction treatment planning for your laser refractive systems or robotic surgery -- surgical planning may be coming, let along electrophysiology.

So the FDA perspective of review challenges is that we lack - there's often a lack of experience for the established device, that is that we really -- this is an innovative world. This is a world that we're seeing for the first time in terms of many devices so there's no clear complete description of the technology or device or there's no experience that we have on how one should describe these types of devices. There is often unclear indications for use and intended us

because it explores new areas. We consider things for eyecare clinical environment versus non-eyecare clinical environment or even non-clinical for devices that are going into the home or going into a school, for example. Those are completely different environments and, therefore, the indications and the intended use would be completely distinct and would need to be clearly specified. But there is often a lack of clearly appropriate predicate so this is, being the innovative field, not devices that are already there. And the risk analysis is inadequate given the risk of the device used and again, the environment where it's -- what population, the environment where it's going to be used is critical to be able to do this risk analysis, and here's -- the limited information is not allowing the evaluation. And this is -- this risk assessment is key to what's, as Bakul has already mentioned, is key to what we need to be doing in terms of looking at devices.

So for lower risk functionality, we find that

the device may not always be enforced in terms of regulatory requirements. These are lower risks and oftentimes are compared to the -- exceeding the limits of exemption but the higher risk, and this is not in other fields like mortality we often don't deal with. But we do risk -- we do deal with risk like permanent vision loss or other aspects of things that are risk functionality and that we need to assess whether or not the safety of the patient… And because of the innovation technology, we're going to find that many of our applications are going to be *de novos* because there's no appropriate -- and once they start coming in as *de novos*, then they get shifted to 510(k)s.

And it is unlikely that we'll see a lot of PMA applications, at least at first, but you may challenge us, those of you in the field developing the devices. You -- we may start seeing some of the future that we're just not foreseeing yet what would be a class III PMA device.

So again, the risk assessment is key and the

premarket assessments are to be able to fully define what are the key functions of the device, what are the aspects that are unique and actually key to this device, what aspects make the device vulnerable, what is the impact of that vulnerability and what protections are in place. And you'll hear panels discuss those protections that should be in place to be able to protect the safety of the patient.

The methods of mitigating the risk are also part -- a response to this, the safeguards built into the software or the hardware, for example, inherent in the digital health device, methods to limit the intended users so that's another approach to be able to say that we'll mitigate the risk by limiting the intended users or labeling provided for patient use. And finally, training modules and tutorials may be relied upon to be able to mitigate these risks.

The medical mobile apps is the area where probably we're going to see the largest expansion, as already mentioned. We don't see a lot in the

ophthalmology world the medical mobile apps that are not considered medical devices. You have your Smartwatches that tell you how many steps you've taken today or the variety of things that help you stay fit. We don't have that kind of equivalent in the ophthalmology world, although some but there's not a lot of them. Rather we see a lot of mobile apps that are lower risk mobile apps that meet the device definition. So what we're seeing a lot of these days are people that take Amsler grids and put them on a Smartphone or people that take visual acuity testing and put them on a Smartphone or a tablet. Those types of devices may -- or other types of devices may border that whether or not they are lower risk mobile apps that meet the device definition. But we don't intend to enforce requirements or that they're 510(k) exempt.

Today what we are trying to focus on is the mobile medical apps, the ones that really truly challenge and innovate and provide us new functionality in the digital health world, and

those types of things you're going to hear about today.

So to give you an example of the types of things that we have seen in the mobile medical app or the software as a medical device world, we've seen diagnostic mobile apps such as the DI or the Paxiscope. We are starting to hear about R&D in tablet video field assessment where you take what is done with a perimeter and you put it on a table and you go to India and/or rural areas of the United States and you do your visual field screening with a tablet instead of a perimeter device.

We're hearing at ARVO there was a symposium on computer assisted detection for diabetic retinopathy and at that ARVO symposium, there were several companies that identified themselves as talking with the FDA already to be able to submit an application. So soon you may see devices that are already CE-marked but they are soon to be on the U.S. market as well for computer assisted diagnosis of diabetic retinopathy.

Then in terms of therapy, we are seeing R&D for dichoptic treatment of amblyopia; people that are developing red-green glasses or virtual reality glasses with mobile displays, you see these in the press, you see these in publications, that this is being developed. These are virtual reality with tablet-type of technology.

There's R&D for wave-finding and object detection and assistive technology for devices for visually impaired.

So to give you a couple examples of disease progression aids and diagnosis, there's the myVisionTrack™ which I showed you before, is the Amsler grid on a Smartphone using circles. There's a Saccadometer Plus which which is an eye movement monitor, EYE-SYNC, which is a nystagmograph looking at abnormal eye movements; ophthalmic imaging systems, there is a large number of them that have come in with the cameras that put their images into an imaging system that then uses software to be able to do additional analysis but these PACS systems are, if they do

not do advanced analytics or other types of digital health, they, too, may be just class I or devices that do not need a 510(k), but that is the trigger that you will need to evaluate.

So the last topic I'd like to talk about is the interoperability. So we have a picture here -- I've tried -- I've stolen, I admit, a picture that doesn't truly represent the types of things that we see in an ophthalmology, but I would ask you to bear with me and think that this could be a laser refractive surgery system, for example. And that little picture of a guy standing -- or sitting at a workstation, let alone all of the devices that are interplaying with laser refractive surgery, are soon to be changed. So we have devices within a company, if they are all talking, it's very easy for that company to be able to interchange information between these different devices.

If, on the other hand, these devices are produced by different companies, then the challenges start increasing in terms of making

sure that all of the interoperability aspects are maintained because of different companies and that workstation is now becoming a tablet. People are walking around with a tablet instead of sitting at a workstation. So these are starting to get into the digital health world.

The interoperability standards are there. There's -- FDA has recognized 14 standards for interoperability. There is -- I'm sorry for the small print but the standards are there to be able to help with the guidance in terms of the standards that would help you evaluate this interoperability . Of course, the standards alone by themselves do not provide all of the information that you should be thinking about, so the standards alone will not answer all questions. And as the innovation increases, the challenges with making sure that all of the information in terms of HIPAA, in terms of safety, in terms of good effective information exchange is maintained is critical.

So in conclusion, ophthalmic digital health is

going to lead to many new innovative devices that will provide diagnostic and therapeutic healthcare. We hope today's workshop will foster this type of new innovation in the ophthalmic digital health. You're starting to see, I hope, a phrase that is used by many people in the digital health world. We believe that this digital health will be able to help the right cure for the right patient at the right time be an appropriate phrase. Thank you for your participation.

(Applause.)

DR. REPKA: Thank you, Dr. Schuchard. So our next speaker will be Ms. -- Dr. Krishna Yeshwant, who is a physician programmer and entrepreneur working with GV. Prior to Google, he worked on electronic data interchange. His background is both Stanford and Harvard. Good morning.

DR. YESHWANT: Hey guys. Look at this. All right. So thank you guys for letting me spend a little bit of time this morning talking about some of the things we're doing at Google Ventures. And I'll talk through some of things that I find

interesting, exciting about data systems more broadly and we'll find some particular applications in the ophthalmology universe.

Just by background, I'm a physician, as you kindly noted. I was a computer scientist before. Had started two software companies, sold one to HP, one to Symantec but am a primary care doctor over at Brigham and Women's now, and I spend most of my time looking at the intersection between technology and healthcare. At Google Ventures, we invest in companies across the space and so if I went through all of my disclosures, I'd spend the entire time here doing that, probably investor a few hundred companies across the space.

But I couldn't be more excited about the moment in time that we're at right now. We have Scott's, we have Jeff, we have Malvina, we have Bakul; it's just a tremendous array of insight and forward-looking thought processes around the regulatory sphere here. So I thought I'd just call out a couple of areas that I think are interesting and I'll actually get my timer here

started so I don't go too far over.

I know that there are some people talking later on around machine learning so I probably won't spend as much time on that today but I'll spend most of the time really on this area that we're seeing a lot of activity in called "real world evidence." And I know there's been a lot of interaction with the FDA over the last few years, really, but increasingly around how we can use real world evidence to help in various parts of the clinical universe. And then I'll kind of interweave some lessons from some of our experiences there.

So real world evidence, this is a graphic that many of you have probably seen before, if you can see it in the back, and perhaps some of you may even have helped create. So it's from a report that was released in 2012 from the Institute of Medicine around the learning health system. And I find -- being somebody who's interested at this intersection between clinical medicine and technology, I find this to be at least one of my,

you know, goals on life so to speak. This is I think at the dream state of what we can achieve in our healthcare system. And today I'll just walk through it briefly.

You know, at the top, you have kind of care delivery and, at least in today's world, thanks to a lot of progress that's happened because of work from the government, we have actually reasonably EMR penetration in the universe. So there is data produced just as clinicians see patients day to day. In an optimal world, that data would be fed into our scientific discovery process, the data coming out of clinical trials, and the process of discovering new therapeutics and diagnostics would be used to fluidly generate evidence that we could then use to rapidly expand clinical limitation of new opportunities and new things that we see

And, of course, also from that same report, we're kind of still in a world where a lot of data is kind of lost. In particular, data from the day-to-day work that we do as clinicians, in that

array, we see a variety of natural experiments that happen every day that we don't really take advantage of in today's universe. And that's for all sorts of reasons that I think most people in the room are probably familiar with. But that's perhaps what I find most interesting, exciting is other ways in whatever format, whatever vehicles possible to kind of close the loop around the subtle data, that we're producing every day that we're paying for every day as a society, as a system, that we don't really bring back to bear on other parts of the healthcare system.

So, to me at least, I think ophthalmology sits in a unique place in being able to reconnect and bring some of that data back to bear in many of the different pieces of the healthcare system. And I think that for a few different reasons. First off, it's unlike,- you know, my practice in primary care. Data in ophthalmology tend to be much more structured. I think part of it is that ophthalmologists are just very smart. They can organize their notes in this really organized way.

I'm always jealous when I see the referrals come in and it's all just, you know, this nice tightly compacted note but on the technical side, it allows for a computer to be able to pull the data out far more easily than what we see out of places like general internal medicine. Also, of course, as we've heard already, there's huge amounts of imaging data in ophthalmology and as we'll, I'm sure, talk about later and I'll certainly allude to, some of the techniques that we're seeing in machine learning, deep learning naturally allow for better analysis of these sorts of graphical and the sort of data that comes out and is fundamental to ophthalmology.

That naturally leads to, I think, another exciting area on pragmatic clinical trials so -- and I think this ties also back into the entire context of there's data being produced, can we use it and more sophisticated ways to give us a better sense of how new therapeutics and diagnostics might be useful across not just the group of people who we might study as we're looking for

efficacy but how it might actually work as we deploy these products and services into the rest of the clinical environment, which is obviously a lot more complicated.

There are a couple of use cases that just over the last few months that I think we've all seen come through the ophthalmology world where once again, I think if we'd had a unified dataset that is tied into the clinical universe could have been, I think, transformative in our thinking. And I kind of went through with that experiment when I heard each of these pieces of news, you know, so certainly the Lucentis versus Avastin, you know, sort of thought process; you know, each of these medications we'd -- I think are all alike, one had the suspicion that they might be similar in use, and it took a long time for us to be able to put together the clinical trial to raise the funds to get the coordination amongst the investigators to ask this questions. And with that experiment, you know, I often go through when we see these sorts of questions is, you know, if

we had this dataset, couldn't we have asked this question. It wouldn't have necessarily be the definitive answer but could we have gotten an earlier read as to whether there's a equivalents and where in our practice do we see other opportunities like this that we're not really able to take advantage of today because of the friction of putting these sorts of studies together.

On the opposite -- and I think -- and I just -- I go back to it and say, you know, I think ophthalmology is well-suited because there are large populations of patients out there we can ask these questions of. And again, day-to-day, we are seeing these patients in clinic, we're taking care of them, we're actually running these trials in kind of a natural experiment sort of way but we're not really yet tooled up as an infrastructure to be able to take advantage of all that.

A look in the opposite end of the spectrum is kind of small cohorts, which, to me, I think are particularly interesting as we're entering a world of gene therapy. You know, clearly, there are

exciting things afoot, even recently with Spark Therapeutics and a variety of other companies. I think we'll see a lot more in ophthalmology just given that the delivery modality into the eye being, you know, perhaps more clear as to how one might do that. But these tend to be small rare disease cohorts and, you know, how do we think about regulating these sorts of drugs when they're small patient populations and can we enable, you know, post-market approval surveillance. And I think once again, this sort of data infrastructure would enable that.

So going back to it, I think, you know, I find this to be kind of my guiding principle as we look at various opportunities in this space. And unlike many other specialty areas, I think ophthalmology is, again, particularly well-suited because there's been some work that the societies in this space have already done in starting to pull these pieces together. You know, in particular, there's the iris database that the American Academy of Ophthalmology has pulled

together. And it can clearly be applied in a lot of different areas. Whenever and however that happens, I think one of the learnings I've had from the technology universe is to rather than try to, you know, take the whole system on at once, you know, it's always better to try and start in one arena.

We spent a lot of time in a company called Flatiron Health, which is a company based in New York working in oncology looking at some of these real world evidence opportunities. And a couple of the learnings that we've seen there as it's applied to the regulatory framework is the deep need for clarity and transparency around what the data is and what the outcomes are. And that sounds obvious but in the end, when we're making these sorts of inclusions around any of these sorts of decisions, one wants to be able to go all the way back to the raw clinical data and any sort of transforms that are done on top of that to be able to understand what is it that's guiding some of this decision-making process. And so -- and

the infrastructure here needs to account for that.

There's a deep need for careful cohort selection. Thinking back to the -- some of the examples we were talking about earlier. It's one thing to compare it to drugs but you have to be very certain that you're talking to similar patient populations. And leading into -- kind of following out of that or as a corollary to that, there's a real necessity for a pre-specified analytic plan. As a computational person, you can sometimes fall into the trap of asking a whole bunch of "what ifs" just to get to the answer that you want and not necessarily to the ground truth, and I think we've all seen some of the flaw with that.

And then, of course, culture incentives are crucial in this whole thing.

I'll spend a couple of minutes just talking about machine learning and I know that there are a few people talking about it in the -- you know, in the subsequent sessions but I guess the the main point I want to make just as someone speaking from

the perspective of Google is that it's been a transformative set of things that have happened over the last few years, and it's certainly been transformative internal to our company. And I think that we're starting to see how some of these areas are affecting healthcare.

You know, the broad painting of it here is that, you know, we historically used "if then" statements to write a lot of our software and over the last few decades, we have seen the opportunities to actually write software that figures out based on exposure to datasets how to classify different inputs.

And fundamentally, the structure of what's happening in deep learning, which is kind of this term that I think we all hear a lot, is not that different than what's happened before, you know, is fundamentally the same sort of neural network architecture that's possible today that was possible 30 years ago.

But what's different is that the infrastructures that we're running these sorts of

analyses on are far larger and far more robust. So when I was a computer science student, we maybe were able to make neural networks that were, you know, single digit sort of layers, so 5, 6, 7, you know, layers in depth. And then you kind of run up against the limitations of what was possible and today we're seeing far deeper sort of structures which will offer far more sophisticated classifications. And clearly, we're seeing that touch on some of the areas. I think we've seen the paper in *JAMA* around diabetic retinopathy project.

But I kind of wanted to call out this interesting effect that's kind of well-known in the artificial intelligence and machine learning world called "the AI effect." And it's a little tongue-in-cheek but I figured I'd bring it up here because we talk about it a lot inside of Google. And it's one of these funny things where, you know, there's this quote, "intelligence is whatever machines haven't done yet." And, you know, when we see kind of some of these sorts of

advances, it's sometimes easy to kind of look at them and say, well, like the world's going to move forward in this particular way and machines are going to take over and all this stuff, and we hear various venture capitalists saying that we don't need doctors anymore and, to me, I think that's -- you know, that's certainly one of the ways that people can characterize it.

But I think what tends to happen far more often in these fields --and there are several examples, you know, that we can go through, but usually what happens when an AI application works is it kind of gets subsumed into the field that it's working in. You know, so to that extent, I think as I look at what's happening in the world of machine learning meets ophthalmology in particular, you know, nothing but excited because I think ultimately, it will mean that we're able to do better diagnoses for our patients and get to the sorts of care that they need. And I think the way the world will look at it ultimately is not machine learning taking over some of these areas

but rather these areas finally kind of getting some of these benefits of computer science applied in these areas and getting to some of those efficiencies.

So just in the interest of time, I'll close the comments there but I'll be around the conference very excited to engage with this audience and thank you very much.

(Applause.)

DR. REPKA: Thanks for those comments. A couple announcements -- just they were asked. In the interest of interconnectivity, there is a wifi pass code for this room. It's an upper case "W-A-S-R-V." Don't know the source of that but that's great. And the slides will be available to attendees subsequent to this meeting.

Our next speaker -- and thank you, Dr. Paul Lee for joining us at the podium. He's the Professor and Chair of the Department of Ophthalmology and Visual Sciences at the University of Michigan and has to direct the WK Kellogg Eye Center and all of its people. Thanks,

Paul, for joining us.

DR. LEE: Thank you and the organizers for the opportunity to be here. I was asked to speak about an introduction to the area of telemedicine in ophthalmology and we have a terrific program that's put together today.

In terms of the rationale for why teleophthalmology is so prevalent and so important today is that we are in the midst, as our last speaker talked about, about a transformation n health and healthcare.

And so you can see the pressures on the left that are forcing us to look at the changes and all the new attributes on the right that, as Eric Topol put very nicely, is leading to a new way of looking at medicine and healthcare. And our earlier speakers have already talked about a key part of this movement is to take what we do out of our traditional offices and clinics and de-marketizing it and moving it into the hands of patients where they live as well as other distributive networks.

In terms of the level of usage that we have right now, Kaiser is a leader in the implementation of e-health. There are many others but Kaiser last year had over 50 percent of their patient interactions done through remote mechanisms. And so if you look at some of the things that they were doing in ophthalmology; for example, they're using glaucoma suspects being followed only by OCTs after initial examination. There are a lot I interesting things going on out there right now.

And the other piece is all the experts we talked to tell us that at least 25 percent of our visits that we do in the office today, within the next few years, will be done by e-health or remote mechanisms.

Patients are interested. That's helping driving the market and why this is growing and there are a lot of folks that look at how we can do this. And we all recognize that there are various ways we can interact, storing forward, live motion. There's going to be some good talk

about that today.

Across the disease areas we have in ophthalmology, there are a lot of different use cases across a lot of different diseases. Pretty much everything that's out there that we do is being investigated and in many aspects, especially in the back of the eye, there's solid evidence for why it works well.

So in today's presentations, you've already heard from our colleagues at the FDA and at Google about all the different things that we need to look at as we're interested in moving these into the hands of real patients. Going forward, we've got some great talks about some key examples where there's rich data about what we do and the issues related to understanding their usage. And you've heard about the deep analytics and the deep learning.

And the meat of today's presentations, obviously, are going to be the panel discussions. The panel discussions will focus on digital health devices as an aid for diagnosis, safety and

effectiveness and the risk mitigation strategies that are out there.

Just a few additional thoughts for us as we move forward in today's agenda; a really important question is "what is the gold standard?" Is it what a group of doctors or an individual doctor thinks? Is it a reading center or is it a machine? If we go back to the original ETDRS papers, there was that beautiful grid or table that had the reading center on one axis, the physicians on the other axis, and the diagonal fortunately matched very nicely. But there were differences and discrepancies across there. And so what's truth? What are we going to use to say this is accurate and this is the way we should go? Is it the machine learning? Is it the physician? That's something that we have yet to resolve.

A second along the lines of validity is how well does this device or software perform relative to whatever we determine to be the gold standard. And then the reliability piece is very important. If we do repeated measurements, do we get the same

result? Within an image, if it's software, do we get the same analysis at different spots with the same characteristics? And also, if it's appropriate, do we get the same result if different people use the equipment? Now this is something that we apply to all the new devices but have we applied it to clinical care as we understand it today?

So let's look at some of that information that we have in the literature because I think this will help frame the standards and the context by which we evaluate the new technologies.

So the ATA has some nice guidance on diabetic retinopathy in terms of the relative ways we can look at standard comparison, but in real -- the real world -- this study is almost 25 years old. It looked at how ophthalmologists compared live examinations for patients with diabetic retinopathy compared to photographs, single site study but you can see that the performance specifications of various technology elements can rival that of ophthalmologists and that there are,

were and potentially still are opportunities for improvement.

This is a paper from the Oats Reading Center (ph) out of Miami that reviewed the literature for the simple vertical cup to disc ratio performance in the literature of ophthalmologists relative to other means of analysis.

And this is a meta analysis just published recently looking at teleglaucoma and looking at the sensitivity and specificity of performance in the literature and those studies that compared teleophthalmology, teleglaucoma to in-person examinations.

In terms of the implementation issues, a key factor is to recall where in the care spectrum are we using this. Is it a new patient or established patient and what's the level of autonomy we expect the system to be able to deliver.

Patients do want to use this. There's good data from the public opinion polling that suggests that patients are very receptive to using these. And the National Academy of Medicine has made it

clear that the communications aspect of what we find is as important as the diagnostic accuracy for what's considered diagnostic error. We know that the implementation of diabetic retinopathy programs -- this one in LA -- greatly increases the screening rates of use of retina, but there's still a problem in terms of getting people into care even after they've been screened. So to close the loop, we have to keep that in mind.

And, of course, there are various legal issues and payment coverage issues, so some questions on the legal liability side on the left for providers and physicians relative to malpractice coverage, to actually use it. On the system side, we were having dinner last night, conversation with Mike Change and Mike Abramoff about if there's an error in a system, who's responsible; is it the system; is it the physician, and it probably varies based on the purpose of the system. Is it an aid in which case the physician's probably going to be responsible. Or is it meant to substitute for a physician in which case it's probably the system.

And then there are a whole host of issues we run into that the FDA Centers have a lot of precedence on in terms of radiological monitoring devices and teleradiology in terms of display characteristics.

Lots of different state laws that need to be navigated relative to the actual implementation as well as reimbursement issues that are for another time.

And just a couple of final thoughts. The first is what's the implication of all this? We saw the reference to the Institute of Medicine National Academies report on a learning health system. As we look at the impact of all of this technology on how we interact with patients, we'll be able to do a lot better. But the essence so far that the system can't quite replace is that human interaction. And so in a way, this has a promise of restoring traditional physician-patient functionality that current regulatory and work pressures keep us from doing as well as we want.

The second you've seen is that there are a lot

of new entrants and so the world's going to be a really exciting place in the next five years. I see many of our pioneers out here in the audience and participating in a panel, and so I'm looking forward to a terrific day. Thank you very much.

(Applause.)

DR. REPKA: Thank you, Dr. Lee. So our next speaker will be Dr. Paul Chan, Professor of Ophthalmology and Visual Sciences at Illinois Eye and Ear Infirmary at UIC and Vice Chair for Global Ophthalmology there with a great deal of interest in telemedicine for ROP. So, Paul?

DR. CHAN: Great. Thanks, Mike and thank you to the organizers for having me speak here today. Wonderful series of talks which I think leads into what I'm going to talk about, which is where did we go wrong, right. So in terms of diagnostic accuracy and things that we don't necessarily do well, what are we having trouble with and also, how do we do better and what are strategies to make that improved.

Here are my financial disclosures. I am a

consultant for Visionex Medical Systems, which does make some of these ophthalmic cameras. And I'd first like to acknowledge the collaborators that I work with, especially Mike Chang, the i-ROP Group and the Gen-Rop Group that focuses on education for ROP. So what do we know? I think historically, there are a lot of retrospective and prospective studies looking at whether or not telemedicine and image-based diagnosis for ROP works well. And I think that we've shown that its' very good for identifying something called referral warranted ROP. We've shown that it's reliable, accurate, cost effect. In terms of physician time, it's definitely more time efficient.

And there are a lot of active clinical ROP programs outside of the context of the study. For example, SUNDROP -- Darius is here -- is going to be part of the panel with Mike Trese and the focus ROP and also, in a development world, which is a specific interest of mine, many, many ROP programs actively in use and they design their own

telemedicine reading centers, so it's going on in the real world.

What I'm going to talk about mostly is, well, okay, it's going on but who's really qualified to do these image readings. And this is sort of a personal experience but, you know, are we good enough? And as I mentioned before, a lot of the systems that we look at today, a lot of the programs that we're focusing on focus on this definition of referral warranted ROP, which basically is what we call "type 2 ROP," so things that need to be referred immediately, something that may progress to treatment sooner than later and needs to be examined very quickly by an ophthalmologist.

One of the potential diagnostic challenges -- I'm not really going to go into image quality, field of view, or go too much into the hardware issue; going to go mostly into how do physicians perform in terms of making the correct diagnosis. We know that experience matters. We also know that potentially, experts have some challenges in

identifying pluses use. That's been well-documented. We know that there is a lot of controversy and a lot of variability in how expert graders examine pluses use and make that diagnosis.

And in terms of training, so how do we certify graders. That's been a particular interest of mine and there's a lot of variability in, you know, I call it sort of telecertification or how do we certify people to actually read telemedicine images.

Many years ago, we started doing studies looking at whether or not board eligible ophthalmologists so these are really general ophthalmologists who finish their residency training, they're going into fellowship; how do they do in terms of compared to an expert grader reading a telemedicine image? They don't do that well, okay, so what's interesting is that they actually misdiagnose type 2 ROP more frequently than not and these are retina fellows who are going into retina practice and when they're done.

And we also looked at pediatric ophthalmology fellows.

Again, in the general community, sometimes general ophthalmologists are screening for ROP so this has some relevance. And in pediatric ophthalmology fellowships in the community, a lot of times the pediatric ophthalmologist is doing the screening. They're making a diagnosis and then they're calling the retina specialist to do the treatment. Pediatric ophthalmology fellows also don't do very well, right, so the type 2 ROP and also even treatment of treatment required ROP were challenges.

So what do we do? Well, we found that there were issues so we created a tele-education system using the system that Michael has with the i-ROP system. And what -- we recruited about 250 ophthalmology trainees from around the world, U.S. and international. And it wasn't just to education them to make them better. We also used the system to evaluate their performance and see what they were doing wrong or incorrectly. And we

presented them with RP case presentations, very much in the same way that someone would read a telemedicine image; give a case, read the image, how do you do, what's your answer.

What did we find? Well, similarly, there are struggles with type 2 ROP, so this referral warranted ROP that we focus on, right; this critical period, this threshold and they're not diagnosing this correctly. Almost 50 percent of the time, they're incorrectly diagnosing this. And why? Well, there are struggles in zone of disease in the diagnosis and there are struggles in terms of pluses used diagnosis. So these are specific categories that they're struggling with that make this difficult to do well.

In the international arena, we talked about global. This isn't just about domestic policy. We have to look at the international role and I do a lot of this. And they're also finding similar error rates, right. So type-2 ROP is difficult to examine if you're not experienced. So looking at inadequacies in diagnostic accuracy for ROP, the

U.S. international cohort, all of them across the board find difficulties. They misdiagnosis almost 50 percent of the time.

What does this mean? Well, we need to improve our diagnostic accuracy. We have to find ways to improve the ability to diagnose referral warranted disease to get these kids to an ophthalmologist, to get get kids examined so they don't go blind. And in terms of just the real world, why is this important? Well, general ophthalmologists, you know, we talk about non-physician readers and so forth and so on, well, are they good enough, right? Well, we have over 250 physicians, ophthalmologists in training who just didn't seem to do very well. Okay. Then that's problematic if we're going to implement these systems in the real world.

So how do we do better? Well, let's first go to experts and what they might do well with. And when looking at the clinical diagnosis and the image-based diagnosis, we started doing some exams about stage four, retinal detachment. What we

found is that retinal detachment seen on a two-dimensional image may be difficult to pick up even among expert graders compared to indirect ophthalmoscopy.

Here's an example. Well, so this patient here, you can see the subtle changes here, some traction and some elevation of sub-retinal fluid. Examiner one, stage 4-a, so diagnosed this correctly, but examiner two actually said that no treatment was required. And if that were the case in a telemedicine scenario, this child may have gone blind and didn't receive a vitrectomy.

What about aggressive posterior ROP; we published some data looking at this and there are some difficulties in identifying this type of disease. Now we could say that ancillary images in other modalities can help improve diagnosis for AP-ROP or other conditions that may be subtle but again, there are certain things that we may not be doing very well.

How can we do better? Okay. So we mentioned the tele-education system. How do we certify

readers. This potentially can improve diagnostic accuracy, definitely has implication for training and we think that it has implications for ROP telemedicine as well. Can we add ancillary imaging, so fluorescein angiography, OCT angiography, OCTs into our paradigm and our algorithm? Can that pick up retinal detachment and subtle changes that improve diagnostic accuracy?

We've show that, actually, FA improves diagnostic accuracy for identifying this referral warranted disease. We've shown that digital mosaic images may improve inter-grader liability and agreeing among graders, which is important, and also improve diagnostic accuracy for certain conditions.

Here's just the data showing that the tele-education system improves performance from U.S. and international trainees for every category of disease.

And in summary, we have challenges, right. So I say this all the time to people who say that

they want to set up a telemedicine system. Well, it's just about reading an image, right. There are all sorts of logistical issues. It's challenging in terms of diagnostic accuracy. You have to be good at this, right. There's a certain level of quality that we have to look at. How do we do better? Tele-education, standardized certification programs to certify readers, investing in potentially improved imaging in multi-modal imaging, and computer-based image analysis and deep learning that was mentioned. And I think that this is really exciting in terms of the future

Now what it comes down to is who should be responsible for ROP telemedicine programs. Now, you know, I sort of a little opinion with some data but in my opinion, I think we're still at a point where skilled ophthalmologists should and need to be responsible for the oversight and the diagnosis and management for decisions regarding ROP care and telemedicine. Thank you.

(Applause.)

DR. REPKA: Thank you, Dr. Chan. So our next speaker is my colleague, Dr. Ingrid Zimmer-Galler from the Wilmer Institute who's going to speak about the diagnostic challenges for diabetic retinopathy.

DR. ZIMMER-GALLER: Thank you very much for allowing me to be a part of this very exciting day today. I do not have an financial disclosures to report.

So just to clarify, telemedicine is certainly used in a number of different ways for diabetic retinopathy screening, but we are talking about here is doing the telemedicine screening in the primary care setting. Remember that the big problem we have with diabetic retinopathy is that so many patients with diabetes do not have their recommended eye evaluation but they do go to see their primary care physician. So this is the perfect place where we can capture patients that are not compliant with the recommendations for a diabetic retinopathy evaluation. Traditionally, this is done with a nonmidriatic fundus camera but

it can be certainly done with a number of different imaging devices.

Traditionally, those images are transmitted to a remote reading center and at least in the United States, typically, the images are then reviewed by a licensed eyecare provider and a report is sent back to the primary care physician generally within one to three business days including whether or not referral to an ophthalmologist or to a retina specialist is warranted for further evaluation. What's exciting is the possibility of using automated image analysis to do this image reading at the point of care at the time when the patient is actually in the primary care physician's office. The algorithms will allow that report to immediately come out and the patient will know whether a referral is needed before they leave the primary care physician.

Some of the diagnostic challenges with telemedicine diabetic retinopathy surveillance that I'm going to touch on include ungradable images, diabetic macular edema, the use of wide

field imaging, and then the concept of other pathology.

So what is an ungradable image? Would this be considered an ungradable image? Or what about this image? Would you call this an ungradable image or is this considered advance diabetic retinopathy with a vitreous hemorrhage. Or would you say "does it really matter if an ungradable image results in a referral as well. So there are a lot of questions that come about with ungradable images. Certainly, image quality depends on a lot of factors, many of which you really can't control, some that you can control include the imaging device. The resolution of most fundus cameras that are available today really is adequate to pick up even the tiniest micro aneurysms, but the field of view comes into play as well. And then we will hear more and more about the use of various handheld and Smartphone adapters to allow fundus imaging with handheld devices.

Keep in mind one of the problems with

Smartphones is that I'm not sure we really have good enough validation yet and also, an issue is that Smartphone platforms keep being upgraded and by the time you have a validation study done for one Smartphone platform, that Smartphone may be one or two generations further along. The acquisition procedures play a role here, too. This is mydriasis dilating the pupils or not dilating the pupils, the number of images, the number of fields that are obtained. The operator experience clearly still makes a difference, someone who is well-trained on an imaging device is going to more consistently get good images than someone who does this once in a while.

And then, of course, patient variables come into play as well; the age of the patient, whether or not they have media opacities, whether or not they are able to be positioned adequately at the imaging device.

So ungradable images really need to be discussed in the context of validation, and Dr. Lee already mentioned the American Telemedicine

Association has different categories of validation. The American Academy of Ophthalmology also stresses the importance of validation.

When we talk about traditional telemedicine diabetic retinopathy systems, the reference standard that we are comparing that telemedicine system to, the "gold standard" is considered ETDRS 7-field stereo photographs. When we talk about validating automated systems, what we are talking about is looking at large datasets that have been annotated, that have been looked at by experts or groups of experts and you're comparing how the machine is reading that to the group of experts.

The problem with have with validation studies is that the validation really needs to be targeted to the clinical outcome that the program is trying to achieve. So the targeted outcome may be presence of absence of any diabetic retinopathy; it may be presence or absence of vision-threatening diabetic retinopathy; it may be presence or absence of specific diabetic lesions so you can't really compare the validation studies

across the board. The measures that we use for validation, of course, include sensitivity, specificity, false-positives, false-negatives and positive and negative predictive values. We need good sensitivity so that we can make sure we pick up all the disease, that we don't miss someone who has significant disease, but we also want high specificity because we want to limit the number of patients that are referred who don't actually need to be referred. This, of course, will increase the cost of the whole screening process and it'll decrease the efficiency.

We draw a lot of information from our colleagues in the United Kingdom who have done an admirable job of setting up a national telemedicine diabetic retinopathy screening program and together with traditional examinations, they have now screened more than 90 percent of their patients with diabetes, and they have been able to, for the first time in five decades, show that diabetic retinopathy is not longer the leading cause of vision loss in the UK

among working-age adults. They first came up with the numbers of a target sensitivity of 80 percent and a specificity of 90 percent, and those are numbers that are often tossed out but we don't really even know if these are the best target number that we should be using.

Coming back to the ungradable images, a validation really is not useful if we don't include the ungradable images. For the traditional telemedicine systems, a lot of the validation studies did not include ungradable images. Having the ungradable images included will certainly likely result in a change in the specificity because you're probably referring patients that don't necessarily need to be referred.

Again, drawing upon the experience from the UK, they give a target ungradable rate of five percent. That's a pretty specific -- that's a pretty high target to have. And again, it will depend very much on what the outcomes are that the particular telemedicine screening program is

looking for. So again, this is not necessarily the best number to be using.

With automated systems, one of the things that we can do is adjust the target, the set -- you can have a set point at a different level so that you can minimize the false-negatives but also have a manageable level of false-positives.

In the interest of time, we can't really talk about QA, about quality assurance but it's very important to keep in mind that the relevance of a program's validation really can only -- you can only keep that relevance if you have a robust QA program in place as well.

A couple of quick words on diabetic macular edema. So clinically, this requires identification of retinal thickening, and this, of course, can be done with stereo viewing or with OCT. And most diabetic retinopathy screening programs don't include stereo images or, obviously, OCT. And without an assessment of retinal thickening, we are traditionally using

surrogate markers, so we use hard exudates, micro aneurysms, and hemorrhages in the macula as surrogate markers. But that clearly doesn't identify the extent of the macular edema and you can have surrogate markers present even in the absence of macular edema. So this is an area where we still have work that needs to be done.

I also want to just point out some of the new information that we have with ultra-wide field imaging. This, of course, gives us a much larger of the retina that is imaged. And studies have shown that for telemedicine purposes, this can significantly reduce the ungradable image rate and it also reduces the imaging time but keep in mind that these imaging devices are very large and usually, they're too expensive to be used in a screening environment in every primary care setting.

What's interesting, though, is that wide field imaging has been shown to, in approximately 10 percent of cases, result in a more severe -- a higher level of severity of diabetic retinopathy

compared to that same image if you look at only the ETDRS field of view. And this, of course, brings up the question again, what is the gold standard; what is the reference standard if we have information that potentially gives us more information than the reference standard does. And also, for diabetic retinopathy, that brings into question how is this relevant with all of our clinical trials that are based on the ETDRS photographs.

Other posterior segment pathology also needs to be considered. This is where perhaps there is a greater question that comes up with a human versus a machine interpreting the images. If there are other abnormalities on that image, how are we looking at that with the machine or do we even need to be concerned about that if the purpose of the imaging is specifically for diabetic retinopathy.

And then I'm just going to end with this question that isn't so much a diagnostic challenge but what about the culture change; what will it

take; will physicians and will patients accept what a black box spits out and says is the result; will they accept that result. So thank you very much.

(Applause.)

DR. REPKA: Thank you, Dr. Zimmer-Galler. Our next speaker, Dr. Michael Chiang, is Professor of Ophthalmology and Medical Informatics at the Organ Health and Science University in Portland, who's been active in ROP and comes to speak today about advanced analytics in ophthalmology or, I guess, how to get the doctors to trust the box.

DR. CHIANG: Okay, Mike. Thanks. So I'm going to focus on this interface between clinical diagnosis and analytics and artificial intelligence. So I -- a couple of disclosures here; one of them is that I manage a group called Imaging and Informatics in ROP and we get some funding from NIH and a staff and have a couple of financial relationships here. But -- and I also have a couple of relationships through AOO. I'm on the iris registry executive committee. I'm on

the board of trustees and manage the -- a data analytics committee but I'm not speaking here on behalf of the AOO. But I think it's a relevant disclosure.

I want to highlight that I've worked with Paul Chan and a couple of others here for years on this and'll be presenting some data from those projects.

So the disease I will focus on here is retinopathy of prematurity, and, you know, one of the reasons I'm talking about that is because it's the work that I do and it's the work that I'm most familiar with and several people in this room have done a lot of work in this area. But more importantly, I think that I'm going to try to highlight some generalizable principles out of this work that I think are going to be relevant for this topic. Okay. So that's what I hope we can focus on, sort of the generalizable principles that come out of this data.

So the topic here is going to be "gold standards." Paul Lee and several others alluded

to this, like how do we classify a disease in ROP and diabetic retinopathy but really not a whole lot of other diseases in ophthalmology. There are very clear classification standards that have been develop din the case of ROP over 30 years ago with what's called the international classification of disease, a standard terminology that, really, everybody else in the world uses. And it happens for ROP that these terms are things like the zone, stage, clock hour extent and something called "plus disease."

And so because of these standards, we can do clinical trials. And because of the clinical trials, we know that presence of something that's called "plus disease" is the most critical thing that determines whether or not a baby needs to be treated. So if you have plus disease, you are at risk for going blind; you need treatment. Okay. So we really need to be good about identifying plus disease in ROP.

So what's plus disease? It means that you've got tortuous arteries and dilated veins in the

posterior pole of the retina. Okay. so remember those terms tortuous arteries, dilated veins in the posterior pole. Okay. and if you've got that, that's bad. So one of the problems is that we're not very good at identifying plus disease.

Okay. About 10 years ago, we worked on a project where we presented the same images to experts around the world. Now these are not trainees These are legitimate world experts who've led clinical trials in the area. And so here's an example where there's a little bit of tortuosity, a little bit of dilation in the retina. And 15 percent of experts called this "plus disease," 85 percent called this "not plus." And the image on the right side, it's split 50/50, half called it "plus," half called it "not plus." And so we've got a situation where, you know, the world -- you know, this is so important that it determines whether or not you to treat a baby, yet the world's experts are splitting 50/50 or 60/40. So intuitively, that's not good.

And so I want to talk a little bit about the

science and the art as medicine, because I think that's going to be relevant to this panel here. So seven or eight years ago, I'm on a panel about ROP and one of the experts on the panel used the analogy that there was a U.S. Supreme Court Justice, Potter Stewart, in the 1980's and so the analogy was that plus disease is like what Potter Stewart -- how Potter Stewart described pornography: You can't define it but you know it when you see it, because it just looks bad. Okay. And, you know, that's what I would call the art of medicine, clinical judgment. And yet that comment bothered me for a few months because I thought if we want to be scientific about it, how can we just be saying that things look bad, okay, but you just get a gestalt about it.

And so the thing that it made, really, us wonder is that -- you know, we see this all the time in clinical medicine, that the experienced doctor will say, I just don't like the way this looks. And so we got interested in this fact. Well, if everybody is looking at different things,

could that explain some of the variability that we're seeing and are these definitions that we come up with post hoc; arterial tortuosity, venous dilation in the posterior pole, is that an over simplification?

Okay. And so what we did is we got who we considered the seven most prominent experts in ROP diagnosis in the world. They were people who, in many cases, had come up with the original definition of plus disease and came up with that original classification scheme and, you know, we got them individually into a room and we collected standardized images and we had them, you know, sit there and you diagnose this image, you annotate the images, we'll videotape when you do it, take us through your thought process.

Okay. And so there were house of videotape record, you know, hundreds of transcript pages here. And so we analyzed them with a cognitive psychologist. And it turns out that there's a disagreement in the process of diagnosis, that you've got one image diagnosed with expert number

one as plus; expert number two as pre-plus, okay, an intermediate state; expert number three is normal; okay, same image and they're all looking at different parts of the retina when they make a diagnosis.

Okay. So not only is the diagnosis different but the process of diagnosis is different. And in fact, if you go through and analyze those hours of transcript, it turns out that it's not just those three terms, arterial tortuosity, venous dilation in the posterior pole that they're looking at, it's all sorts of different stuff. Okay. So these terms that we use in ophthalmology are, in a lot of ways, oversimplifications. And so I think that Krishna made a really good point about the ophthalmic exam being structured, you know. But opinion the other hand, what we do as ophthalmologists is we look at images and the counterpoint is that those images are inherently unstructured. And with these classifications (inaudible), we try to create structure out of that but it's not perfect. Okay. In fact, in a

lot of ways, its' far from perfect.

And that's one of the reasons that, you know, we've gotten interested in things like computer-aided diagnosis. Can you, you know, use machines to, you know, try to quantify these areas and make it more objective and quantitative instead of subject? And so, yeah, we've done some work in this area. There have been a couple others, some in this room, like Mike Trese has done some beautiful work in this area with a guy, David Wallace at Duke University.

And, you know, for our team, the data that I'll be talking about represents the work of, you know, us together with Paul and several computer scientists, Jayashree Kalpathy-Cramer from Harvard, (Inaudible) and Deniz Erdogmus from Northeastern University and we've got a team with two post docs, four PhD students and two master's students who've worked on this for about six years.

But anyway, we found -- you know, come up with -- you know, we've looked at two different

approaches and the ones that Paul and, you know, others have talked about; you know, number one's a classic machine learning approach, and number two is a deep learning approach with convolution on neural networks.

And, you know, Paul mentioned the idea of reference standard and I think that's a huge challenge. And the way that we've dealt with reference standards where is that we've captured -- a clinical exam did; in other words, what did the real ophthalmologist diagnose at the bedside. We've taken photographs of every retina and we've had a series of several experts look at each photographs and come up with consensus reference standards that blend, in this case, four different evaluations into a consensus reference standard. Okay. So that's how we evaluate these systems.

And I just want to present some data about what we're -- you know, what we're getting here because again, I think the generalizable thing is, you know, what's our concept of how you evaluate these systems and, you know, how you validate

them. And Paul, we used terminology. So this happens to be a system that classifies images using machine learning approaches. And so there are 73 images and we're comparing diagnostic accuracy of how well do you classify plus versus pre-plus versus normal, okay, compared to that reference standard diagnosis. And you've got eight experts and a computer system and the eight experts here are between 79 and 99 percent accurate; on the average, 87 percent accurate. And the computer system is 95 percent accurate, okay, for classifying plus versus pre-plus versus normal.

Okay. Second approach here is a deep learning approach and, you know, this has gotten a lot of press recently. In our case, we've trained a convolution neural network on a series of about 6,000 RP images. Again, every image has a reference standard diagnosis. Okay, so very painful to come up with that for 6,000 images in a consensus way, but "a" under the arc, "c" curves her about***98 for diagnosing plus disease. And

so really, really high. If you divided them on independent data sets, the system outperforms most experts. Okay. So in this case, 91 percent accurate compared to between 77 and 94 percent accurate.

And I just want to make one point about this black box concept because I think Ingrid made a really good point about, you know, do people trust these systems as black boxes. You know, one of the things that we've looked at is occlusion analysis; in other words, you feed thee systems into these deep neural networks and based on what part of the image you don't feed into the network, it can highlight areas here shown in purple that the machine thinks where most important for clinical diagnosis. So in other words, it's a process of working backwards. Okay. What can the machine tell us what the doctor might have been thinking; you know, because if you take that piece of information out, the diagnosis changes. Okay. So I do think that there's potential for these systems to work backwards and tell us what we were

thinking in a way that doctors are actually not always able to articulate, because we've done the cognitive psychology studies. So maybe some potential for that.

And I want to close just with a couple examples looking at variability because again, I think this is going to be generalizable. So here what we've got is data, in this case, from eight different experts looking at 100 different images. Okay. SO here's one, two, three, four, all the way up to 100 and if the box is "red," that expert diagnosed it as plus disease; if it's "yellow," that expert diagnosed that image as pre-plus; if it's "green," that expert diagnosed it as normal. Okay. So point number one is that experts seven and eight diagnosed plus disease six times more frequently than expert number one. Okay. So that's not good. And if you go to a different dataset, it's that same six to one ratio. Okay. So this phenomenon of under-callers and over-callers is a real thing. Okay. We all know this clinically but, you know, I think this presents it

graphically.

The second thing is that for every image, if you give it a score, one point for a "green," two for "yellow," three for "red," and if you average that score for each image, convert it to a color, you've got a continuous spectrum. You've got the very abnormal over here and the very normal and then every color in between. Okay. So that's that graphical represent -- what we do in ophthalmology is a continuous spectrum and what we do when we treat disease and diagnose it is that we draw those lines; okay, are you plus or pre-plus; are you pre-plus or normal.

And as ophthalmologists, we've got data that I haven't shown here that ophthalmologists are very good at comparing; okay, what's better, number one or number two, you know, very consistent. But we are not consistent at drawing those lines and I think that's a huge problem. And I think that's where these systems can really help us make better decisions.

And so, yeah, we've done some work here

choosing sets of standardized images where this is very, very abnormal. And how do we know it's abnormal? Eight experts called it plus, nobody called it pre-plus, nobody called it normal. And this one's very, very normal. Everybody called it normal, nobody called it plus, nobody called it pre-plus, everything in between. And in fact, if you feed these images into that computer-based system and give it a score, it falls on a straight line. Okay. So I think, again, that's where computer diagnosis can really help us as clinicians.

And so my -- this is my last slide and these are some points that I think are useful for discussion later Number one is that I think that ophthalmic diagnosis is innately subjective and qualitative and, you know, we've seen that in diabetes, you know, with Ingrid's story and Paul Lee's story. We see it ROP and they're significant inconsistence, even among experts, in terms of drawing these lines. And my suspicion is that in the real world, they variability is even

more than what we're seeing here. I do think there's a role for expert systems and improving that consistency. I think the bar for these systems should be that they're human-like and that they're not going to be perfect but, you know, they should be as good as humans.

And I do think that validation requires transparency. I don't think it's enough to use a single expert as a reference standard like Paul was saying. And, you know, we've tried to use consensus panels, maybe there's a better approach. I think this is a rapidly changing field and, you know, this point was made before. You know, these systems inherently learn from their mistakes, you know, with this concept of the learning health system.

And so I hope that in coming up with these rules, you know, we can take that into account where, you know, the cycle time for updating these systems, you know, whatever we can do to try to decrease that I think is going to help the field. And I do think that the intended use of these

systems matters. You know, do they a), give advice to physicians in a decision support manner or b), are they closed-loop systems, you know, like (inaudible) for primary care where there's no ophthalmologist involved? And I hope that the FDA's going to consider variable levels of regulation based on the intended use. So thank you very much.

(Applause.)

DR. REPKA: Thank you, Dr. Chiang. Our next speaker is Dr. Linda Zangwill, Professor of Ophthalmology at UC San Diego and serves as Director of Clinical Research in the Glaucoma Center and Director of Imaging Data Evaluation and Analysis. Good morning.

DR. ZANGWILL: I want to thank the organizers for inviting me here today and I want to acknowledge my financial disclosures. And I'll be talking about machine learning in general and, obviously, the applications in ophthalmic diagnostics.

Machine learning, obviously, is changing our

lives on a daily basis with recommendation engineers, with autonomous driving, and we've heard a lot about it in ophthalmology. And there's terminology that we've already heard about but I just want to emphasize the difference between machine learning and deep learning where the deep learning can really -- the instrument, the machine, the algorithm learns from deep layers and sees the patterns within the layers.

There are different types of machine learning tasks, most of what we heard about is supervised learning where you have data and a label. In this case, for example, glaucoma or not from visual fields, the processor looks at the data and the label. You have an outcome, glaucoma or not, and the accuracy is compared to the expert or the label data. Unsupervised learning, you just have data and the machine looks at that data and sees patterns; in this case example, visual field patterns. Sometimes these patterns are very similar that a clinician might identify as a nasal step or paracentral scotoma and sometimes they are

not.

In terms of machine learning applications, we've heard a lot about today software as a medical device and there are different categories of software as a medical device. And machine learning applications are relevant for informing clinical management, driving clinical management, and treating or making a diagnosis or referral.

So there's a long history of machine learning in ophthalmology and it started, really, with the supervised learning and the most applications have been in retinal disease and in glaucoma. Here's an example from my colleagues, Mike Goldbaum in the early 1990s, from UCSD before I arrived, looking at visual fields, and the conclusion was a neural network can be taught to be as proficient as a trained reader interpreting visual fields for glaucoma. So that was, you know, many years ago, over 30 years ago almost.

There's lots of work in this area looking at visual fields, looking at fundus photographs, really early with the machine learning, with the

neural networks, for glaucoma damage and detection and progression. Similarly, with retinal disease, detection of retinal lesions started very early. Later came detection of diabetic retinopathy. There are numerous challenges to the community for automated detection. And there's -- in diabetic retinopathy, we've heard that's really the most mature, I'd say, in the ophthalmic diagnostics. And once again, there are differences between strategies that wanted a design to detect micro aneurysms, hemorrhages -- this is an excellent review article from 2013 -- and also detection of diabetic retinopathy.

Also, there's, as I mentioned, unsupervised learning in ophthalmic diagnostics. Here are some examples where -- mostly in visual fields in glaucoma, some of our work and others where you put the visual fields points in the machine learning algorithm and these patterns are quite remarkably like some of the patterns the clinician identifies and others are not. And we can even see the progression of these patterns that really,

the diagnostic accuracy is similar to more standard progression algorithms that we are using.

There's been tremendous progress in the last three to five years due to deep learning, due to the computational resources that are now available and also due to available data sets for these algorithms. This is a slide about deep learning and health informatics, the tremendous growth of published articles through 2015. I think if you went to 2017, they graphs would be off the chart but look at where imaging is in here. And in ophthalmology, we use imaging on a day-to-day basis.

So can -- deep learning can be supervised or unsupervised and here's an image and it uses the patterns of the image to recognize a face, recall this particular photograph, identify the specific person. Once again, deep learning with convolution neural networks, we can identify specific lesions, micro aneurysms, etcetera as well as diabetic retinopathy or different diseases classifying the severity of the disease, as we've

heard, from different speakers to date.

I want to highlight three -- oh, here's an OCT for segmentation deep learning that's trying to target some of the more challenging aspects of segmentation with macular edema, exudates, and detecting and measuring the fluid in these lesions. Competitions have spurred machine learning progress in general and in ophthalmology in particular, in 2015, there was the Kaggle competition where they classified five levels of diabetic retinopathy using 100,000 images from 50,000 patients with the EyePACS database from California. There were over 661 contestants. The winner, as we've heard, did better than the experts and this was a professor -- or I think an assistant professor from the UK with absolutely no ophthalmology experience.

I want to highlight three recent -- very recent papers of deep learning for diabetic retinopathy detection. I apologize for these slides, the legibility, but I want to highlight these because they were done with deep learning

with very large data sets with independent validation datasets. We heard a lot about validation and the importance of validation using more than one grader, etcetera.

And what's unique about these three studies, and there are others, is that all used, at least for one of their independent validation sets, this method or a dataset from France that had over 1,700 images that were graded by numerous experts as one of their validation sets. The first paper used with a lesion-based approach -- this is Michael Abramoff as the lead author here -- and a lesion-based approach with lots of images, and the area under the ROC care for -- curve for referable diabetic retinopathy, was quite high at .98.

The second, Gulshan and colleagues at Google used EyePACS database and an Indian database and they also used "transfer learning" where they trained the system on non-ophthalmic images at first, and this tends to boost the performance of deep learning algorithms. They had an area in the

ROC curve of .99 for referable diabetic retinopathy. As we've heard, it depends on the target, the objective. The last study was detecting diabetic retinopathy "yes" or "no" with also a very high diagnostic accuracy.

We heard about opening the black box and I think this is really where there's' going to be a lot of work in the near future. We heard about it in the last presentation. Here's another example where the automated generated heat maps identify the regions for closer examination by the clinician. This is -- these are the areas where the deep learning algorithm was focusing, at least in part, to detect the disease in these particular cases.

Also, other areas that haven't been yet touched upon here; pediatric cataracts; these are very high diagnostic accuracy for not only detecting the lesion, measuring the density, the ilea, etcetera.

So where are we today? Well, Google DeepMind, as many of you know, is working with Moorfield's

Eye Hospital using OCT images in the macula, and their work is, I think, being submitted very soon. It's going to be detecting not only diabetic retinopathy is my understanding but other retinal diseases.

Using fundus photographs, IBM Watson is working with IDX and colleagues at University of Iowa and they're algorithm has been approved in the Europe economic area and Google Brand and Eye Research Group is using their work to put their algorithm in India. So that's where we are today.

Obviously, there are many advantages and limitations to AI. We've heard a lot about the advantages, objective reproducibility, tends to do better than the experts; you can modify the sensitivity and specificity to the specific application, and you can -- the model can be trained and it can be relatively inexpensively deployed.

Many limitations; large datasets are needed; Gulshan and colleagues die a post hoc analysis and found that 60,000 images were optimal with 17,000

images for referable diabetic retinopathy. We need well-labeled datasets that we're going to be -- there's also weak labeling is possible. There's -- the black box, I think, is being opened and obviously, there's a lot of regulatory, legal and other issues that I'm really looking forward to discussing today.

There's also unintended consequences in machine learning. One of them is the context. There's a well-known example where a machine learning-based decision support system determined by patients with pneumonia and asthma were at a lower risk of death than patients with pneumonia and without asthma. Well, how did this machine learning algorithm come to that conclusion? Well, what happened was it was trained on a dataset where patients with asthma and pneumonia were immediately sent to the ICE so they had better outcomes. So the machine was accurately learned but the treating set was flawed, so the context matters. And these are things that we have to be area of when we're applying and texting our

systems.

Other unresolved issues that others have mentioned, the patient and physician acceptance of these models; and analogous to the pneumonia example, are these classification systems for diabetic retinopathy that are doing so well, are some of them, because these -- they're detecting eyes with a small pupil in cataract, which is also more prevalent in eyes with diabetic retinopathy, and how much does that matter if this person will be referred, if the objection is referral diabetic retinopathy, does it matter?

So I think with the future with AI and deep learning, there's going to be a general algorithm for diagnosing some retinal diseases. There's going to e new clinical and scientific insights. We're going to be really reinventing the eye exam and possibly allowing more time for that patient interaction where there's going to be seamless integration perhaps with EMR, with instruments, with cameras. The black box is already beginning to be opened and is going to be the eye as a

window into the body. There's already deep learning algorithms for predicting cardiovascular risk factors from fundus photographs.

I look forward to discussing the constraints and unresolved issues in the panel discussion. Thank you.

(Applause.)

DR. REPKA: Thank you, Dr. Zangwill. Our final speaker of this session will be Mr. John Reites, a partner and Chief Product Officer at Thread where highly involved in digital health platforms to enable patient research. Thanks.

MR. REITES: Great, thanks. And while he gets that loaded up, thanks for having me today. I'm going to shift gears a little bit and we're going to talk about this really big topic called the patient interface in digital health. And we're going to try and do it in like nine minutes, so I'm just going to warn you we're going to blaze through this. And we're not going to capture everything, but one of the things I've really learned -- I've connected a few hundred digital

programs with patients all over the world and one of the things that we continue to find is that the interface, the engagement, the interaction that people have with these technologies is just as critical as the scientific validated measure we're trying to get out of them.

And so one of the key components that we have to keep in mind when we're looking at digital health technicians is what is that interface, how does it work, and how does it produce value for patients. So let me give you an example. So I just bought a new TV and I put it in my living room and my four-year-old walked up to it -- so just imagine with me for one second, this is the new TV on the wall and she walks up and I said, "What'd you think?" And she said, (off mic/nonverbal gesturing).

(Laughter.)

MR. REITES: So I think we'd admit that the world's changed, right? We're all carrying these. I'm sitting in the back of the row, I'm seeing everybody on iPhones, iPads, computers, we're

typing, we're engaging but it really has changed. And one of the perspectives I want to do is I want to kind of step back from the science, from all the work we're trying to do and I want us to just take a patient, a consumer's perspective for a few minutes and maybe take some takeaways home from this to apply to all this really innovative scientific work we're doing.

So let's talk about these evolutions happening really quickly. So remember that this digital evolution is not just happening to us. It's actually happening because patients and consumers are pushing it forward in the market. And so let's not think that we're all smart creating all these great devices It's actually that the devices out there are helping people, patients, consumers to see that there's more out there that can be done with digital technologies. And so if you look at this evolution we've been under, there's really four key areas we're in. First is we've been digitizing stuff, right; we've been taking everything we've been doing on paper for a

long time and putting it in digital, right. And everybody kind of knows that's happening but I got to tell you I was at a research site a couple weeks ago, and there were still paper forms being collected outside of an EMR.

So the reality is we all know that we're still in this movement. But the second piece of this movement is really important and it's this made up word called "remotidization" (ph). Remotidization is where we start taking digital things and we make them remote, right. We let patients do them in their homes. We let patients do this as they live their lives.

This third movement though that's really happening, and one that's taking place, is contextualization. Contextualization is not just in the data we collect. Contextualization is actually when I'm a patient and I'm on my phone and I've got something digital and I'm doing it outside of a clinic, and then I'm in the altitudes of Denver, Colorado, the barometric pressure where I'm at may impact data that I'm providing. So

understanding the context of when, how, and why it was collected is really important and that's becoming a variable, actually, in the interface that we're collecting this research data with.

And then last but not last, we've talking a lot about the automation today, right. We look at deep learning, machine learning, AI. These start to take the data, the interface we're collecting and starting to give back insights and information to people. And so remember that evolution; this is what consumers are actually seeing in many other consumer engagements they're having. We're just finally getting to it in our industry.

So let me give you an example. Anybody ever spent your life savings at Disney World by chance, Walt Disney World. Okay. So I have three kids, 10, 7 and 4 and like 25 percent of my salary in QuickBooks is like Disney and Disney products. But one of the opportunities I got because I'm a tech nerd and because I've been involved in a lot of these sort of technology innovations was I got a chance early on to try this thing called a

"MagicBand." You guys ever heard of MagicBand? So I'm wearing one today. I won't tell you how many of these paid for but I have one today, and this MagicBand is essentially what Disney was trying to create as a digital health device. Now it's not the digital health devices we're creating but I want you to just take this example and think about it in your perspective and know that really what patients want is not another device. They don't really care about the data you capture. They care about the experience you're giving them and they know that if you take something from them, you should give something in return.

And so what Disney figured over a long stretch of working through this problem that they had was that their parks were expensive, that actually survey results showed that some of the most stressful situations beside a hospital setting were going to Disney World with your kids. There are a lot of similarities, actually, to the work we live in that when you look at that, they were trying to alleviate and provide a support and

structure for people to make this experience easier but also to collect data along the way.

And so what you need to understand is that Disney is actually having an impact on the work we're doing in digital health. There are two impacts you need to be aware of. One is they've introduced an omnichannel experience and we'll talk about that in a minute. But the second thing they've done is they've actually raised the bar in what consumers and people see as a good experience in digital health.

So it used to be when I did an early sort of mobile app like eight years ago, I built it and it was the ugliest thing you've ever seen. It was ugly and it worked and it could be validated but it wasn't very engaging. It really just took a lot of data from patients but the reality is, is because it came out of a research institute, nobody even gave it an issue. Actually, the patients are like, okay, great. They kind of said oh, I expected this to look like this because it came out of a -- out of your practice.

And fast forward to today though, because patients have been engaging, consumers are engaging in all these experiences, they're bar is raised So they see thee apps, and they go "that is an ugly app," close it. Wow, this app's asked me for nine reminders today and not given me any value; close it. So that same thing, I'll just tell you, from all the data I see every day with tens of thousands of patients across the U.S. shows that patients do the exact same thing even when there's altruism involved, even when there's a medical device involved. So we have to be cognizant of this omnichannel experience.

So what's an omnichannel experience? It's just a key word that I could leave you with one word today to think about. An omnichannel experience is using multiple channels that integrate together to provide one seamless experience for a user. So what that means is we -- tend to focus on the point solution, right; you focus on the digital health device or what it does but remember from a patient's perspective, that's

just one of a lot of things that they're experienced with, right. So if we go back to that slide before and we think about Disney, they don't just have a wearable, they have a mobile app. They have a web experience. There's a location they're going to. There's all these different sort of locations and places they're experiencing with and all of them are coming to the same value and to the same goal. And so this omnichannel experience means that whatever you do in digital health, make sure that it's connected to all the other things that a patient is experience. IT should be linked to the location therapy go to. It should be linked to the mobile app and the web experience and to the medical device.

So don't just think about the point solution you have or you've made. Think about how it integrates into overarching experience because that's actually what patients want. That's actually what patients are being trained to do in the consumer world.

So on this little slide that I know you guys

just all want to punch and make bigger, right; am I experience receptive? I want to leave you with just a couple points to keep in mind. So I know that we're really training to come up with these digital health solutions and to really focus on the data collection and the validation in those pieces. But one of the things that, frankly, really hit me really hard about six years ago was I got this chance to enroll in a clinical trial myself. And I won't belabor the story because I don't have time but I'll just cut to the chase. I dropped out of a clinical trial and I've been running clinical trials for 15 years. And I dropped out because the experience was not great and one of the things that I really took away from that learning was that even though I'd been doing clinical trials for so long and then I was a patient myself and experienced it, the same issues were coming up whether I workshop surveying patients or whether I was the patient.

And these four issues I think we can really dial into this area. When we think about a

digital health interface, there are four key areas I'd like for you to keep in mind. The first is -- first and foremost is value. So if you're going to do something in digital health, don't just think about the value we get out of it from the data, think about the value that the patient gets out of it making sure we instruct them that when you do this, we get this data, this data does "x." I can't tell you how many apps I've seen this year that don't do that. They just say do this, collect this data. They don't take advantage of the opportunity to help apt understand the positivity of what we're doing.

The second piece is the experience, like I talked about this omnichannel experience, bringing together your solution into a mix of other things that a patient engages with.

The third piece is what we call balance. Balance in digital health is really important because a lot of the things we've been building have been very active, right, very activity-based. Please click here, do this, touch that but

remember that there are a number of different passive sensors, clock alarms, all kinds of things that are interrupting our day, and we need to make sure that we have a good balance of active versus passive things we're having patients to do.

And then last but not least, again, is the channel. And when you think about the channel, we're not just thinking about -- this is not a TV channel; this is what channels are you using to get people to engage with the digital platform you have or see the results. And remember that most patients are using, just as a basics, a mobile app and a web experience. So if you're not -- if you don't have at least those base minimums, you're not reaching the majority of the population that would want to engage with your digital health solution. So again, try to rapid through a lot on a patient interface, we're going to talk more about it on a panel later today.

But hope that gives you some thoughts to think about. We take the patient's perspective whenever we're implementing and starting to coordinate our

health efforts. Thanks so much for having me. I really appreciate it.

(Applause.)

DR. REPKA: Thanks, Mr. Reites. Thanks to all of the speakers for their engaging comments. It's 10:20 so we are going to go to break. We do reconvene at 10:35 so just 15 minutes. If the panelists for the first panel could just stop up real quickly so we can make sure that they have a plan, that would be great. Thanks.

(Whereupon, off the record at 10:23 a.m., and back on the record at 10:42 a.m.)

DR. REPKA: You can start with --

MALE SPEAKER: Michael, microphone.

DR. REPKA: Oh, sorry. Is that better?

MALE SPEAKER: Yes.

DR. REPKA: Okay. Dimitri, please just say a few things about yourself and we'll go around the table.

DR. AZAR: Hello, everyone. My name is Dimitri Azar and and I think there are -- I apologize that the conflicts of interest go beyond

what is listed there. We have to add Verb Surgical and Novartis. I'm on their board, so I apologize for that.

I am currently the Dean of the Medical School at the University of Illinois. in a very unusual arrangement where I'm spending only a day a week as a Dean. We hired an Acting Dean, the Chief of Radiology,, and I spend the balance of ht time at Google as the Verily Life Sciences Senior Director for Ophthalmic Innovations. So nice to be here. Thank you so much for including me.

DR. ZIMMER-GALLER: Ingrid Zimmer-Galler. I think you already heard a little bit about me. I have a long history in the past being involved with diabetic retinopathy screening and currently, I split my time half between Wilmer in the Retina Division and the other half is running the Office of Telemedicine for all of the Johns Hopkins tele -- all of the Johns Hopkins health system so not just tele ophthalmology but all of telemedicine.

DR. MOSHFEGHI: Thank you for having me here

today, Darius Moshfeghi. I'm at Stanford University. My areas in telemedicine surround pediatric retina, specifically retinopathy of prematurity and also universal newborn screening. I have numerous conflicts with Visionex, which is a camera company. I am involved in a screening company, an artificial intelligence company, and I serve on the board for 1-800-contacts where I work on their telemedicine outreach.

DR. WOODWARD: Hi. My name is Mia Woodward. I am cornea specialist from the University of Michigan and I co-direct the Kellogg Eye Center for eHealth. I also serve on the Academy of ophthalmology's Telemedicine Task Force and have an NIH grant to study telemedicine for anterior eye diseases.

DR. TRESE: I'm Mike Trese. I'm a pediatric retina surgeon in Michigan and have done quite a bit of work, as Darius has, in ROP telemedicine type things. And I'd like to introduce my co-moderator.

DR. AFSHARI: Natalie Afshari, talked earlier,

Professor of ophthalmology from the University of California San Diego. It's a pleasure to be here and I also wanted to let you all know that the audience can ask questions once the panel discusses a question. So question one and four will be our charge and please feel free to ask questions once questions one is done. Thank you.

MS. BOTTORFF: And I'm Leslie Bottorff. I'm with GE Ventures. I've been in the venture capitalists about -- capital business about 20 years, the last four with GE and we're invested in a number of portfolio companies across digital health as well as some other areas and pleased to be here.

DR. MORRISON: Good morning. My name is David Morrison. I'm a pediatric ophthalmology and I'm the director of the Telemedicine Screening Program for Retinopathy of Prematurity at Vanderbilt University.

MR. PATEL: Hi everybody. This is Bakul Patel. I'm the Associate Center Director for Digital Health at CDRH and I lead sort of the

efforts on digital and various aspects of how emerging technologies are coming together in this space and how they're cutting across every aspect that we have regulated in the past and what those connections really mean. So I am also leading the pre-certification program, as you heard me talk this morning. So happy to be here. Thank you.

DR. TRESE: Well, we have a very exciting thing. We have some new technology that was not discussed this morning and won't be this afternoon. And it's basically made for people my age that are getting into the digital age, and that is this is Mahmud, who is your personal digital health advisor. So you may want some of these as you go along.

So what we're going to do is discuss two of the questions that the committee came up with. And the first question is one that deals with safety concerns. I think we had a really nice discussion this morning relative to control of risk and benefit of efficacy and what gold standards are and what they may become and things

like that.

And I think these -- this first question has some interesting implications and let's just read it together. A digital health device provides a diagnosis, a computer-assisted diagnosis for screening diabetic retinopathy by adding on software to a fundus camera image in comparison to a digital device that provides information as an aid for diagnosis to the healthcare provider. I think you have an answer, perhaps, Mr. Patel, but why don't we start with the Dr. Azar and give us your opinion on that.

DR. AZAR: I think that we want to advance on both areas and here the question is about safety and effectiveness. Questions at the early stage are going to be where do we draw the line along this continuum that we've all heard about; referable versus not referable in areas where you're trying to -- let's take diabetic retinopathy as a good example; that's at least a group I'm involved in is focusing on and there it's relatively easy to go ahead and let the

machine learning algorithm, for example, the machine make the diagnosis because it's a low level impact. You can always increase the sensitivity at the cost of effectiveness and at specificity as a result of which there will be more costs but greater safety. There's always this balance to draw.

Now if you want to go into various subgroup diagnoses, it's going to be very difficult to do that. The gold standard approach has become much more difficult. We know that today's gold standards are -- need some alchemy, but you can't do that today. You have to look, I think, at the end at outcomes meaning whatever has come out of the early studies were based on the different subgroups and the machine learning algorithms of today have to simulate that. But at some point, new categorization has to come out looking at long-term outcomes of patients who have been diagnosed and we're really far from doing that.

And from an FDA perspective, I think categorizing some of these rapidly-evolving

systems is -- may slow down the pace to improve safety but at the same time may, if the pace is not very slow, we may have new developments but then there's a potential downside of having potential serious unintended consequences that may end up stopping many of these processes, and that's where the FDA has to draw these lines.

DR. TRESE: So Ingrid, how do you feel about those things?

DR. ZIMMER-GALLER: So I basically agree with everything that's been said here. Diabetic retinopathy is certainly a great place to start with automated image analysis or computer-aided diagnosis, because unlike, as we heard earlier with retinopathy of prematurity, there is much more variability in how experts read those images, and we really have very good consensus on diabetic retinopathy. It's a much better or much more easily-defined disease state.

I want to step back for one second and just also remind everyone why this is something that is going to become more and more market participant.

Obviously, the diabetes epidemic globally is -- that's going to continue to get worse. We clearly need to do a better job of evaluating patients with diabetes 4 retinopathy because we have fantastic ways to treat the disease, and we can't -- we cannot afford to continue to have patients come into our practices that have traction retinal detachments and, you know, come in at a point where we really can't do anything to help preserve their or maintain normal vision.

And I'm sure everyone has heard some of the numbers that have been tossed out but I think, for example, if every patient with diabetes in the world were to have the recommended eye examination, we would have to do one every seven seconds and, you know, we clearly don't have the workload either to examine all those patients in person but we also don't have the workload if we have millions of images -- we don't have the workforce if we have millions of images that need to be evaluated. So I think this is something that very much we need to continue to work on in

advance. I think it's certainly, for a disease like diabetes, imperative.

DR. TRESE: So I know that I'm in Washington, DC now because the question really deals with risk analysis between the risk of a device that gives you a diagnosis and something that gives the doctor an aid. So what would the two o you -- how would you -- would you grade one of those as more risky than the other and if so, why?

DR. AZAR: I think there is a happy solution by thinking about the context in which the machine is providing either help or diagnoses. For the screening, I think especially in diabetic retinopathy, the referable versus non-referable, I wrote down what the ROP people use -- referral warranted, it's the same idea. The referral -- I mean there, as I said earlier, you can increase the sensitivity but what you could do for aiding a doctor in instances where there are potentially more difficulty where you're trying to make a diagnosis, etcetera, I think at this stage, we would need to go into a technology that's easily

available but that's inside the black box which is the heat map assistance, meaning the reasons sometimes these programs do better than the experts is that they can focus on 50 notes in one image whereas an expert, given their limited time, even under time when they're given time, can focus on four or five different areas. So by identifying heat maps, you can assist the healthcare provider, the ophthalmologist most often, to actually make a diagnosis. We leave it up to the ophthalmologist and the higher the risk in missing a diagnosis, the more likely there will be a need to assist the physician rather than replace the physician.

DR. TRESE: I agree with that. I think that's a very good point. Ingrid, what is your opinion on risk assessment.

DR. ZIMMER-GALLER: So I think it basically does boil down to validation and the validation has to be appropriate for what you're trying to accomplish with that -- with the program. You know, it's interesting we talk about validation

but we don't validate physicians and, you know, so you can -- you know, clearly there are physicians that do a not very good job of diagnosing any of these diseases. But I think with proper validation, I think you can very definitely keep the safety issue -- you can control that very well.

DR. TRESE: Darius?

DR. ZIMMER-GALLER: And that also includes you need to have ongoing QA.

DR. TRESE: Yeah.

DR. ZIMMER-GALLER: You need to be continuously monitoring all of these programs to make sure that things don't change over time.

DR. TRESE: Go ahead, Darius. What's your opinion on this? You may take the ROP point of view.

DR. MOSHFEGHI: So when we look at the safety of diagnostic-based systems versus diagnostics-assisted systems, it really comes down to what Michael Chiang was referring to earlier, is what is the intent of the system, because at the system

level, the difference between them is really kind of arbitrary. It's what are you trying to do.

For example, when we were looking at a diagnostic system that's going to work independent of a physician, that may be appropriate for certain low risk situations where the rapidity of which the disease onset can occur and the magnitude of the bad thing that can happen from the disease are not very large. And so an example of this may be glaucoma screening in a general population. You're not going to go blind immediately and if you -- you know your risk of vision loss is very slow over a long period of time on one missed examination and so I would feel quite comfortable using an independent diagnostic system in that sort of situation.

When we go into the opposite, which is retinopathy of prematurity, both the disease severity, you can end up bilaterally blind and the speed at which that can happen can be within 24 or 48 hours. And there I'm more comfortable using an ROP-assisted sort of diagnostic system.

Then you get into these intermediate areas of where things can go wrong which is such as diabetic retinopathy where the disease, obviously, you could end up with bad macular edema, proliferic diabetic retinopathy, tractional retinal detachment but clearly, the screening burden is very large and we can tolerate a lot of macular edema and a lot of diabetic retinopathy for a long time and still come in and end up with good visual acuity outcomes. So the overall risk is low but it's higher than what we see in the glaucoma situation so I would be more inclined to go towards using a -- I'd be a little bit happier using that, a diagnosis-only system in that sort of situation than I would where the rapidity and the magnitude end p being a lot worse.

DR. TRESE: So your risk analysis really is based on rapidity of disease progression and severity of outcome --

DR. MOSHFEGHI: Sure.

DR. TRESE: -- and that time. And I think that's what Dimitri said first of all was that the

time feature with diabetic retinopathy is a lot less than ROP. Mia, what is your opinion?

DR. WOODWARD: Well, I'm very happy to follow those comments. So as an anterior segment specialist, you know, the diseases in the front of the eye that are population health level diseases that we really should be focused on here today are ones that are urgent and, you know, ones that don't have a really known underlying condition. You know, we don't know the patient has diabetes; we don't know they were born prematurely; we don't know they have macular degeneration. So that's our problem with the anterior segment diseases.

You know, people come in because they have symptoms, they have eye pain, their hurts, and where they show up is also very different. So, you know, the problem is is about two million people come to the ER per year for eye complaints and half of those eye complaints are nothing, they're -- well, they're not nothing but they're not things that need me that day, right; so they're dry eye; they're --

DR. AZAR: They do not need a retinal consultation?

DR. WOODWARD: Correct.

(Laughter.)

DR. WOODWARD: Nor anterior segment consultation urgently. You know, they're dry eye, they're a floater so they're -- but not a retinal detachment so -- but what we worry about is any one of those people, are they angle closure glaucoma; are they a corneal ulcer; you know, are they diseases that could be very severe and could progress very rapidly. And for anterior segment diseases, you know, I think -- so time is very important for us and humans.

So I interpret this question to say like what's the value of humans and, you know, when are we useful. And I really enjoyed Dr. Yeshwant's comments to that effect, you know, because I hope that machine learning will help us not have the burden of things that we don't understand and be useful as human beings. And I think that my added value as a human from a rapid standpoint is I can

tell if -- you know, cornea patients and people in the ER walk out the door and then you tell them to come back in one or two days. And so I can tell if that person's not going to come back in one or two days versus a, you know, a diagnostic independent device, right. So I can tell if there's alcohol on their breath; I can tell if they don't have a ride to come back the next day and so I think that's the value added of a human.

And I also think geography is very important for your anterior segment diseases, like these people have symptoms in their home. The young people are going to go online first to maybe triage their symptoms. They're not going to -- you know, they're not even going to think about going to the ER unless it really, really hurts and -- but then they're going to go to an ER or a primary care doctor. They're not going to come to an eye provider. They're not going to be in a hospital getting ROP screening because they're a young baby. And so, you know, there' a huge role of devices.

I think, you know, what's better in an ER setting is also interesting. You know, only one-fifth of medical schools now teach any ophthalmology training whatsoever, so the primary providers, whether they're in an ER or an ED do not know ophthalmology. So these devices that can say this is a bad eye thing, please find an eye provider really does have huge value and has huge opportunity to triage patients appropriately to get the right ones to us.

And I also wanted to bring up one separate point. You know, I do think that -- my concern -- and I like that we talked a lot about trust of systems earlier in the morning. I have a concern about how devices are being built around young versus old patients.

You know, the statistics were thrown out but 77 percent of people have Smartphones but all of those people are -- like they're not the older patients and they're not the poor patients, and those are the people who are the sickest and, you know, old people get sick and so if you're 65 and

older, you're more at risk of having diseases.

And so applications that are Smartphone only and not -- no, don't panic if that's you -- because I mean ultimately, it's all of us, right. You know, I'm going to be 65 someday and I will get sick and I hope that the device that's built and, you know, it's about the sort of patient experience, the omni -- what was it -- I know I learned a new term this morning -- omnichannel, like that's fantastic, right. You know, it is about that experience that anyone can relate to. I mean we all went to the airport and half the people still have paper tickets because they don't trust, you know, that they can do it on their device. And so we have to do that -- you know, we have -- whatever we do, we have to design it for all patients.

DR. TRESE: You know, I've often thought that 90 percent of that slide that shows the percent of people that have cell phones are in the East and West Coast and not in Michigan, that --

(Laughter.)

DR. TRESE: And then in addition to that, you not only have to own a Smartphone, you have to be able to turn it on.

DR. WOODWARD: Right.

DR. TRESE: And that can be challenging. Darius has trouble with that sometimes. So Natalie, do you have an opinion there?

DR. AFSHARI: Well, I think one of the most important things that was brought up by Mike Chiang was -- and also Paul Lee -- it's what is the gold standard, you know, if you're going to have this safety and efficacy that is high and that is our charge, to really decrease our error rate and increase our efficacy and safety over time. You know, when experts don't agree, then what? So, you know, while we have started a great road, in some fields, we're not quite there and while it's great in anterior segment, that many of our patients otherwise wouldn't get diagnosed or wouldn't -- you know, in the ER, they show up and that's the best thing that we have, we still have some roads to really decrease our error rate and

our safety and efficacy.

And there is so much more in retina, as Dr. Azar said. It's not all about referring to the retina doctor, so what about these anterior segment diagnoses. And many of them can be really crucial right there and then; you know, is it just a regular red eye or is it a corneal ulcer as Dr. Woodward said. And that could have devastating visual consequences. So we have a little work to do but we are, as I think Dr. Bakul (sic) said, we are in high road right now, so.

DR. TRESE: So Leslie, you bring a little different perspective to us and for those of you that may not be aware, Leslie has done a lot of work in radiology. And so I would think there'd be some of the same type of concerns in terms of radiologic diagnoses being made by machine as opposed to the doctor. Can you address some of that?

MS. BOTTORFF: There are absolutely the same types of concerns and radiology may be even a little bit ahead of this in terms of the

proliferation of the number of companies that are doing various types of image analytics and combining that with other analytics. And what's happening in that field is that initially what we're seeing is the adoption is about automation and efficiency. And, you know, that because these radiologists are getting so many more scans per study and so many more patients and volume that they just can't do a good job, just like what you were talking about in looking at these images. They just can't do a good job in terms of the time and money allowed for this. And so they're using these tools to help them scale basically and to have greater efficiency but yet they're still making the diagnosis in the end.

And what the physicians have suggested here I think is exactly right, that depending upon what that information is going to be used for and, therefore, what safety risks and what kind of time scale do you have to work with with that patient should be the factors in determining what level of sensitivity and specificity that these devices

have because they -- you know, the other constituencies here are these people who wouldn't normally even get any care.

I mean, you know, scale and access and reach is what these new technologies will do for you, the amount of patients that you can see before, you know, like let's say that are - that haven't -- are going to have a diabetic retinopathy problem, that you can see them while you can still do something with it, they are your constituency also in terms of safety and thinking about those.

And so I think that the direction you're going here is exactly right. And the other thing I would say is that all of these technologies are great but you have to get them paid for. The economics here are key because the fact is is that, you know, it doesn't get adopted commercially if you can't get reimbursed for it and if it doesn't make sense on time and money allowed.

And so I think that it's fantastic to see the FDA and the physician groups and the industry

groups all working together, you know, in this ophthalmology area. And in fact, it's one of the specialties that can really take the ball and run with it on digital health because of the ambulatory nature generally of the practice and also because these digital technologies lend themselves very well to people with eye problems.

So your group, really, in this room could be the leaders in this but you have to bring along the economics of this, the reimbursement parties and not that any of these -- any of you, you know, of those -- these groups control that but have to make them part of the conversation and also make efficiency and resources allowed part of the consideration in terms of, you know, how are we going to use these devices and what can we really use them for.

And then later on, you're going to be able to out of the efficiency and automation level and into the multi data source diagnostic predictive value of this, and that's going to be incredible also. And the data that you're going to collect

as these things progress are going to really help you to get to a higher level of diagnostic power as well.

DR. TRESE: David, do you have some comments?

DR. MORRISON: Yeah. I'd like -- actually like to put on my pediatric ophthalmologist hat to answer this because I think it emphasizes a couple of different things that we've touched on so far. I think vision screening is really the original telemedicine in ophthalmology. We've been doing it for years and years and it was initially software in a medical device and as we move forward, it is not software as a medical device.

And I'll kind of hit on a few points of how this could be positive or negative. I think the benefit of making a diagnosis with a machine is that you can absolutely improve care and improve the finances of care. In addition to that, I think that you run the risk of excluding the physician completely in certain circumstances. And so let me give some examples. About 20 percent of kids will have amblyopia risk factors,

so refractive error, anismetropia, different things like that; about two to three percent of the population will have amblyopia and if not discovered and treated, then you can have permanent visual loss, and so that's obviously a significant problem. We do have a large window to treat but it's there.

So with vision screening, initially, we were taking photographs of the eye and were looking at the red reflex and on a film-based camera, someone was looking at the photograph and determining whether that's normal or abnormal. As we moved forward, now we autorefractors and different levels of technology that can absolutely diagnose relatively accurately the refractor error itself and not just say "yea" or "nay" but this is your diagnosis.

Further, as we develop this technology in apps, there is an app that you get on your phone called "go check kids" and that's the red reflex test that basically says "positive" or "negative" and there are multiple other apps. I wrote them

down. I Googled it while we were watching so there's one called "the eyes can" app, Blink Netra. MIT has an app that they've developed that can actually diagnose refractive error on your Smartphone in an autorefractor-type setting.

So the pros; if you have this technology, you can reach a bunch of people at once. In Tennessee, with our outreach program, we've screened over a half a million children that likely would never have been screened or had their amblyopia diagnosed. We also went back and looked at kids who had normal exams that ended up getting glasses and it was shocking.

So if a child had a normal exam defined as the absence of a post-amblyopia risk factors, if a child had a normal exam and they saw a pediatric ophthalmologist, about two percent of the time, they got glasses. If they saw a comprehensive ophthalmologist, about 12 percent of the time, they got glasses. And if they saw an optometrist, about 35 percent of the time, they got glasses. I'm not going to comment as to why that may be.

I certainly am not going to imply anything by it. I will say that there's obviously a wide variance in how we treat these kids. But if a child has a normal exam and 30 percent of the time in a state-mandated program, every child in that state is getting glasses when they don't need them, that's poor care and it's poor use of healthcare dollars. So the technology does have the ability to improve care and certainly improve finances.

But let's look at the other side of that coin now. Say we have a parent who screens their child with one of these app screeners, finds amblyopia, goes to an autorefractor, diagnoses the prescription that they think the child needs and then they go to Zenni Optical or one of the other online stores, enter in those numbers, they can theoretically diagnose and treat their child's own disease without ever having seen a physician. and I don't think that's good care. I think it's the opposite of good care.

And so I think that as we move into this new

realm, there's definitely positives or negatives but I think that this specific example shows us kind of the extremes of what it can be on either side. I think in the end, big data and the ability to improve care will win but there are some pitfalls.

DR. TRESE: I think you bring up a dilemma that's common to probably everything we're talking about today and that is that I notice as I'm dealing with residents that as soon as I say anything, they take their Smartphone and they get on Google. And so I think we've talked about artificial intelligence. We've talked about deep learning. What we need to do is try and structure what Paul Lee so nicely showed as the new medicine. How do we use these things to our patients' advantage? I'm very happy about the program so far because it's clear that the message is to try and develop better care for patients, that the FDA is delivering and that I think that all the speakers have.

And I wanted to ask Mr. Patel this question

because you're a little closer than Malvina. That's the only reason I'm asking you. Has any of this helped you that you've heard so far?

MR. PATEL: I would say "yes" and a resoundingly "yes."

DR. TRESE: Because you're a nice person?

(Laughter.)

MR. PATEL: That could be one of the reasons but I just wanted to make a couple of comments and just hearing people's opinions here, I think what we are seeing is just not about them in the questions phase as raised is about what are the concerns, right? So I think every conversation that I've heard so far is about benefits and risks. It's not about concerns only but there are some benefits that come with that. And it comes from the aspect about can it deliver care at the right point, at the right time, to the right patients in the right way. It sounds very "right" but let's just leave it at that for the moment.

But how do we think about those new benefits that are sort of coming into play with the new

risks that are coming into play as well? And there are some social risks with it; there's a provider risk to it; there's actually some transforming that's happening because what if those gold standards that were are all accustomed to is different today. I think that's really what we are asking in terms of, in my mind -- for me, it is fascinating to see that we are talking -- we are having a conversation about gold standards. And I think fundamentally, in my mind, gold standards are being changed with this technology. And it happened in the imaging world. The gold standard was changed. Radiologists were the gold standards and then they had aids to help them spot things that they couldn't spot when they had the volume come across.

So that's how I think about these worlds. So we can lose the sight -- lose sight of the fact that there are benefits and there are some new risks that we are not quite there in terms of understanding how big or small a risk they are.

One quick thought I think from a positive

perspective and a negative perspective is when I heard and I was sort of putting notes down, there's this opportunity to sort of detect early, so early that interventions can be extremely small. So that's a big opportunity I see with this technology.

And then on the risk side I see as negative is we are not being trained to recognize when things are not what we expect it to be. So how do we change that equation from, you know, from med schools to engineering schools to delivery and etcetera to -- even for FDA for that matter, like how do you start recognizing that detectability of error, which we all know is really easy when we know a fundus camera with whatever it puts out, it puts out, and we know what those readings mean, right, and that this technology, when it's in the hands of patients, how do you sort of allow the detectability to be there that's ubiquitously and doesn't require, you know, eight years of college to go to -- so I'll leave it at that.

DR. TRESE: So I -- when I first looked at

this question one, I said this is so trivial, I can't believe it's actually a question and how in the name of God are we going to spend 30 minutes talking about this. But -- and originally, I thought well, it's simple. If you're aiding the doctor, that has to be less risk than if a machine is making the diagnosis. That would be my initial opinion.

But then I agree exactly with what Ingrid said earlier about validating the physicians. And so we have an ROP software program called "FocusROP" and in it, we have an education module. And if you're OMIC insured ophthalmologist, you have to take a test and you have to pass it with an 80 percent and you get three tries. Okay. So we first launched that maybe 2.5-3 years ago, something like that. I think one out of 19 people got an 80 percent the first time. And obviously, ROP is the worst thing that OMIC insures.

So I think that the education component -- the gold standard thing I agree is changing but I think the education component is really very, very

important. And to get the doctor that just what Paul was talking about earlier -- to get the doctor to be educated in terms of either diabetic retinopathy or retinopathy of prematurity or anterior segment disease or any other -- glaucoma, any other thing that lends itself to telemedicine is still extremely important.

So I can cross off my list here payment, Leslie. Thank you. That was on my list to do. And the gold standard, I think we've discussed pretty aptly. And I think that we can probably switch now to question four. Natalie, you're going to do question four.

DR. AFSHARI: Are there questions from the audience?

DR. TRESE: Yes. Do we have any questions from the audience? Yes.

MALE SPEAKER: (Off mic.)

DR. TRESE: Oh, there's a microphone there if you wouldn't mind.

MALE SPEAKER: (Off mic). (Inaudible) something about computer systems (inaudible) and I

remember being (inaudible) New England Journal (inaudible) screening program, and it showed docs assisted by a computer in their work in (inaudible) actions and then (inaudible) regular(inaudible) alone. And (inaudible/off mic) rather than computer-assisted so I'm just wondering what (inaudible).

DR. ZIMMER-GALLER: I can just -- I'll make one comment. As -- just from my past experience, I can say that for diabetic retinopathy, having -- and I haven't -- I have no experience with computer-assisted reading but the reading center that we had, when I read images when I was doing over-reads, no question; when somebody pointed out already -- had already circled the lesions that were there, it was infinitely easier. Literally, you take a quick look and you immediately say, yes, I agree, that's a hemorrhage, or yes, I agree that's NBD or yes, I agree, you know, whatever the lesion is and it clearly, as far as reading diabetic retinopathy, from my standpoint, having something already pre-read that image made it much

quicker for me. So I do think that that can be very helpful but I also think that for diabetic retinopathy, I think that we are at a point where if it's properly controlled and validated and has QA, I think very clearly, it's an area -- and I strongly believe that we are at a point where we can use automated analysis.

DR. MOSHFEGHI: I think those are very excellent points. One area that I would like to differentiate a little bit is that there's a difference between doing a screening one off for, let's say, glaucoma or diabetic retinopathy and then monitoring an active disease like retinopathy of prematurity. And so I'm a little more happy using diagnostic-based systems for one off screenings and a little bit more concerned with the rate of progression in diseases that we're actively monitoring.

MS. BOTTORFF: Yeah. I just wanted to point out one difference in the CAD, like the R-2, when that came out, that was CAD design. You know, they were like 80 percent or something like that

sensitivity specificity and they never got any better. It didn't get any better no matter what. It was not a learning technology. And I think that's' one of the big differences in what's happening with the new technology wave that's happening today is that it does get better with -- as it learns, you know, because it gets compared with the outcomes and they get to add that knowledge back in.

And then the other thing that's really different is that with these new digital technologies like "Mobility," then you get to continue to collect and monitor that same data and so you get a lot more data streaming in, you know, than what they ever got with the R-2s of the world. And that technology is still used to help mammography, for instance, but the radiologists complained that it took them actually more time because they had to go through there; whereas -- so that was an efficiency problem but it helped the sensitivity and specificity.

And so with the new technologies, I think that

some of those problems, you know, are -- go away. And the numbers you saw presented earlier were much higher in terms of what kind of sensitivity and specificity that these are able to get.

DR. WOODWARD: Can I add one comment? I think it also very much depends on the user, right. So it's all about Bayes' theorem and, you know, what's the user's pre-test probability; what does the device add; does the user understand how good that device is in terms of sensitivity? And I just wanted to tell an anecdote.

When I was pregnant with my second son, I developed my second kidney stone and when you're pregnant, they can't do the same imaging. And so I got an ultrasound to detect my kidney stone, right. And then the resident comes into this pregnant kidney -- active kidney stone me and says, "We don't think you have a kidney stone; the ultrasound is negative." And I was like, "You've got to be kidding me, right?" You know, so I was like Bayes' theorem says I have a kidney stone. My pre-test probability is like 90 percent. This

diagnostic test is very inaccurate. I have a kidney stone, you know.

So I mean I think that like this is the point, you know. If you don't know that the accuracy is, you know, 80 percent sensitivity specificity and you are confident in that test and you're not a knowledged user, like you're -- it's a higher risk situation.

DR. ZIMMER-GALLER: But I think we can also add to that and think of what we can potentially do in the future. I know -- I don't think we're there yet but at some point, if w can plug all of the -- looking at analytics, if we can plug all the data in, if we can add how long the patient has had diabetes or when it was diagnosed, if we can add what their A1C is, if we can add all of those things in and include that in the analysis I mean the -- you know, from a safety and effectiveness standpoint, you know, in the future, these are things that will be tremendously valuable.

DR. AFSHARI: Mike.

DR. CHIANG: Good afternoon. Like I was -- (inaudible) panel. I think that there are two issues within the culture of medicine that are relevant to telemedicine and sort of computer (inaudible) -- you know, number of (inaudible) by virtue of being board certified, I'm considered competent to make diagnoses and manage things that I'm probably not competing. And I think that applies to everybody.

And then the second is that we've got a culture in medicine where when we make mistakes, we get punished for them, and I think that it could e argued that both of those are barriers to quality improvements that really could be addressed, you know, with these sort of technologies that we're talking about.

And so my question to the panel is what do you -- how do you think we can address these issues in terms of the culture of medicine and is there anything that we can do from a regulatory standpoint that can sort of promote that gradual -- that cultural shift?

DR. TRESE: So I have a comment. I think that's a great question but I think maybe Dr. Repka could answer it better than most any of us, because I think it deals more with legislation; it deals with the doctors of the world wanting their licensure to be so broad that I can go home and do a breast biopsy or an appendix. Are you kidding me? And I meant that's a licensure issue. the licensure issue, I think, needs to be broadened relative to telemedicine but you're absolutely right. To be board-certified -- I don't want to put a tube shunt in either so, you know, it's -- I think it's a very, very good question. I don't know an answer.

DR. AFSHARI: Other panelists?

DR. MOSHFEGHI: I actually like that question a lot because it kind of goes towards the whole problem that we deal with in retinopathy of prematurity and a lot of this stuff that you brought up with experienced physicians having a lot of change there.

Roughly, we have four million live births a

year; 400,000 are premature; 80,000 are eligible for screening which comes out to be about 10,500 week. And you could have 15 highly trained individuals using telemedicine with cameras distributed over 1,000 different NICUs around the country doing three days a week, you know, eight hours a day of reading, and then if you have assisted device using ROP plus algorithm, evaluators and, you know, trying to -- you could really eliminate a lot of the people who shouldn't be screening and take it from referral warranted ROP up to treatment warranted ROP up to the really, this is the one that needs to be treated in an hour kind of ROP.

And you can avoid putting people into positions where they shouldn’t necessarily be. We have general ophthalmologists doing screening for retinopathy of prematurity; we have other people screening in other areas that they're not necessarily experienced in. And I think this is an area where telemedicine could actually enhance what our safety network is overall by bringing the

very best trained people to the diseases that need them.

DR. AFSHARI: Great point and an excellent question. Other comments? I think Bakul --

MR. PATEL: I was just going to make an observation and perhaps this is more to the combination of the first question and the second, I think what we are seeing and witnessing is a need for sort of one technology aiding to the right points or clinicians who got validated once in their life and got their license. And to your point about being continuously validated, so as humans, we get validated once, get licensed to go practice and then we rely on something that's validated continuously. How can those two things come together? I think that's really what the question comes down to is when we talk about aiding and making people make choices that are right in terms of patients at the end of the day, I think that's really where it comes down to.

So my observation was more about I don't think it has to be one or the other. I think it has to

be n combination of like how can technology, schools, licensing boards, and other things can come together to figure out what the right, you know, in machine learning terms, what the right minimum or the right maximum can be reached. So we need to maximize this and I think what technology is sort of enabling or getting us there to think about is like how do you maximize those, you know, positives and negatives.

DR. AFSHARI: Great. So --

DR. TRESE: Can I have one more comment?

DR. AFSHARI: Oh, yes.

DR. AZAR: I think this came up before and this follow-up on the issue of trust that today I think any of us asked the question, and it was asked in previous sessions, do people, do the patients, the doctors even, trust that black box. And it seems where asking the question, the implication is "not yet" but I can see a day when we combine this with the question that will just ask -- that we'll be asking -- we trust the black box now, can we trust the doctor, because -- and

that's going to be the status of affairs of the future and we have to prepare ourselves for it at all levels, technology, medical student education, and residency and fellowship education.

There's going to be a lack period between the two but I think that's a question to have to be prepared at the educational level as well.

DR. AFSHARI: Great point. So we'll move on to question number four and this focuses on patient privacy and there are three prongs to this; one regarding electronic medical records; second, about storage; and third is about patient behavior and locations. So let's read the question together. What are the assets, traits and vulnerabilities that should be considered and identified as a threat to the privacy of a patient by ophthalmic digital health device developers. And there are three sub-questions; transmission of information to electronic medical records or other databases; b) storage of information on the personal device or cloud devices; and c) monitoring patient behavior and locations.

So let's start by -- from Dimitri.

DR. AZAR: Well, I don't think I'm highly qualified to answer this so I'm looking at it from perspective of this similar to the HIPAA compliance days of the late 1990s. There's going to be two issues I think; one is the intentional misuse of information whether it's because of sloppiness or because of malice; and the other is unintentional, meaning despite all the safeguards that are applied, there could still be some unanticipated problems that hackers can go into the systems and work on them. So that's a potential difficulty because then whoever develops the databases or the analyses or the way to guard against the issues have to be knowledgeable enough. And you would think you can do it.

I remember -- I'll give you an anecdote here. At the university level -- and I'm not judging universities versus private business -- I thought we had a good enough, very secure IT system because the number of hackings that, as a Dean of a medical school, I was aware of was very small.

But I think it's because nobody was interested in hacking us. Move to a private company on which I'm on the board, most of the discussions at the board level are about how do we avoid this from happening and it's a fear that they have that --

(Leaf blower noise interruption.)

DR. AZAR: -- is this a hacking?

(Laughter.)

DR. AZAR: And you can tell it's in transition; the systems are there. The expenses are numerous but it's a problem that these companies are dealing with. Now you move to the new place where I'm now spending most of time, at Google, and I wondered about how do we use these videoconferencing between multiple offices. You feel you're in the same room whether in you're in two neighboring buildings or I'm in Chicago and somebody else in San Francisco are conversing. And I was told there are 600 people who are on the payroll who are hacking the system on a regular basis. Whether that's a rumor or not, I don't know but that's what I was told by an outsider --

to make sure that these systems are secure. And this is just for conversations over the phone.

Imagine the level of security that's needed and the risk that many beginning companies trying to get in this field would be facing if you didn't have that infrastructure of IT security that's needed. That's a fear, a big vulnerability but that comes from the unintentional component.

I don't want to talk about manners and sloppiness but those are other issues also that you can have one company that's really advanced, they do everything in a good way and some others collect data and there's fear that that data can be used for other purposes for secondary gain, etcetera and again, you can -- in a way that may inadvertently lead to loss of protected information.

DR. AFSHARI: And I will give an example. I'm at the University of California San Diego and there is a system that you could access any medical record from another University of California campus by just going to the electronic

medical record. Well, so the broader the access, the possibility of threat is larger. So how do we dealt with that over time as the access would be anywhere, anytime for any patient, any place in the country or the world that we would have this systematic access? So Leslie.

DR. ZIMMER-GALLER: So I'm certainly not an expert in this area either. Going back to, again with diabetic retinopathy, I think certainly the programs are -- that are in place, really, by and large, images are being transferred in a HIPAA-compliant manner. They're being transferred securely and, you know, the technology is there that if you are using a Smartphone, you can take images and they can be transmitted to an electronic medical record and instantly be deleted from the imaging device. So there's a lot of technology out there. An evil-intended -- intented (ph) person probably can hack just about anything but I think the technology is there to keep medical records, to keep personal information relatively secure but I don't think it's 100

percent guarantee no matter what you do so.

DR. AFSHARI: Dimitri.

DR. AZAR: I was going to add before the blower came and --

(Laughter.)

DR. AZAR: -- a point about --

FEMALE SPEAKER: Before the hackers.

DR. AZAR: -- before the potential hackers -- it just made me forget a second important point and I don't know how to -- this whole group will address this. At some point, the retina becomes an identifier of a patient so if we're dealing with retinal images, we may be -- again, this is inadvertently but as a group, it has to be decided; it's different than a radiological image. It's -- whether it's a pathologic or normal retina, this is going to be a very difficult aspect to be dealt with and the stricter the regulation, the greater the impedance on the advancement of the technology. So I think the FDA is probably spending a lot of time thinking about this and the guidance about it is going to be a

major determinant of where these technologies will be going.

DR. AFSHARI: And then ethically, are we responsible somehow if there is pathology in the retina to diagnose it, to do something about it when it's being saved in some company for some security detection? Darius.

DR. MOSHFEGHI: So Dimitri brought up the area that I was thinking about which is biometric identifiers and it's not just the retina, it's the iris. I have cleared. I go to the airport. They have me scan my iris. It gets me through. I go to the front of the line. It's really fun but it used to be that we worried about hacks where people are going in to get your information just to expose a big healthcare corporation or something like that.

But really, now we we have two areas with these informations. For ROP, we get the iris photograph and we get the retina photographs. My database is very valuable, particularly when you're going to use facial recognition, eye

recognition, retina recognition to get into this and increasingly into your ATM account or anywhere else.

And there are two ways to go after it. One is the kind of like pickpocket approach which patients, customers, people out there, the 77 percent who have their phone, they're easily hackable. We don't all have robust software. But then that's kind of like the small bear approach if you want.

If you really want to go for robbing the Brink's trucks, you're going to go for these database systems because you can download incredible amounts of data which can then be deployed across that person's entire financial life. So that's where the concern is on my part, less on a patient privacy issue. I'm fully cognizant of privacy issues and a big advocate of them but more on the potential for mischief in other areas.

DR. AFSHARI: Okay, Mia.

DR. WOODWARD: I think this is a really

interesting discussion. Again, you know, I think that as an anterior segment specialist and again, looking at this from an urgency, you know, emergency room perspective, it think that privacy is often compromised when we're trying to deal with something fast. And so, you know, either because of privacy, we can't get the information as quickly as we needed it.

You know, a lot of solutions right now are we get a PDF dumped into our media tab of our EMR and that has the most up-to-date information and that's really not accessible for the doctor who's trying to use it, right; or a separate app that I then have to download if I want to get the information from this emergency visit. So I think the how privacy interplays with interoperability is very important. I don't have a solution for that. That's your job.

And I guess the other thing is again, about older populations and really giving away privacy and not recognizing it, I can guarantee you that, you know, a lot of older patients are on Facebook;

you know, they have their kids, they want to see their grandkids on Facebook and so they don't -- probably aren't aware of the privacy settings that they are giving away. And, you know, again, I'm not immune to that. I will not be immune to that n the future as well.

So those are my main concerns when I think about privacy is we're -- you know, are the older populations giving things away and not realizing it and that the burden's on us to make it private for them.

DR. ZIMMER-GALLER: But it's not just the older patients. Nowadays young patients have a rash and they "I don't' want to see a doctor for this, let me just send a picture to my doctor," and they'll even send it through email or, you know, any other way that they can. So patients really aren't aware of it either or don't think about it I should say.

DR. AFSHARI: And then there are some patients who ask their primary care doctor for certain things not to be written in the electronic medical

record because of the fear of the privacy, so we have that opposite aspect. There are patients who email us all of their medical records and all of us, I think, have some patients records in our iPhones; even though they are protected by some password, is that enough? So Mike

DR. TRESE: Well, that actually raises an interesting point because the rules -- one of the things that I thought was interesting this morning was that there is an international device group that's going to set standards, but I get a lot of both in-country and out-of-country requests like that. And the out-of-country requests I handle pretty easily because I don't know the rules. The in-country requests, I handle by saying, you know, you're an American citizen --

DR. AFSHARI: You are on the watch list of many countries and you don't know it.

(Laughter.)

DR. TRESE: That may be true -- but I ask -- I think the -- I handle it the way I'm supposed to for the United States but it is definitely -- it's

an important feature when we think about global telemedicine for at least rare diseases, you know.

The other thing that I wanted to bring up -- I share the concerns about the patient privacy and the idea that you could be hacked from almost any position. But I think one of the things that's a big deal in a lot of places is where you're data is stored; is it local server stored; is it cloud stored.

The iris registry, which was mentioned earlier -- I think you said you're on the board of the iris registry, is that right -- so that data is cloud is my understanding and it's -- I mean it's carefully cloud stored; its' cloud stored with the same security level that is like high security military, just one below that is what I was told. So -- but that type of consideration is a limitation for moving telemedicine some places, because people insist, universities, hospitals, sometimes insist that their data only be stored in their server.

And you lose the capability of secure access

wherever you are to patient data, like you were referencing for the people in your hospital system. So -- and that can be a big advantage for patients that you can deal with from here that have problems wherever your home practice is. So I think that's an important issue.

I don't know are any FDA regulations on storage. It that correct or no?

DR. AZAR: Don't give them ideas.

(Laughter.)

DR. TRESE: Try Denmark.

DR. AFSHARI: Leslie.

MS. BOTTORFF: Yep. So of course, GE and also a bunch of our portfolio companies are doing things that are cloud-based and we have discussions every day with, you know, integrated delivery network big systems as well as smaller entities about this very issue, because everyone's quite concerned about moving to the cloud and having it off premise. But, you know, that presumes one thing, that their on-premise security is better than Amazon web services. Really?

Probably not, you know. Not even for, you know, a big IDN necessarily but certainly not for an individual practice person. Their security is probably not nearly as good as the cloud. So, you know, everyone's working through that issue, of course.

And you're right that the big sources of data storage are a bigger target but nevertheless, if it's an easier target, then the people who are trying to get at this information are going to go for smaller -- a lot of smaller sources. So I think that's still an issue.

Some of the things that -- even though the FDA does not regulate this, there certainly regulatory entities that are regulating cybersecurity. And some of the advent of things that are coming about is there are some guidelines like the NEST guidelines, but they're pretty big-level guidelines. And one of the things that I think could be helpful, especially to small companies who, you know, I invest in and represent, is, you know, what are some of the ways that we can meet

the requirements of these guidelines in a more specific way.

And some of the things that I know are coming up are there's a number of companies working on some metrics for cybersecurity, both inside-out and outside-in type metrics. So there's a lot of activity, companies like FICO, RedSeal Happens to be one that my husband, you know, is involved with. He's a cybersecurity expert. There's a number of them though that are working on these concepts and that could the regulatory bodies that are working with this sort of come to some kind of agreement or guidelines just like they do for medical ways to monitor people; here are some ways that you could report this metric, and it's not perfect.

It doesn't, you know, cover everything but at least ways that companies that are manufacturers that are trying to produce these medical devices and services and systems can, you know, have something to go to and say, okay, we did that test and here is our metric so that you have some way

that they can tell you, you know, that they're cyber secure. I mean I know, like for instance, the MACRA guidelines which are not the FDA's but are CMS guidelines, have a component upon which bonuses, incentive bonuses are based for providing you have a certain level of cybersecurity. But they're not very specific about like what -- how do I report that to you and how do I prove that to you.

And so -- but I think we're getting there and that's something maybe that the FDA can think about in terms of working with the other regulatory bodies on your piece of, you know, those privacy concerns about how can we provide some really clear and specific guidelines or goals to the big companies and small companies alike out there of here's the kind of level you have to reach. I don't know.

DR. AFSHARI: David.

DR. MORRISON: I think that my perspective is going to be a little different, mainly because things have been covered, to a large degree, and I

can't actually define what the cloud is so I can't think I can be of much help there But I think "c," monitoring patient behavior and location is a very interesting and concerning topic. So as we move forward with telemedicine, there's a significant chance that patients can be monitored in a way that they've never been monitored, or at least that data could be collected, continuous blood sugar monitoring, IOP with glaucoma may have some ability to locate the patient to a specific geographic area to a specific behavior that could then, in some way, represent a b reach of their privacy if it were disclosed to insurers or anyone else.

Physicians certainly have been led into the concept of pay for performance in addition to outcomes measures; is it possible that insurance companies would be interested in trying to obtain any of this data to change premiums based on patient behavior and risk of having higher cost of healthcare So I think that we haven't touched on that as much but it certainly a potential for the

technology moving forward and something that I think would have to be regular and would be very challenging to deal with.

DR. AFSHARI: Bakul.

MR. PATEL: I think we should ask the patients what they care about. I think that's really the question right here. I think this question is not the role that, you know, this community played at one point in time about, you know, privacy and privacy of -- or it's preferences that people had for privacy. I think that's fundamentally challenged wit this technology.

I could argue one hand that we can secure the data that people have in hospitals all the time but then you allow them to download from and to their "MyPatient portal" and they put it on a disc somewhere or a paper somewhere and just they'll forget about it.

So I think that's sort of the realities of where we are today, the realities of being able to glean into retina images and identify a patient as things are becoming more and more ubiquitous and

sort o almost common. But so is being on Facebook at the same time. So I think the fundamental question in my mind -- I wondered about this -- is it may not be just one of us where you're able to take the fully responsibility of what we consider privacy or patient preference for that matter. I think previously in our -- even in this space, there is (inaudible) of civil rights that does HIPAA, FTC, the Federal Trade Commission talks about security as well from a commerce perspective. And just to take the example to the international front for a second and the EU, pretty much every state has their own security requirements and they're' all different.

So I don't think one universal answer sort of serves in these purposes. I think we had a tolerance at some level before I think that tolerance is changing rapidly because people are just donating. I don't know if some of you guys have heard about the non-profit organization, Tidepool. They're actually one of the pilot participants, they just advertised to help

patients who are collecting their CGM data, continuous glucose monitor data, to be donated to their website. And if you just go to their website, you will see tons of people just willingly give that data away to that website for the greater good.

So I think the concept about protecting patient data and, you know, clinicians have the responsibility, healthcare institutions have the responsibility to do so, it's being sort of questioned at this time. So that's how I would start thinking about it.

DR. AFSHARI: And the patient's desire, depending on the data point, may be different and each patient may be different, or there are ways to study the identified data like dbGaP we have for -- from an NIH, that whole genetics database.

So are there comments or questions from the audience? I know we are getting very close -- yes.

MALE SPEAKER: (Off mic) ophthalmologist (off mic). So thank you for this conversation. This

is actually what I believed for the last five years trying to get a single device (inaudible/off mic) -- so when I deal with the vendors, the vendors put their hands up and say, "It's your problem, it's not our problem" and I think we, as a community, have to recognize that it is our problem, for ophthalmology in general. In fact, radiology is much better at governing and pressing on the vendors to improve their (inaudible) posture. In ophthalmology, we know the vendors have (inaudible) whenever they want.

So a couple comments about the cloud. So there is an accreditation-certification process. It's (Inaudible/office mic). Now it's government accreditation. I think there are (inaudible) out there that are a part of that, Microsoft, Amazon and I think Google and that's it. So Amazon has one part (off mic) so have to make sure when you're talking to your vendors or you (inaudible) "yes," which is part of the design. And I (inaudible/off mic) who use the cloud services are not (Inaudible) certified.

So if you're going to place (inaudible/off mic) on a system or (inaudible/off mic) on a system that's not (Inaudible) certified, just understand that there's risks there. (Inaudible) startups themselves who work with the (Inaudible) rule (inaudible) or PHI so if you're going to leak a data, that is (inaudible). So if you're not going to use a (Inaudible) certified system, you lose data. There's potential risk (inaudible/off mic).

I could use the help from the DoD because I can't use any of these new emerging technologies unless it's accredited. I can't get it accredited if we won't have (inaudible) efforts (inaudible). So I just want to thank everyone for talking about this.

DR. AFSHARI: Great point. Comments from the panel or from the audience? Now I believe everybody is ready -- yes, please.

MALE SPEAKER: (Off mic) question back to the first one (off mic.) My name is (Inaudible/off) mic, very small startup so this is very

appropriate for us (inaudible) discussions. (Inaudible) I've heard a lot this morning about physicians really (inaudible) the gold standard. And if you looked at gold standards as (inaudible), gold standard seems like it's the physician. And then I heard -- we went all the way back through well, our physicians, most of us, is a black box. And so we're back to this question of well, how do we then -- you know, what is the standard or how do we test or, you know, what is enough data for us to have, you know, a clearance for a medical device? So I don't know if the panel could help address that in terms of, you know, how much data do we need in -- you know, in order to get our clearances for a screening device in particular, something that is going to be an adjunct to the physician, not to replace, not to be viewed as a diagnosis but just an adjunct? Any thoughts on that?

DR. AFSHARI: Thoughts on that -- and of course, it depends on the device. I'm sure we are going to hear that. Michael.

DR. TRESE: I think we're really close to the end of time. I know that Mike made a big point about that. I'd be happy to talk to you about that afterwards because there is a lot of info -- or was that Mike to go ahead and answer it?

DR. CHIANG: I think we should answer his question.

DR. TRESE: Okay. The -- I think that there is no real gold standard of to what a gold standard is. It comes in terms of physician acceptance. In my mind, someone that does a drawing of a retina for retinopathy of prematurity might as well hand me a cartoon about peanuts or something like that, because I can take a picture and get an exact image, okay. That's not the mentality of everybody that does ROP. And some people think that there may be reasons to draw something that may be wider even though you can image really from aura to aura with some of the cameras now. So -- but I think it is -- in all parts of medicine, it seems to me it takes to me it takes a really long time to get to -- things to

change and a gold standard to emerge.

So years ago, Bill Rich (ph) told me that there was some data that suggested -- this was maybe five-six years ago -- that it took almost 20 years for clinicians to accept randomized controlled prospective data. And so I think it's hard to determine. I think the black box is opening. I think the gold standard is evolving Bureau you could get arguments on both sides, I think, of diabetes and ROP for sure.

DR. AFSHARI: Michael Goldbaum and Bakul.

MR. PATEL: Oh, I'll just make a point. I think it all depends on the device -- the device that we're looking at specifically but I think it really -- and this is the way I think about it. It is what you really want it to be and it's not about what we want as an FDA review product, right. So it has to match the claim that the evidence is there for. So it's not about what it is. It is about where you want to be as a product and we can talk in detail about that and Malvina and the team is there to think about -- think

through that but it's about matching that. In my mind, that's really where we need to sort of get to and not -- it's not about you need 50 patients or 500. That's really not the discussion.

DR. AFSHARI: Mike Goldbaum.

MR. GOLDBAUM: Yeah. So I'd like to pretend that the physician is not a black box. So the AI is a (inaudible) -- you have to come up with some techniques to try to figure out how it got to where it. And there's a bunch of (off mic) inside this device, and you can't really figure it out. They're all (inaudible). You can talk to a physician. You can get the reason. You may not agree with it. And there's (inaudible) if the reasoning is wrong. You can teach them the correct reason. So there is -- I think you -- to say the physician is a black box is more than (inaudible/off mic).

The other thing is that the seclusion is one of the methods to try to (inaudible/off mic) machine learning classifier got to where it got and got to -- and it's a variant factor of

elimination. I mean it's an old technique. These are a lot of techniques to try to figure out what's going on inside a machine unclassified and in factor of elimination is one -- it's -- and it's --actually, it's quite effective. And as you can see, (inaudible/off mic).

DR. TRESE: One more.

DR. AFSHARI: Comments? Dimitri.

DR. AZAR: Can I comment about this because the black box issue was raised by me. I was actually agreeing with you that the worry that people have now is about the black box not being trusted but in fact, we need to be focusing for the future that the question should not be about physicians. We need to establish the trust in physicians and that requires recognition of that potential risk and preparing our doctors today to be -- for that future that if we don't, they will be questioned. Trust in the machine, I dread the day if the trust in a machine is greater than trust in your physician.

MR. GOLDBAUM: As an example, the testing for

colleges, the scores went up when (inaudible). The scores went up as we went on the in years and the reason why the scores went up is that the young people were using the internet to teach themselves a lot more than people of previous generations were able to do so that -- so my point is that all of these tools can help us to become better at what we do so that we don't need to be abandoned and replaced by AI.

(Laughter.)

DR. AFSHARI: So Leslie, the last comment.

MS. BOTTORFF: I just had one more comment because I just want to say to the folks of the FDA that the pre-certification process and the collaborative, you know, things that you suggested are really the way to go to make this thing go fast. And I just congratulate you on that and say, you know, that needs to speed ahead. And it might not be perfect the first time around. Just get it going because the data is going to help it. And so congratulations on that. It's a fantastic concept.

DR. AFSHARI: So on that topic of the black box, we have boxed lunches outside. Please come and join us so we can talk about privacy or less of it this day and age. Thank you to all of the panelists, Dr. Trese. Thank you for a great morning and session.

(Chorus of thank yous.)

(Applause.)

DR. BLUMENKRANZ: Afternoon. Hopefully, everyone had a good lunch and you didn't carbohydrate load so much that we'll all be sleepy here. But we're already a few minutes behind so we'll try to catch up, and we're going to be doing a panel two. This is safety and efficacy concerns for ophthalmology digital devices in differing settings and there's an emphasis on differing use settings, and there's an emphasis on differing use settings. The same tool might be more or less useful in the clinic versus in a primary care setting or in the workplace.

And I thought we ask the -- I'm Mark Blumenkranz, and I'm the immediate past Chairman

of the Department of Ophthalmology at Stanford, and I've co-founded the Optomic Innovation program with David Myung, and I am a a practicing vitreoretinal surgeon, and I have worked in technology development in Silicon Valley.

I'll ask everyone to introduce themselves and then we'll move on to the presentations and questions and so forth, and this is my co-moderator.

DR. NISCHAL: Thank you. I'm Ken Nischal. I'm representing the American Academy of Pediatrics and Section of Ophthalmology. I'm Professor of Pediatric Ophthalmology at Pittsburgh. We've been involved with a lot of global health in terms of telemedicine, not just for ophthalmology but also for cardiology and hepatology, and we provide real time ICU surveillance for several centers in South America, and it's sort of a pivotal role of what we do in terms of tele medicine.

MR. OSWALD: Good afternoon. My name is Quinton Oswald. I traveled a long route in

ophthalmology. Currently, I'm the the CEO of a company called Notal Vision, which is a device-based platform which we'll talk about briefly in a second.

Right from the the existence of vision died through the launch of Lucentis were experiences I had, and now I moved to the light side. Dealing with the side of the business has been a fascinating change in my life and certainly the way that I think about measurement of the patients.

DR. BODNAR: Good afternoon. I'm Zach Bodnar. I am originally a software engineer and I transitioned to become a physician. I've completed my residency in ophthalmology at St. Louis University and over the past year, I did the ophthalmology innovation fellowship under the tutelage of Dr. Blumenkranz and David Myung, and I'm currently a vitreoretinal surgery fellow at Stanford.

MR. PATEL: Bakul Patel. I'm Associate Center Director for Digital Health at CDRH.

DR. ZANGWILL: Linda Zangwill at UCSD. I'm an epidemiologist by training and Director of the Diagnostic Imaging Laboratory and the Hamilton Glaucoma Center.

DR. GOLDBAUM: I'm Michael Goldbaum. I'm also at the University of California San Diego, and I do -- I practice in retina but my research, a lot of it, is also in glaucoma. And I've had an interest as far back as 1987 in imaging and machine learning. So that's where the original work came from.

DR. ABRAMOFF: Yeah, I think we're the oldies. I'm Michael Abramoff. I'm -- I was trained in the Netherlands as a vitreoretinal surgeon, came to the U.S. 15 years ago. I'm a Professor of Ophthalmology at Iowa. I'm also the founder and president of IDx which is getting ready to put its first submission for an automated diabetic retinopathy detection device for primary care into the FDA.

DR. CHIANG: I'm Michael Chang. I'm a pediatric ophthalmologist by background and also

board certified in clinical informatics. And I run a research group at Oregon Health and Science University that deals with various aspects of applying information technologies to eyecare.

(Off record comments.)

DR. BLUMENKRANZ: Good. So this is our listing. So I thought by way of introduction, it might be worth just talking a little bit about the fact that this is an area that benefits all constituencies. Digital health is -- has a value proposition for patients in terms of engagement, better care, and convenience for physicians in terms of workflow, expanded reach, engagement and research for payers, in terms of, hopefully -- I won't say reducing costs but appropriating cost and better outcomes, and finally, for industry and pharma and the device industry for data-based insights and value based analysis.

I think our -- there we go -- it's interesting, I think, to get to look at the evolution of the Smartphone and as you saw earlier today, although 75 to 80 percent of Americans have

Smartphones, it's true even in the 65 and plus category, in which it's now approaching 30 percent in 2016, so it's -- they are becoming ubiquitous across all segments of society although maybe a little more prevalent on the coasts.

Smartphone capabilities are increasing at a similarly rapid rate if you compare the original Mac in 1984 to existing iPhone or Android, you can see the -- whether it's the pitch or the DP, whether it's the degree of memory in terms of images, in terms of the speed, and in terms of overall memory, they're all dramatically greater for small Smartphones now than they were for the very best Macs that were available at that time. And we can expect that in keeping with Moore's laws to continue.

Smart devices are now being used across all of all of medicine, not just ophthalmology but whether it's asthma, cardiovascular, ENT, oncology, diabetic management, they are becoming an accepted standard of practice and I think this workshop today, it just goes a long way in terms

of bringing ophthalmology into the forefront of this area.

These are just a few of the companies that I am aware of as of last year that are already offering products, some approved some not, that are being used by inpatients are being put through clinical trials to be able to seek approval. If you look at it just in terms of cameras alone, there are more than seven cameras that are either approved or in the process of being approved to take advantage of the degree of the high resolution camera found in Android platforms and in iOS devices that, coupled with some sort of an optical device, can produce very high quality images.

The FDA has cleared -- as of 2014, the FDA had already cleared more than 100 mobile health apps for medical use and that's been increasing at an increasingly rapid rate. If you look at it from the total market perspective, although the amount spent in the last few years is still significant, it's expected to grow by more than 6x, according

to market research, between 2015 and 2020. So it's certainly a very important area of the medical economy as well as the value-based creation for patients and for physicians and for the industry as a whole.

With that I'd like to go on to our questions. Now our questions really revolve around safety and effectiveness and in specific situations. So the first question is for Quinton Oswald and that is, What are the important safety and effectiveness concerns for an ophthalmic digital health device for the screening or monitoring of progression of macular diseases?

MR. OSWALD: Thank you, Mark. I'd like to handle the effectiveness piece first because safety, I think, is a relative issue in the device space. But in reality, I think insufficient focus has been put on the effectiveness and the clinical utility of these devices as we go forward.

And for example, before we're able to get reimbursement at Notal Vision for our device, we had to involve ourselves in quite an extensive

clinical study to prove the relevance of our device in terms of a reimbursement environment. So, therefore, we think that that clinical utility and the evidence based of that clinical utility is a critical element in the development of the device situation.

So I think that the bar for approval of a device is relatively simple today. We think that maybe we need to expand that from a point of view of the clinical evidence that supports that, and being, you know, well-controlled clinical trials as we would thought of in the drug space. As I said, with regard to safety in particular area, it's benign by virtue of the fact we have no real impact on the eye, so that's not something we certainly think about.

DR. BLUMENKRANZ: Linda, perhaps you could address that with regard to glaucoma to the extent that there are similarities or differences between macular disease and glaucoma relating to the particular type of pathogenesis of those diseases.

DR. ZANGWILL: Well, I think there are a lot

of similarities and it's really the attended -- intended use is what's critical here. And in terms of, for example, screening for glaucoma, I think detecting the earliest disease is not going to be effective or efficacious, but treating -- detecting moderate disease that we can do well -- with early disease, there is -- the clinicians disagree and the machine learning will disagree, but for moderate disease, we are fortunate that glaucoma is, as mentioned earlier, slow progressing. So if we can detect moderate disease, that would be -- and I think we have tools that we can do that -- is -- would be a target for specifically screening or advanced disease. Lots of advanced disease is undetected.

And in terms of the safety as well, if you're thinking about the safety of the visual field test or a photograph or something, the safety issues are not that difficult to deal with in terms of safety in a similar way to to the macular disease.

MR OSWALD: (Inaudible.)

MR. OSWALD: I think it does fundamentally

different medical diseases versus glaucoma relative to the time of onset, because we know that the earlier you get a patient in a switch from dry to wet, the more -- the better the outcome is going to be for the patient in terms of treatment. If you can get them with a relatively good vision and small lesion or fibrosis, the outcome for the patient is going to be superior. So, yeah, different from glaucoma, we think that's a really critical element.

DR. GOLDBAUM: If I may add something? As well as diagnosis for glaucoma, diagnosis is important but what the glaucoma clinician needs day-to-day is to determine whether the disease is stable or progressing. And we have found that that machine learning classifiers or a hybrid system using machine learning classification is quite good at picking up and detecting progression.

DR. BLUMENKRANZ: I just want to remind the fact that the panelists that -- and the audience as well, please, before -- if I haven't called on

you, if you could just state your name so the transcriptionist will be able to do their jobs in terms of having the proceedings of this publishable.

DR. BODNAR: I have something to add to that. This is Zach Bodnar. I think one of the differences with glaucoma and macular disease, glaucoma, it depends on the type of testing you're doing but if you're talking about visual field testing, one of the big issues is the reliability and the reproducibility of the test. And there are some questions about an increase in the variability if a patient is testing their visual field at home. Are they going to do it in the same setting where there's the same background luminance each time or even across different patient populations?

So some of those things need to be validated. And I think there's a risk that in some of these initial projects, we'll see kind of a negative result where we find they're not useful, but that may just be because there's a lot of noise and we

haven't taken into account the amount of variability that's in these other environments.

DR. NISCHAL: Before we go on to the next question, Linda, I wanted to ask you when we talk about safety with glaucoma, you know, there's the issue of self-monitoring with some of the companies that give you the self-monitoring equipment, and there's the issue of implanting a device in the eye that gives you constant monitoring. You know, if there is that kind of. implantation of a device, that must have some safety issues.

DR. ZANGWILL: Yes. You know, that that's a good point and I was not addressing -- I was focusing more on the screening but that's a good point. I believe the -- one of the implants, IOP monitoring, has just been approved by the FDA for continuous monitoring. So that is absolutely a safety issue and my understanding is it went through the proper -- the -- it has been approved but there are safety concerns. It really depends on the intended use, if it's screening, if it's

monitoring pressure, etcetera, that is but that is -- that's a very good point.

DR. BLUMENKRANZ: Just for point of clarification, Bakul, was it -- I know there was an external device, the SENSIMED or -- but I don't -- was there an internal implant as well or --

DR. ZANGWILL: Oh, go ahead. My terminology --

MALE SPEAKER: Yes. Actually, I'm going to see if Ron is in the room to answer that. He's probably better-suited because his branch and his division sort of looks at that.

DR. SCHUCHARD: This is Ron. As far as I know, the SENSIMED is what's being referred to, and that is a device that is monitoring relative change rather than absolute IOP.

DR. BLUMENKRANZ: I think also the point being the risk-benefit ratio. I think that the physical risk that is there -- we'll set aside the risk of getting the numbers wrong or things being imprecise or not reproducible. But there is there is a fundamental difference between a surgical procedure to implant a device in the eye and one

to monitor externally in terms of the risk profile. And I think that was the only reason I asked the question, was because I think you might value them -- evaluate those in a different way fundamentally related to the risk-benefit ratio, not to the accuracy or reproducibility per se.

DR. ABRAMOFF: Mike Abramoff. I wanted to say something about diabetic retinopathy where the presence of preferred practice patterns. For example the American Telemedicine Association's guidelines for DR detection have been invaluable, so if you have these guidelines to align with, and we prefer a practice plans to align with, that makes it much easier. All these debates about which level of diabetic retinopathy to detect, you know, what to do what the different levels, that has sort of been decided already by a professional society. So -- and, you know, if we're looking at glaucoma, it's a little bit more tricky just because the guidelines are a bit more vague so, you know, professional societies can really help by making, you know, the guidelines as specific as

they can be.

DR. NISCHAL: Okay. So we're going to move on now and one of the other questions that we have to cover is alluding to what Zach was saying, you know, do the use of these digital applications change depending on the environment that you're using them in. So we're going to start off with the use in an eyecare clinical environment.

And the first question, really, I want to discuss is the fact that if you -- do we need to specifically train somebody in the office to have the responsibility to look after the data that comes in? This is a real live, practical problem at Children's Hospital of Pittsburgh. We have a portal for patients to be able to contact us, and I sometimes find it difficult to return emails on the same day within a couple of hours rather than having a number of patients go home and then send me a message. So we've actually trained somebody who is -- was an ophthalmic tech, to take the responsibility to look at the portal messages of all the attendings. That's one question and the -

- the one thing to discuss.

And the other thing that's quite important is that in my role with Telehealth in helping develop surveillance for pretty sick children in another part of the world, if something goes wrong with the data that you're getting, you need someone to recognize that it's a data acquisition problem and not a problem with the data itself. And so I'm becoming more and more concerned that rather than having somebody who's just an ophthalmologist or an eye healthcare person who we may give the responsibility to do this, that you need someone who has adequate training in IT and software so that they can then assess whether the data they're guessing is wrong or indicating that there's a problem with the patient or whether there's a problem with the acquisition.

DR. BLUMENKRANZ: So I'd be interested to see what the panel says about that. Michael, you've done a lot of work with Telehealth. What do you think about the specific roles?

DR. CHIANG: Can I have a comment about the

first thing that you mentioned in terms of that monitoring? So this is in response to your comment, not an answer to your question. I think that we have a fundamental model in clinical medicine which is that we will see a patient every "x" number of months; you know, every three months we'll see them and we'll say, you know, take your eye drops, take your blood pressure medicine. We'll come back three months later and then they either have done it or haven't done it, and we'll repeat that cycle like every three to six months/ With a lot of these Telehealth things that we're talking about in this panel. we're sort of fundamentally changing that model where the idea is that instead of every three months, you go do this at home and you'll do it every day or every week, and you tell us or, you know, somehow we've got to figure out when there's a problem with that.

And so I think that number one, I think one of our challenges, as a community, is to demonstrate that that second Telehealth model provides added

value over doing it every three months. And I'm not convinced that we've answered that question, but I hope that's something that we can all kind of go back and tackle, you know a little bit more, because you know, Ken, for the reason that you're mentioning, you gave an example, I think, of how that might hurt us because you can get bad data and we may make the wrong decisions based on something that came up. And so I hope that's something that we can, you know, address a little bit as a community.

In terms of -- Ken, in terms of the question that you asked me, I think that the way that I sort of put these technologies together is sort of who is responsible for interpreting the data and who is going to be that decision maker. And it seems like there are three options; one of them is that the managing ophthalmologist takes that data and then they make a decision on it, which in that case, it's a decision support tool and so that seems easy.

Number two is that like it's basically a --

some sort of remote reading center where, you know, some expert remotely, you know, makes that diagnosis and you've got to have faith that that was done correctly; in other words, it's a black box kind of approach and I think that there are some safety issues with that. Who is doing the readings; what's the validation?

And and I think the third issue is that the system does it automatically, which I think is a whole 'nother level of of oversight that that requires. So I think that in a lot of cases, how we deal with them is got to depend a little bit on how those systems are architected and what that sort of model is.

DR. BLUMENKRANZ: Let me let me push back to you a little bit, Michael. It came up in the course of we did some work on home monitoring for macular degeneration with patients generating visual plots and data that are then able to be shared with in a secure and private way the physician. But when -- in talking to other physicians, they say we're already horribly

burdened by information overload. I'm getting emails, questions from patients through a portal. All of a sudden somebody sends me an alert that says that maybe a vision is dropped; I'm not checking that every day amongst all of my other email accounts, my other correspondence, and so forth.

Question comes up; does it put a burden on physicians if they don't have some procedural process that -- so that somebody actually is looking at these results every day, even if it's not a physician, if it's a paramedical personnel and so forth? I mean what are your thoughts about that.

DR. CHIANG: Mark, my first thought is that I think that's a really good point, you know, that you made. You know, I think that doctors feel pushed to do more in less time, you know, often with fewer resources than ever before and I think that's an example of that. And so I guess my first thought is that if the data is going to be generated, there has to be somebody who is

available to interpret that data and ideally, it would mean that there's somebody who can get paid or whose job it is to interpret that data and, you know, get paid for doing it.

You know, I think one of the challenges is that if we get patients who are generating more and more data and you've got 2,000 patients in your panel, that's an enormous amount of data that's going to be coming back, and I don't think any human can do that. And to me, the natural extension of that is that we need to develop automated systems to do exactly what you're saying, which is to screen at what point does a doctor need to get involved. And so I think with those issues, the challenge is going to be number one, how good are those systems and how good can we make them at distinguishing the sort of normal from the not normal.

And the second is what is our threshold as healthcare providers for being bothered by the system. In other words, do you set your threshold here or here or here, you know, to get bothered a

lot versus a little based on your tolerance for potential errors. And so those are sort of the frontiers that I see that we'll need to address to try to deal with this question of basically what I would call data overload.

MR. PATEL: So can --

MALE SPEAKER: Yes, please.

MR. PATEL: I was just going to make a point. I think this discussion is fascinating for me because I think we're really talking about data prominences, like what -- where is the data coming from, right? So -- and this is the cross-walk of interoperability where things can -- things -- devices and other products would have to sort of declare itself the performance, not just put it in a label someplace but actually be that -- the data performance and the validity of it should travel with the data stream itself. So that's one issue.

I think then we talked about the data volume itself, like how do we sort of take care of data volume. And that's where -- so in the discussion we had earlier, thought comes into play where

automation is probably the only solution that sort of exists in terms of taking in -- taking data volume into insights. And that data insight is really where we are talking about does it really mean meet the screening threshold. Does it really meet our sensitivity-specificity sort of threshold? And how does it actually aid?

So if we take this string of requirements, so to speak, you start with what generating; how transparent it is; how does it report back so people -- either it's man or machine interpreting the data, knows what kind of data they're interpreting. And then you walk down to the next step of how do you take that and turn that into insights or information that can be used in practice. I think that scale will be something that we need to evolve.

FDA just put out or are in the process of putting out the interoperability guidance which talked exactly about that. It's about how do we sort of get people who are generating this information and data streams to be very

transparent. So it doesn't really matter whether man or a machine sort of uses that data. They know how much to trust it, what -- how much to rely on it. So we should think about that even in this space.

DR. NISCHAL: So now that we've sort of discussed that, you know, we either have somebody who has a a specific role where you go for automation, the question really is if, Mike -- Michael -- going back to what Michael was saying, if you're going to have a physician look at this data, how do we reimburse them, because I can tell you right now if we do a tele ROP screen for Inaudible), the middle of Pennsylvania from Pittsburgh, it's been a real struggle for my doc to get the appropriate reimbursement either from the insurance or from the hospital that wants to do it. So the question, really, that I wanted Quinton Oswald to track is how do we -- should we tackle that question now and, you know, how do we do it.

MR. OSWALD: Thanks, Ken. In 1974, the

Harvard Business Review published a one-page article entitled, "The folly of incenting x whilst was expecting a y." and really, we have to think about the alignment of incentives through this process. And you think about the ophthalmic space, the reimbursement environment and the rapid change of telemedicine/telehealth are traveling to different paths. And today we're having difficulty aligning those incentives. And we really need to set up a platform and a process where we think about how to do this so we're not expecting x whilst incenting y.

DR. GOLDBAUM: So I think what we're hoping is that computers will help physicians or healthcare providers to be able to do more with less time. And so if I understand Michael Chiang correctly, then his answer to nagware is triageware and that may be the answer actually, because then the physician is presented with what's important and can concentrate on either the patient or other parts of the patient care.

DR. BLUMENKRANZ: Okay. Can you turn the

slide back on? We're going to move from the ophthalmology office or the optometric office to another side of medical service; in this case, let's just say hypothetically either a primary care center or the emergency room where we're using telehealth and digital tools to try to expand the reach and improve efficiency and outcomes.

What experiences do we have now for interfacing between eye health professionals and primary and urgent care providers? Michael, I'm going to start this off with you, and then I'm going to move to -- I'm sorry -- Michael Goldbaum. We have three Michaels on this panel. I don't know what that means.

DR. GOLDBAUM: And we're all right next to each other.

DR. BLUMENKRANZ: It must mean something. I think we need to do deep learning to figure out.

DR. GOLDBAUM: So this is the Michael cluster over here.

DR. BLUMENKRANZ: Yeah. Well, some names are

in vogue like Ambrose at the turn of the century but is -- you don't hear much anymore. So we'll start with Mike Goldbaum since you've got the mic, and then we want to move to Michael.

DR. GOLDBAUM: So we can go to the --

DR. BLUMENKRANZ: Yeah.

DR. GOLDBAUM: -- communication or -- communication slides.

DR. BLUMENKRANZ: And while you're doing that, would each of the kindly brought along some visual materials and so please feel free to bring those up when I -- when we ask you to be able to answer the questions. And we'll -- if we have to work through them a little, hopefully, that'll be worth it.

DR. GOLDBAUM: I think we're starting at 46?

MALE SPEAKER: Yeah. Mike's at -- Mike starts at --

MALE SPEAKER: (Inaudible) --

DR. BLUMENKRANZ: Why don't you hoof it while he's looking for it here.

DR. GOLDBAUM: Well, anyway, I can start by

saying --

(Off record comments/adjusting slides.)

DR. GOLDBAUM: It's labeled. It says "FDA workshop" and it says "interface" on the -- "interface between eyecare."

And anyway, so the goal of interface is to overcome incommunicable silos in medical records and I think what you're talking about is, really, peer-to-peer communication, not necessarily eyecare to non-eyecare and it's a generalized problem. And there are methods of communication or hard copy with letters and -- or a patient can carry information in either paper or a thumb drive, and that helps to overcome the HIPAA because you don't have to get permission, the patient's already got the data with them in hospital consult and -- and one of the things that's good about a hard copy is that you can -- if it's in another language, you can use something like "Google Translate" to translate for you.

Phone calls requires that somebody be there to take the phone call but its benefit is that it's

interactive and you have proof of receipt that you know that the person getting that information -- who needs to get the information has gotten it. And you can send messages by phone which is invariant to time, place, or geography and that -- but it can be interactive if you get somebody who's responsive at the time.

DR. ABRAMOFF: So it's encrypted or secure?

DR. GOLDBAUM: Yes. Well, I think if you're using cell phones, it's -- well, I don't know how the digital -- I think that's person-to-person and not -- and reasonably secure. But if you send a message, there's no proof of receipt. Email also invariant to time, place, and geography. You need to use a secure system for that. Electronic medical records, you have professional-to-professional notes within the electronic medical record system with -- or you can autopopulate a report which take -- which is not as time consuming. I mean you make your report and it automatically populates the report for you so you don't have to spend the time doing it. And with

the EMR, you have a holistic view of the patient. And the social networks, not very good for communication about a patient but a good way to communicate to distribute knowledge.

So in summary, we have hard copy, cell phone, email, EMR, and social networks as ways that peers can communicate.

DR. BLUMENKRANZ: Thank you. I'm going to move it over to Michael Abramoff now and just -- he's starting at slide 25. Okay, thanks.

DR. ABRAMOFF: Since I made the slide, I want to show at least one. And the rest we'll just forget about it. But as I mentioned, we finished this -- sorry -- I will show only one, this one because it sort of sets the context. And so remember, we just finished the clinical trial for automated detection of diabetic retinopathy, not FDA cleared; we don't know where it will go but I'm just saying.

And so one of the sights for the trial was in New Mexico, close to the Mexican border where there's no ophthalmologists and, in only four

hours, we install a DRAI system there. And then the question is, you know, the family physicians don't have experience with EDT arrests or ICDR, and so, you know, the question is should it be actionable for this primary care physician because they don't have any experience or context of knowledge of what to do with a patient with a certain level of diabetic retinopathy. So I think it should be you know very much dependent on the context, but at least in primary care, it needs to be, you know, very actionable rather than some abstract disease level that they need to look up.

DR. GOLDBAUM: Presumably, if you use a dichotomous system, either refer or not refer, or refer with urgency, then you can cover that. I mean that -- I guess the question is how do you set up the workflow? What have you done? What have you seen done that works to sort of simplify that so that they don't have to make a judgment about whether level 35 or 43 retinopathy is or isn't?

DR. ABRAMOFF: Well, other slides will show it

but we won't show it because it takes too long, but, you know, the preferred practice patter from the American Academy of Ophthalmology has very clear levels of which is actionable by, you know, maybe see earlier, which needs treatment, which needs close management, like you know follow-up much sooner. And so you try to stick with those preferred practice patterns if they exist. If they don't exist, you know, it's much harder to automate any of this

DR. GOLDBAUM: But in other words -- but in terms of producing that report so that a busy primary care physician or any other physician knows what to do, I mean isn't -- shouldn't we be making that part of the workflow and --

DR. ABRAMOFF: So that's my answer. So, yeah, you have, you know, for example, the output is more than mild diabetic retinopathy and/or macular edema, refer patient and --

DR. GOLDBAUM: Right.

DR. ABRAMOFF: -- that's the output.

DR. GOLDBAUM: Right, perfect.

DR. ABRAMOFF: But again, not clear.

DR. NISCHAL: So there will be time for questions once we finish the environment section. We got one more environment section to go to, but it seems to me what you're saying is that protocolization is actually the first step, that the more protocolized we are in medicine, the easier it's going to be for AI devices to follow those algorithms that our professional bodies have created.

DR. ABRAMOFF: Absolutely. It's so much more difficult for glaucoma screening let alone AMD screening where we -- you know, there are so many rules for when you should screen and not, it needs to be a treatment, it needs to be efficacy, it needs to be equitable; there's so many, you know, rules for when you should screen and when not. But it definitely helps for someone making an automated system or an algorithm or a black box or whatever it is, an AI system that there is something that you can guide yourself by. And gold standards, we have been discussing this at

length this morning, are really important. The more (inaudible) on a gold standard, the easier it is because then you know what to choose. And similarly for what the output should be like and how it should fit in our healthcare system, it's so much easier.

DR. BLUMENKRANZ: So, Michael Goldbaum and then we go to Michael Chiang.

DR. GOLDBAUM: So we're talking about gold standards and we're talking about a lot. Is gold standards where the physician or the expert is the gold standard and you're trying to make the classifier approach what the physician is doing, but it can never be better than what the physician is doing if you use that. So another gold standard -- so that's expert-driven.

Another way of having it would be outcomes-driven, so you can follow patients and see if they had a certain outcome over time that says that that patient needed to be referred and you use that. That doesn't require physician or expert input. What you're doing is you're looking at the

outcome and then you're teaching based on the outcome.

And so the machine learning classifier may become better than the expert at determining which ones need to be referred.

DR. BLUMENKRANZ: Michael Chiang.

DR. CHIANG: Ken, I wanted to comment on the idea of protocols that you and Michael Abramoff talked about, because I think that they're really important and I'm a big -- part of my career is based on developing and implementing, you know, protocols.

But I wanted to talk about the limitations of protocols, because I just want to get that on the record here. In ROP, we had done some studies -- and I want to use an example -- where they're very, very clear protocols based on tens of millions of dollars of NIH money about who gets treated and who doesn't. You know, we studied who gets treated and who doesn't and about 10 percent of the time, the kids who got treated were treated outside the protocols, and it was not because the

treating doctors were not aware of the protocols. Something made them nervous that didn't fall within the protocols.

And that's where I would consider the art of medicine. So you've got the science which is the protocols, and the art which is sort of clinical judgment and what makes someone nervous. And so I wanted to say that, you know, just because we've got a protocol and we've got a machine that can do that doesn't necessarily mean that there's no role for the doctor. And I hope that that's something that we can consider as a community, that there still is a doctor to interpret what we're seeing and make their own individual sort of a clinical judgment, and the systems are tools to help the doctors do that.

DR. NISCHAL: Michael, I totally agree with you but I think what's really important is that if you have protocols, when you look at the protocol deviations, you then look at the outcomes that Michael was talking -- Michael One -- One, Two and Three, okay -- Michael One was talking about. And

and I think those protocol deviations actually may give us more information than the actual people, the ones where you followed the protocol.

And -- but the reason why I'm raising this is I'm glad that you're you're invested in protocols, because having come from Europe six years ago, it was a real mountain to climb to convince my attending colleagues to put together protocols, because they felt it was taking away their autonomy, and I don't think it does. I think it actually protects you. It still allows you deviations but you can now quantify and qualify those deviations to look at outcomes.

DR. CHIANG: Ken, I completely agree. And just a couple follow-ups to that. One of them is that this particular study that we did was looking at real, you know, investigators who were really intimately familiar with those protocols. And so I think that's different from the population of real world ophthalmologists who the protocols are intended to target, you know, to standardize to standardized care.

You know, the second thing is that earlier this morning, Krishna made a great comment about the learning healthcare system. and I hope that this is something that -- and, you know, I think that that's relevant, Ken, to your comment -- that hopefully, you know, we're going to have a situation where the, I'll call them errors or protocol deviations or something, can maybe feed back into the protocols to say well, what was different about this that made the expert nervous and how can we develop a better protocol based on that constant feedback.

And where this ties into the regulatory thing is that I hope that in developing the rules for this that we don't have a system where it takes like a year to refine the algorithm if somebody comes up with something new, because some of these data may be generated in real time. And, you know, hopefully we can come up with a rule that lets these systems get better as they learn more from the data, because I think, in a way, that was the whole point of the National academy of

Medicine Report, the learning healthcare system.

DR. GOLDBAUM: Well, this is the point where regulation steps in because now you've come up with something that looks like it makes the system better and you have to retest it for so many millions of dollars or can you just make an adjustment. And so that's a question for the FDA to answer

DR. ABRAMOFF: I think Ingrid Zimmer-Galler -- I don't see her right now but she (inaudible) this morning, right, that even after approval, you need some form of continuous monitoring or, you know, whatever, post-market surveillance, to make sure your system, a, is indeed doing what it should do, and if there is are exceptions like Michael said, find them and try to improve your algorithm. I've done the regulatory. You would then go back to the FDA say, hey, you know, we approved it because you need to prove that but yeah, definitely that will be very important for these systems.

DR. BLUMENKRANZ: Could you go back to slide 14 for me, please? I'm going to go back to

Michael Chiang for this next question which is we're going to talk about the non-clinical environment. We spoke about the ophthalmology office. We've spoken about other primary care and urgent care settings. Now we're going to talk about testing in the home and what unique sorts of concerns and considerations we might have in that environment.

So I'm going to start with you and ask is symptom diagnosis and triage analysis safely left to the potential patient, or does it rest with someone else? In other words, where is the responsibility, both ethically and also legally?

DR. CHIANG: Yeah. Mark, it's -- I -- you know, one of the things I do is I teach user interfaces to grad students in Oregon. One of the things that we sort of use is, you know, could your grandmother use this system. And, you know, I think that, for example, in glaucoma, since that was the example, we've got evidence that patients cannot put eye drops in their own eyes let alone use the system. So I don't know how they're going

to perform when they're asked to do -- you know, Bakul, the point that you made -- doing home visual field testing. And I think this is a huge problem. You know, it's what the engineers will call, you know, quote, "garbage in, garbage out."

And, you know, I think, Bakul, your point was some method for assigning quality of data. And, you know, we do this for visual fields where there are metrics to assign the quality of that field. And I hope that we can come up with something to assess the quality of data that are obtained outside the clinical environment, because I think that ties into the issue of, you know, as the doctor or, you know, somebody is going to have to review all these data and then figure out is this going to be my trigger point for taking some action.

And the whole purpose of these systems is to save us -- you know, to lead to better outcomes, to save money for the healthcare system. And I can think of some scenarios where, you know, where these systems could have an unintended consequence

of bringing more patients to the office because of bad tests and there's going to be a way of distinguishing. We've got to have some way of distinguishing, I think, good versus bad. And I think it's completely solvable, you know, just as long as we think about that and, you know, figure out in advance.

DR. GOLDBAUM: On the other hand, I think one of the things that the Kaggle competition showed is how these systems work with bad data, because a lot of the images in the Kaggle competition were atrocious and some of them were very good. And your system had to learn on the whole complex, the whole cloud of data, and they managed to learn pretty well. So I think these systems can look past some of the bad data and still learn how to classify as we need it. So, yes, good data are important but at times, the real world doesn't have a lot of good data and so it's nice to have a system that can survive in that environment, too.

MR. PATEL: Just one comment. I was --

DR. BLUMENKRANZ: Identify, Bakul, just for

the transcriptionist.

MR. PATEL: Bakul -- going with Michael Three with your (inaudible) train here. I can also see we're all, I think which is a unique opportunity for us in the digital health is to tailor solutions for other population types, right. So, you know, in the hardware world, there's one thing about, you know, you need to make a product that sort of spans across the populations you're intending to use. But in the digital health world, I think with software especially and when you're looking at screens, easily malleable to make it to the population data -- population set that we really intend to so it can evolve, it can be personalized.

So when you think about personalization, and I think we need to think about those population types that can provide the same experience that we talked about earlier, right, because without that, you know, the efficacy or effectiveness of the products will be diminished or sort you'll be leaving stuff on the table that we shouldn't so.

MR. BRITTON: I'm going to move to the second part of this question which is, are there digital pharma innovations that could be applied in these circumstances such as tailoring of return visits or modifying treatments based on this information that's gained in the home? And I'm going to direct that to Quinton because you have some personal experience with this.

MR. OSWALD: Thank you, Mark. Can we go to slide 19, please? So Notal Vision introduced a device that monitors the switch from dry to wet AMD and was faced with a number of issues by virtue of the fact that dry AMD can be anything and the switch to wet can be anything from a 3 to 10-year journey which required frequent patient monitoring. And we realized that it was important to create an ecosystem that interfaced the patient with a doctor. And this is a real challenge going forward and we certainly have learned a lot from this process.

So on the left, you see -- on the top left is the ForeseeHome device which is a little difficult

to see, and I'll talk about the next device as I go through my presentation. But basically, what happens is a patient tests on a daily basis or every second day, and we're finding very good compliance because patients are really scared about losing their eyesight so we don't have compliance as an issue, although we have a compliance loop built into the process. This data is then fed to a cloud-based platform which sits in our independent diagnostic testing facility in Manassas, Virginia, tied to the cloud. And basically, it's reviewed by ophthalmologists and ophthalmic techs.

If the patient is not compliant, there's a feedback loop, as I indicated.

And on a monthly basis, we supply reports to the physician which we're learning need to be a lot more decision-based. Yeah, we used to supply a lot of data and we need to turn this a lot more into information, and we went through the process that we've started to do that more efficiently.

Should the patient or when the patient

switches from dry to wet AMD, an alert then is sent by our ophthalmologist, both digitally as well as a telephone call, to the treating physician to say the patient has switched from dry to wet AMD. That's really important because we find that in the study we did to support the reimbursement of this product, we were getting patients on average of 20/40 or better. And as Michael will tell you out of the iris registry, 650,000 patients that switched from dry to wet AMD, the mean presentation of patients at first treatment was 20/80 or worse. Think about it. So you've got a functionally blind person arriving for treatment on an expensive AMD drug so obviously, treating earlier is going to be far better. So that's the mission that we have embarked upon.

The second is we're developing a home-based OCT because we think this is critical for the next phase of the treatment of wet AMD, because if a patient comes in even on a monthly basis, you have no idea what happened to the patient between day 1

and day 30. As we know, that treatment interval is extending to two and three months and we really don't know what is happening to the patient.

So basically, on top of the platform we've built with ForeseeHome, which is we've just completed our 3 millionth test, so it's a pretty validated platform, we're introducing -- we're in early clinical trial on a home-based OCT. Now the three components about a home-based OCT; first of all, what's it going to cost; how reliable is the machine; and how do we present the data?

So basically, we're busy developing it but probably the most important element goes to the next slide. Now we basically developed an algorithm that basically automates the outputs from an OCT, and we conducted a study with 142 eyes, and the top left-hand side is we identified fluid and lesion activity of the 128 scans -- RP scans from the machine. The algorithm then categorized, as you get into point two, and unfortunately, the slide is compacted a little bit but basically, it priorizes (ph) from 1 to 10 the

likelihood of fluid being present. So really, it becomes a decision support platform for the ophthalmologist or the retinal specialist. So we have this device that categorizes where the patient is.

We then validate that test against two groups, compared (Inaudible) to a reading center and we compared (Inaudible) to three individual retinal specialists. And you can see on specificity, accuracy, we pretty much were comparable to that reading center and the three independent retinal specialists. Why is that important? Now obviously, Krishna talked this morning about machine learning and basically, we've run this device through -- this algorithm through about 100,000 scans. We're looking to push that a million so that we continue to learn and improve the accuracy of the algorithm.

So in summary, coming back to the question that Mark asked me, I think it's important to create an ecosystem with not only just the device.

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

lines?

DR. MYUNG: Yeah. Really glad you brought that up because I feel like human factors is a huge area, important area that I think maybe in the next digital health workshop, will be of a major topic. But yeah, I think this is important issue. First of all, for this camera or any system, you kind of describe an indication of use. Is it to be used at home; is it to be prescribed by a physician; is it to be used only in the clinic? So that's, first of all, prescribed, I think, for this device, but there are other devices under development that are intended to be used at home, and that's where the human factors comes into play.

I think the FDA -- and maybe the FDA, maybe Ron May want to speak about this is there's a whole human factors testing that's sort of like a subset of a clinical trial where people can take the device home or actually patients are to take a device home and describe their experience. And all that is supposed to be recorded because then

there's certain feedback, you know, and like a device -- for instance, if you're trying to measure your own IOP, I think there was a device approved recently that does that. I'm sure there was a lot there that -- about hazards to the patient's eye, you know, causing harm to your own eye, all those ergonomics and things like that that are important. So Ron actually just stood up.

DR. SCHUCHARD: So real quickly. There is a guidance document, a human factors guidance document and I would point you towards that but in terms of human factors testing, it falls back to what you've heard Bakul and I say. It's all based on the risk. So indeed if there is risk, you got a device at home and there is a risk for safety or risk associated, it all plays into how much of a risk. And that's part of the human factors testing, is to assess the usability risk.

DR. MYUNG: Right. So go ahead.

DR. SPRUNGER: Thank you. Our next panelist will be Eitan Sharon who is a founder and CEO of Mode AI, which has developed artificial

intelligence-powered visual chatbot. He's co-founder and CTO VideoSurf which will acquired by Xbox, so that means probably all our kids know your name very well.

(Laughter.)

DR. SHARON: At least I'm in the box. So I'm going to to take you to a little bit of a different perspective. If we moved from software to hardware, I'm going to go back all the way into algorithms, so all the way back to AI. My background is academic in math, computer science, computation, vision and learning, been teaching in Brown University and other places and then moving into the world of startups, entrepreneurialship, and building machines and code that can see in real time and analyze visual things and provide things like, you know, intelligence on the Xbox or even more AI power, what we call -- and I'll break that down for you to see the relevance -- AI visual bot for conversational shopping.

So it sounds colorful but let's break it down. So AI, we've talked about it so many times here.

That's what we actually do, deep learning AI, visual. We do that's for images for visual things which is very much relevant for medical imaging as well as other images.

Bot is the aspect of actually talking to some machine, which many of us could be confused as to whether they're talking to a person or not. Our conversation is just the, again, back and forth so we didn't talk about that a lot, but AI -- conversational AI is becoming a thing quickly because it's back and forth with the AI, not only on one off provided image. So you've been asked -- you are being asked questions and according to the responses, there's a conversation.

And shopping, well, it's not not a medical but many of the aspects that we care about in terms of safeguards apply to the financial information, other shopping patterns that you do. For instance, the GDPR European standard -- it was just mentioned -- we're heavily into that and reviewing that with legal all the time is forcing

some compliance with the information that is very much relevant to the privacy that you would want in your medical record, because that would be what you shop and how you pay and what you're interested in, so very much relevant for that.

I'll about three of the questions that were raised for our focus. I'll do it pretty quickly and then we can move to the discussion. One is the safeguarding software, so machine learning can be helpful very much in that; for instance, things that we've been doing for many years, unit testing and holistic testing of the software. We can make that dynamic. Yeah, that would be -- they actually --thank you very much -- we're already there. Thanks.

So machine learning could be applied to testing units and (inaudible) in whole, and it's a dynamic thing. And again, talking about risk, it basically looks forward and it behaves in a way that we have expected it to behave but not in a rigid but rather in a dynamic learning way. So as the unit progresses, the checks progresses as

well.

Abnormalities is just another manifestation of that. So just things that deep learning, I've learned that our standard behavior, once you exceed, they raise a situation which we need to intervene.

And the human factor that was discussed before -- so I don't know if you're aware but currently in companies, the office sales of security are for social are running many tricks, like, you know, they put in front of you an email with a screen that you should be familiar with but it's not the real screen and you would enter your password and your information and you'll just give access to anyone. Or I can go right outside here and set up a wifi that says digital health, you know, conference and you go on your phone you would see that wifi, no one would check anything, would just go on it and put your passwords, and everything on that network is exposed.

So the human factor, in terms of adversary and how to monitor all those things is very much an

relevant in order to, you know, protect the information that we care about.

The other point of patient behavior, again, an interesting point of view. Any one of these machines may not be able to know exactly where it is right now, but it does an accelerometer. So it can to the millimeters know where it's moving, which means that actually, it can start from a place and know exactly where it is. So if it's in your pocket in your home and you allow, as a patient and you're interested in, we could know actually quite accurately the pattern of movement of your day, the locations, like the kitchen, you know, or the bathroom or your living room or other places, a walk, and monitor movement, location, and functional behavior, like where you would spend your time; in this regard, also be alert for problems. I know of people who actually watched the Nest machine on their parent's house, and it has some sense of motion and, you know, when they don't see a motion for a day where there should be one, they'll get concerned and they call their

parents. That's something that is of service if we are able to monitor some behavior.

Things -- and the last point to talk to is just safeguards in storage of information, so it was mentioned here before as well. We have the cloud. We have the two-factor. We have the instrument that we're actually holding, so once I approach the cloud, I can get pinged for validation that it's me through the hardware, and then I can choose which information stays on the cloud and which stays on that phone. Actually. there's a debate where is it better off, to be on the cloud more exposed or actually on this machine in which case we need to retrieve data if I use the machine, but it may be, in some situations, actually safer, believe it or not, to be on the machine than on the cloud. But in any case, the two factor helps us in the end-to-end encryption that we see these days, you know, with things as simple as "what's up" and other conversational all things that I've mentioned before is also very critical. We talked about that before but there

are protocols in which me and the other end are the only ones to be -- you know, have knowledge of the conversation and no interference in the middle can know what was sent.

This was a brief brush-up on the three points and then we'll, I guess, move to the discussion. Thank you.

DR. HUMAYUN: Great. Thank you. Could you talk a little bit about, you know, what your experience with AI is and managing some of these areas, you know, your personal experiences. Has it been something that's been, you know, easy; has been -- taken, you know, a million reiterations to train? I mean where is it in terms of your experience with AI?

DR. SHARON: That's a good point. I'll use to talk a little bit to AI. It was mentioned before that the main -- you know, main things we are aware of may be aware of are the data. Data is in abundance or there's a lot of training data these days, and machines have become very powerful; right, it was mentioned before, so graphical

designers and GPUs. But also, there was a kind of a mini revolution in the space of deep learning and machine learning in the last couple of years, that deep entered deep; basically, mean (inaudible) hierarchy entered into the system and currently, the results have surprisingly good. So if before we had to worry about these types the things that I've been seeing, your features or reasoning or things that we pay attention to, the beauty in this AI is that oftentimes the holistic view of the samples does not easily expose what those features are, which could be considered as a black box but is also something good because oftentimes they are not breakable to a few simple things.

Having said that, there is a walk around localization, identification of the features that matter to a system that has exerted successful behavior. So generative -- adversary networks, for instance, a network that works against unit can provide samples to -- again, was mentioned before -- to train against and then two networks

encrypting and decrypting the same thing in order to improve. It's not -- it does not require many iterations, actually surprisingly efficient. If you have grunt trust data to a sufficient degree and you use -- and, you know, you use a good system, results are surprisingly good and they are not relying on features, and I think there's a breakthrough.

Are we close to machine awareness? Not anywhere more than flying with our hands, but it's also impressively better and and allows us to do many more things and training is straightforward if you have someone who is sharply focused, you know, about what they're doing and feeding the machine.

DR. KARANDIKAR: So I did want to add to that briefly or give a different perspective on AI if I can.

MALE SPEAKER: Please.

DR. KARANDIKAR: So, I mean the previous panel there was a lot of discussion, the Michaels were arguing, right, which is -- you know, there's this

concept that, you know, AI is certainly magical and, you know, were are like solving everything. And there's no question that machine learning specifically, especially supervised learning but also a deep learning, has produced some amazing results. There's no question about that.

But at the same time, there's two major issues, right, that are solvable that I need to -- think we need to look at. Dr. Yeshwant mentioned this morning with the advances in computer science now and hardware technology especially, you know, you're going to have these neural networks that have hundreds of thousands of layers. And if you do that -- and it is going to be a black box for the foreseeable future -- understanding, you know, in these hundred thousand layers how the AI algorithm arrived at the conclusion it did is going to be essentially unknowable for the foreseeable future, because there's just too many layers to be able to go back and look at it arrived where it did. That's one thing.

Second, when the algorithm looks at a certain

set of data -- and I think Eitan kind of alluded to this -- it doesn't have any of the context, so it might look at an image, and it does a really good job, as the Kaggle competition showed us, about specifically doing deep learning and it -- deep learning especially good for image analytics. But, you know, a typical physician would say, oh, but I happen to know that this patient is older and, you know, of a certain ethnicity and maybe, you know, whatever else. And so they use that information to qualify what they get from the images. And so that's a second thing.

And Dr. Woodward, I think this morning, mentioned Bayesian learning, right? I mean the thing is, right, you want to take priors into account as well and really, if you just have a deep learning algorithm focusing on a very specific, although very, you know, wide data set, it still can go wrong because there's no sense of what's happened before. And so, you know, if you -- if we can think about how to combine multiple AI approaches to get to a better, you

know, sort of intelligence, right.

So imagine that you have, you know, the expert systems of the past have been discredited but use some aspect of that, use certainly machine learning and deep learning or both, you know, two sides of the coin we need to look at, but also look at probabilities where instead of the famous Google algorithm that says is this a cat or dog, what if instead you could say it's a "yes" or "no", it could say we think there's a 70 percent probability it's a dog, 20 percent it's a cat, and 5 percent maybe it's a mouse, right? And that's a -- that may be a more accurate representation of how a physician thinks about an image. I -- and I'm not a physician, right, but I don't -- I imagine that, you know, people think "I think it's this but let's get some more data and find out." And you could imagine that AI algorithms in the future could, at some level, formulate the premise, design some experiments, get additional information to validate the results you get from the previous AI.

So sorry -- I mean, you know, that's a long spiel but, you know, I think we should consider multiple approaches coming test procedure a larger AI conclusion. So that's sort of a thought process.

DR. MYUNG: I have a question for Eitan and -- or anyone else who wants to answer this, but I think chatbots are amazing. I think that's sort of the -- we talked a lot about AI in terms of image analysis. With chatbots, it's really you can have a conversation with this -- with the machine basically. And, you know, for instance, I voted last -- I registered to vote last year through a chatbot. I thought, oh, this is cool. So I texted this number and then they asked me my name; then it asked me my -- where I was born; and then it asked me all these really personal questions. But then at the end of this, "congratulations, you're registered to vote" and I just see there's so much power that could be used, because I mean compared -- I think most people would say -- maybe not now but soon -- that they'd

rather text, almost prefer to text than talk to someone on the phone maybe in some cases. It's faster in some; you know, maybe the younger generation. It seems that way.

And so I don't know if you want to mention like what's the future of chatbot in medicine as far as communicating, maybe talking to a device, a chatbot device nurse, that type of thing. I personally think it's really powerful but what's the horizon there?

DR. SHARON: Yeah, I fully agree. I -- you know, if you talk to some of the leadership in the big companies, some is real, some is not. They'll you basically it's already happened. If you just look at the East in which are the nine billion billion of people are just doing that, and they voted with their fingers, not their legs this case, that are shopping and other sensitive information, will be conversation or that will talk to the machine, and it's a very powerful thing. And if you experience any of those like you did, the first time you get some kind of

flowing experience, you just see it's natural. There's no buttons to know. It's not an app. It's not site.

DR. MYUNG: No forms to fill out, right?

DR. SHARON: And look, look, thing forward like next year, you know, all the big companies, you put some glasses on our heads. They're already there, all of the companies, and they'll give us some augmented reality. And now, you know, these things are connected inherently to the chat platforms, right; Facebooks will be connected to the Messenger and Google to the Arlow (ph) and whatnot, and Microsoft will be connected to the Skype with (inaudible). And can you imagine a website or one app on your glasses to talk to? I mean this would be inherently a conversational, social the bot thing, so it's just unavoidable and people love it, and it's natural and I think we're going in many of the cases, and it's better to be ahead of the curve with these things.

DR. MYUNG: Yeah.

DR. SHARON: It will be taking us further

(inaudible) AI just to -- we'd like to meet in, you know, one world as well. I think you're right to think -- but I think AI is acknowledging and getting there. For instance, in a bot, as I said, it's not a one off so it's not "what's this image." It's like are we talking now; what did you say before to talk to a question of a buyer; are we in a conversation; what was said, typed, which is NLP versus visual AI; and which image will click to interact with or the machine has seen? We're doing all that. And the way to go about it is math and vision and whatnot. You just take the vectors from the deep learning and you have a bunch of PhDs and you build upon the probabilities of the AI interweaved with NLP and other priors. So that's something we know how to do. You're and not limited and I think so many resources are going there that there is a lot of activity, and there will be innovation. It's just the beginning of something so certainly getting there within the framework.

DR. HUMAYUN: John --

DR. KARANDIKAR: Absolutely.

DR. HUMAYUN: -- I have a question --

DR. KARANDIKAR: I'm excited about the future.

DR. HUMAYUN: I had a question for John and, Nitin, since you're out there, you know, for a device you verify and validate and you put it through stresses so you get device failures. In this case -- I mean I heard Dimitri talk about this. I think maybe -- you know, I don't know who's doing it outside Google but people are actively trying to hack in. So, you know, small companies can do this. I mean how do you tell somebody that your system is safe when you haven't had somebody for -- put it through that rigorous testing or some sort of testing where somebody is trying to hack into it? Have you guys thought thought about it? I mean, you know, is it -- is that how you test a system and its proprietary and safeguard nature of it, by hiring, you know, 10 people to hack into it? I mean I -- you know, could you comment on that?

DR. KARANDIKAR: Absolutely. And so there's a

couple of different things that we are combining here; right? One is -- I think Dr. Humayun actually alluded to testing and making sure the app works the way it's supposed to. And separately, there's a security aspect of it. And thirdly, I think, you know, I want to talk a little bit about penetration testing which is kind of what Dr. Humayun alluded to as well.

So first of all, for verifying the quality of the system in terms of one of the challenges -- and it was really refreshing for the FDA to sort of talk about agile software development, and it's impressive, right. I mean this is where software is and it's amazing to see the device industry and the health technology industry getting you know at the same level. So for software quality, you can actually have a -- you know, we have a fairly comprehensive QA process that looks -- and this is fairly well understood in software; you know, there's aggression testing, unit testing; you know what should I say, action testing, all of those things, There's a whole bunch of QA activities

that happen for every release. So that's one part and I can go into that in more depth.

Second, I think you kind of eluded to security and so security is (inaudible) each release to do some quality testing or security, you can actually design your security architecture, and you can design the software so that you have a higher level of security regardless of those -- you know, those individual releases don't necessarily have a huge impact on the overall architecture and security framework. And so, you know, you can you can kind of manage that more longer term. As an example, you know, you think of the idea architecture and you have firewalls, you have a DMZ interface where all of that stuff is set up, and that does not change release to release. So you can actually test it very, very thoroughly, although for each release, you still need more testing.

But, third, I think, you know, in terms of hacking into the system, there are these software consulting firms that essentially offer services

called "penetration testing pen test," and so you can actually get these folks to come in as a tiger team and try to hack into your software. And I worked with them at -- you know, these folks are really, really good. No matter how well we kind of protect our app, they find a way to get in. And -- but that actually gives us really valuable information to figure out how to address some of the security issues. One thing I learned is that security is never done, right, so no matter how much you protect it, there are always ways to get in, but the challenge is to make it progressively harder so that at the commercial level, it's too much work for the attacker to come in. So, you know, in other words, you know, they -- you rather have them, you know, burgle a different house because yours is too hard to break into. That's not a really good analogy but that kind of is how it works in security so that's -- I don't know if that answered your question but that's the different parts.

DR. SPRUNGER: So this being a very up-to-date

technologic meeting, we have someone who sent a question from the webcast as a webcast attendee.

DR. HUMAYUN: We have a number of those but -- so it is being well-attended.

DR. SPRUNGER: So this question comes from Ron Cummings Kralik, who is a principal network engineer, surgical equipment at Bausch & Lomb. His question is, "What are the FDA's and doctors' thoughts on storing non-patient, in other words, machine data, on public cloud space?" Who would like to address that?

MALE SPEAKER: So that's maybe --

DR. KARANDIKAR: So (inaudible) the question is what are the FDA's thoughts on that so --

(Laughter.)

DR. BIN-NUN: So this is a GDPL; the European standard for information, requires that the information of the users will stay local to their activity or their country, but this local, to their regulations, could be Amazon's, be public, could be any company's, could be Microsoft so as

long as they keep it local, for instance, the regulations require that but do not require it will belong to the company with (inaudible). It will require -- it requires that it be protected, deletable, sent back to you upon request and many other things but not that it will be off of a public cloud, just under some regulations. So that's a cue from there (inaudible).

DR. SPRUNGER: I think these questions are probably exclusive to this panel so if anybody has any answers. Second one from the same person is, "Would the doctors find value in being able to merge anonymized treatment results from their EMR back into the public cloud space to allow analytics to rate treatment effectiveness? Any takers on that.

UNIDENTIFIED MALE: I mean I do have an opinion here but -- so I do have an opinion which is if it's -- if the question is for de-identified data, then by definition, that is de-identified, there's no HIPAA concerns about putting it in the public cloud. Sorry (inaudible.

DR. SPRUNGER: I think Michael has a comment. Please.

DR. CHIANG: Derek, just in response to that question, I think in a lot of ways, that's the premise of -- this is Michael Chiang from Portland, Oregon -- I think that's what the premise of iris registry is really meant to do, from that big data paradigm, the data gets out there and everybody sort of benefits from the knowledge discovery that can occur from that`. And I think the challenges are, you know, getting the doctors to buy into that, getting doctors to send the data, and then finding people to do that analytics and, you know, perform that knowledge discovery.

DR. EYDELMAN: And we at the FDA are very interested in utilizing all of the data in the registries for -- as post-market data, as data that we can, hopefully, utilize to expedite getting new treatments and new devices to market. And we are exploring collaborations Iris in a number of venues.

DR. SPRUNGER: So while you're there, please don't move. That led to question number three, Would the FDA allow this sort of sharing, which I think you just said yes.

DR. EYDELMAN: We actually would like to propose that more of that occurs as that really is a way to move forward and that is one of our strategic priorities, as a matter of fact.

DR. AL-ASWAD: Can I --

DR. SPRUNGER: If you don't mind staying there for number four. Oh, we --

DR. EYDELMAN: Okay.

DR. AL-ASWAD: -- can I make a comment? I don't know if you know, there is a big study. It's called "all of us." It's five big centers. One of them is Columbia University in New York. And basically it's a genetic study that patients register in it and they get their blood tests and urine tests in addition to their health information. And it's the start of the part of President Obama's initiative to collect information to use it for precision medicine in

the future. And this is currently happening, so it's collected somewhere. The -- well, not identified fully because you have the patients' health information with it, and you can actually do a lot of studies as a participant. You can have an idea for a research project and you could utilize that data, and it's called "all of us" and it's different -- five centers

DR. HUMAYUN: Okay. Well, thank you. The panel has been very interactive. Hopefully, you got something out of it and we're moving on to the next part of that program. Thank you.

(Applause.)

DR. REPKA: So we just have a very short remaining portion of the program, and there will be the three summaries from each of the panel chairs or co-chairs and then some concluding remarks.

Those of you that weren't here at the beginning, Dr. Shuren mentioned that the Commissioner would be unable to attend this afternoon, so that's an unfortunate but realistic

outcome.

So the first panel, I think when Natalie's ready, she'll start.

DR. AFSHARI: Well, great. We had a great day with lots of discussions and exchanges.

So we had the panel one, Mike Trese and I, and the panel was for Dr. Dimitri Azar, Leslie Bottorff, David Morrison, Darius Moshfeghi, Mia Woodward, Ingrid Zimmer-Galler. And the main question and discussion was safety and effectiveness concerns when a digital health device provides information as an aid for diagnosis and the assets, threats, and vulnerabilities to be considered and to be identified.

So we talked and touch base on several items but one was the tempo of the disease may reflect on how important the device contribution is to the diagnosis, and a physician and user may be able to override any machine inconsistencies. It also came out that some of the safety factors depend on the end user, that there should be some assurance

that the doctors are up-to-date in the management of that disease, and it came up that some of the diabetic retinopathy patients may read the images even better than some doctors because they're focusing so much this day and age on their images.

Also, regarding privacy, we have come a long way in privacy but have some ways to go in digital health. It came that there is evidence that local storage and cloud-based storage have similar security profiles and industries are being created to assess the level of cyber security.

And that was the summary of our panel. Thanks again, everyone, for a great day. I'll say goodbye from here, so we won't come back and take a minute to save a minute at the end. Thanks again to all of the organizing organizations and as well as FDA.

(Applause.)

DR. NISCHAL: So this is a synopsis for the second panel. What I found really interesting was that there was a real variety within the panelists. You know, we had known physicians who

were scientists and epidemiologists, physicians who were retina guys, and pediatric ophthalmologists, and also somebody from the FDA, Bakul, in particular, of course. And it was a very global represent -- if you listen to the accents, you know, we went from South African to the UK, Europe and the U.S. So I think that gives you an idea of the backgrounds coming into the panel.

You know, what I took away from Mark Blumenkranz's introductory slides was that digital healthcare actually is a response to the increased connectivity that the Smartphone has brought to us and our society. It's a societal response that we're gaining from entrepreneurs and the relevance of the clinical utility of the digital digital healthcare applications requires large data analyses, which Quinton Oswald was was talking about.

It appears that whatever the environment that the digital application is being used, in we need to identify a workflow that designates who reads

the data, who acts on the data, and who's responsible for that data and the actions that you have to take on that data. I think we don't want to be in a situation where the data ends up in the physician's office, we don't act on it, and a patient comes to harm, because I think we then are in an enormous amount of legal problems.

The development, therefore, of specific roles within the health care environment is probably essential and whether those health care roles have an IT background or a software background, you know, I leave the audience to ponder and to think about.

I think that -- I got the overall sense that the safety of data storage didn't seem to be so much of a problem, that HIPAA compliance with these storage was there; if somebody wants to hack into it, they're going to hack into it no matter how hard you try but that there is this issue of the human engagement or the human abuse of the data that might be more of something that we have to look into.

As big data and volume increases the data that we get, it seems to me that artificial intelligence is almost inevitable. And I was really interested that there was no question that AI had to be involved amongst all the panelists. In order to get it to be effective, however, I think that needs to be a culture change. That's what I got the impression of where protocolization becomes part of the medical culture more than it perhaps is now and that the deviations of -- from the from the protocols are important because they may act as an echo or a feedback so that we can look at the protocols again and see how we can make them better so that there's reduced deviations.

I think as we move to digital health applications more and more, defining who owns that data is also going to be an extremely important point of reference, because is it the patient; is the institution; is it the doctor; or is it the company that's made the application? And I think that these areas perhaps today have given us food

for thought because they're gray areas; they're not black and white. And the arguments that were made amongst the panelists in the second group, in our panel, I think, just highlight that it's exciting but there are a lot of issues and questions that we still have to answer.

Thank you very much to the FDA and everybody else who was involved. Thank you.

(Applause.)

DR. SPRUNGER: So I think Ken and Natalie went over most of the major points, don't have a lot to add. Couple of things that Mark and I wanted to pass on. The importance of software as far as the password protection, employee training, and encryption, I think we need to emphasize all those things.

As far as hardware, we heard about the light standards, the ANSI and the electrical standards, the ISO, again, being very important.

And lastly, I think, for me, personally, I got out of this more than anything is the human factor again, and that was mentioned by Ken. That's very

important. I think a lot of us tend to overlook that and I think it's very important.

So with that, panel three, was happy to be here today and thank you to the FDA for allowing us to participate.

(Applause.)

DR. EYDELMAN: Well, apparently I get the last word. Thank you all so much for coming, spending the day, and sharing with us your thoughts, your knowledge, and your, most importantly, your enthusiasm for helping us expedite ophthalmic digital health.

All of the slides from all of the participants today will be available at FDA's website pretty soon as will be the complete transcript of today's proceedings.

The a goal is for us to write a manuscript summarizing highlights of today's meeting, and I believe that just sharing the knowledge that we accrued today with the general public will help us in our goal of expediting ophthalmic digital health. Thanks to all.

(Applause.)

(Whereupon, at 5:12 p.m., for above-entitled meeting was concluded.)

CERTIFICATE OF NOTARY PUBLIC

I, KeVON CONGO, the officer before whom the foregoing proceeding was taken, do hereby certify that the proceedings were recorded by me and thereafter reduced to typewriting under my direction; that said proceedings are a true and accurate record to the best of my knowledge, skills, and ability; that I am neither counsel for, related to, nor employed by any of the parties to the action in which this was taken; and, further, that I am not a relative or employee of any counsel or attorney employed by the parties hereto, nor financially or otherwise interested in the outcome of this action.

KeVON CONGO
Notary Public in and for the
State of Maryland

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