>> Mark Eichorn: Good morning, everyone. Could I get your attention, please? Thank you all for coming today. We'll be welcoming Chairman Leibowitz in a moment, but I just wanted to give some administrative logistics announcements for starters. First, if you go outside the building without an FTC badge, you'll have to go back through the magnetometer and the X-ray machine before you come back in. Secondly, if there's a fire, please go outside the building. [Light laughter] When you go to the Georgetown Law Library and Georgetown University, right outside the main entrance -- so, if you take a right and go across the street, there's a staging area where you ... is being recorded. Finally, if you're on Tw

and tweeting about the workshop, please use the hashtag #ftcpriv. Now it is my honor and plea

to introduce chairman Jon Leibowitz. Chairman Leibowitz joined the Commission in 2004 and

was named as Chairman in 2009 by President Obama. Throughout his service on the Commiss

Chairman Leibowitz has made consumer privacy a top priority. He's strongly supported the agency's efforts to make Do Not Call registrations permanent, to get spyware off consumers' computers, to encourage self-regulatory efforts, to give consumers notice and an opportunity to control collection and use of data about them, and to bring actions against companies like Twitt Google, and Facebook. He's also led our fundamental reexamination of how the agency approa privacy issues. He's pushed the agency not just to address problems that occurred the past, but think about how changes in the marketplace or in technology will affect consumers in the future So it's particularly apt that he would kick us off today as we examine the development of facial recognition technology. With that, Chairman Leibowitz. [Applause]

>> Jon Leibowitz: Thank you, Mark, for the kind introduction. I've always had an interest in privacy, but one of the things that I was fortunate to have happen to me when I first came to the Commission was Mark Eichorn became my attorney adviser. And he just inculcated sort of the values and the importance of privacy. You were sort of a mentor to

great gifts with one hand, and it stabs you in the back with the other." [Laughter] That's Mark. Do you want to sit, or do you want to stand here while I speak? What would you like to do?

>> Mark Eichorn: No, I have to do one thing before I return to my seat.

>> Jon Leibowitz: Oh, you're doing the -- Oh, go ahead. That's fine.

>> Mark Eichorn: When you start the clip, I'll run the clip for you.

>> Jon Leibowitz: Oh, you're going to do the clip? Okay. So, I'll start.

>> Mark Eichorn: Yeah. But I'm standing unobtrusively back here.

>> Jon Leibowitz: So, let me -- [Laughter] Mark has always had my back Td   -1(S)2(o)1(,P6I25 0 Td   0 TTyp

>> Man #4: Stressed out, John Anderton?

>> Woman #1: Get away, John Anderton. Forget your troubles. Escape from it all.

>> Jon Liebowitz: So, I think we can stipulate that our technology is a little clunkier than facial recognition technology. [Light laughter] So, that is the future DreamWorks imagined for 2054 in 2002, just 9 years ago. Facial recognition technology then was the province of science-fiction writers and futuristic movies. In fact, were we having this conference in 2002, we would probably be holding it in Los Angeles. The audience would be full of science-fiction buffs. Steven Spielberg would give the keynote, and we might actually have the budget to provide you with lunch. [Light laughter] Sorry. It's part of or government austerity program. But here we are in 2011, a full 43 years ahead of schedule, with reports of companies beginning to roll out smart signs and tailored messages based on passerbys general attributes, age, and gender, all gleaned through

>> Mark Eichorn: Thank you, Chairman Leibowitz. We'll jump right in to our first panel. We wanted to set the stage for today's discussions by talk

just deal with face recognition, that aren't used for any other object recognition tasks. And as a consequence, humans are very good at recognizing familiar faces. So, if you look at the images that are shown in the middle here, you probably all recognize them, even though they're incredibly low-resolution images. However, if you look at the image pair below and you ask yourself, "Is that an image of the same person?" it's probably a harder task because humans are not very good at recognizing unfamiliar faces. This has been shown experimentally even for people with extensive forensic background. Now let's look at how automatic face recognition developed. By now, we're looking back at more than 40 years of research that started in the 1960s with a system built by Bledsoe and colleagues, which wasn't an automatic system yet. In that system, a human operator had to enter facial-feature point locations for the system to work. The very first automatic system was described by Takeo Kanade in his PhD thesis in 1973. And he reported results on a grand total of 20 subjects. So, we've come a long way since then, as evidenced in thousands and thousands of academic papers, across a range of academic disciplines. When we are looking at face recognition, we are oftentimes looking at one of two application scenarios that I indicate here -- identification and verification. In the identification case, you're given an image of an unknown subject, and the task of the algorithm is to compare it to a set of images of known subjects, oftentimes referred to as "gallery." And the algorithm then outputs a rank-ordered list of matches, hopefully with the accurate match coming out at the top. In the verification case, a subject claims an identity -- in this case, Robin Williams -- and the input image is compared to historic representation of that person and the algorithm gives an "accept" or "reject" response to that claim, depending on if it finds the match to be close enough. Now let's go to the steps that a face recognition system typically goes through in order to process a face. There are four distinct steps -- face detection, normalization, feature extraction, and matching. The task of face detection is quite obvious -- find all the faces that are in the image. This is, again, a difficult problem because faces are complicated objects. What makes it easy, though, is the fact that faces look very similar to each other. So, we all have

distance. So, we're adding mathematical representations here. We can compute a distance. And there are different options of how you compute that distance. There are other ways, too. Hopefully you come up with the right result here. There are advanced models too that you can use, especially if you have multiple images per subject that usually help you in doing a better job in recognizing. So, now these are all the steps that face recognition systems typically go to -- face detection, normalization, feature extraction, and matching. Now let me conclude here by pointing out two recent trends in the field. One has to do with data. For most of its history, face recognition was lacking data because it's very cost- and time-intensive to recruit subjects to come into a lab and sit down and possibly come multiple times to be recorded. And so, as a consequence, there were only a few databases available of reasonable size. Now, this has all changed in the last 5 to 10 years, with the increased use of digital photography, as well as the users of online social networks to share images. One of the recent examples that we're probably going to hear about more today is from face.com, who say that they have indexed 31 billion -- yes, that's "billion" with a "b" -- 31 billion faces in their processing, which I'm sure helps them tremendously in building very accurate face models. The second trend has to do with pose. As this video here indicates, faces are three-dimensional objects. So they look very differently, depending if I look at the face from a profile view or from a frontal view, which makes it very hard to match faces across pose. On the other hand, in real-world images, you oftentimes find pose variations since people don't always just look directly at the camera. So, what's been interesting in the field is that over the last few years, a number of systems have been proposed and tested that are able to compute a full 3-D model from a single input image completely automatically. And what that helps you to do is to essentially re-render a face in any pose you want. I showed in the example here where you take a frontal-input image, you build a 3-D model, and you compute a profile view of it, Which helps then tremendously in recognizing faces across pose. And with that, I'd like to conclude. [Applause]

>> Mark Eichorn: Dr. Phillips?

>> Dr. Jonathon Phillips: Thank you, and thank you for inviting me here and giving the opportunity to speak about the basics of facial recognition. So, I'm going to start off, and a lot of the first couple slides echo the beginning of Ralph's talk, but I think it also emphasizes the importance of some of the points that he made. I should say, when I go through here, I do ask

questions to try to keep the audience awake and keep them participating. So, you'll see the next slide. Hopefully the format comes through. And please tell me who this person is. You have half a second. Oops. There goes my slide. The next slide -- [Laughter] Who is this? I'm sure nobody in the audience has trouble recognizing this person. Of course, this is President Obama, and this is a famous person or familiar person. We are very good at recognizing people we know well. But a lot of the work we have been doing or I've been participating in over close to 20 years has been unfamiliar. So, I'm going

I'd like to, I guess, move to question time, but I think I would just like to summarize the two key points, is that there are two fundamental classes of face recognition. One is recognizing familiar faces. We're all very, very good at that, and we see new faces or make decisions. Sometimes we interpolate our decisions or extrapolate our decisions based on our ability to recognize familiar faces versus unfamiliar, which is a different problem, which a lot of algorithm-development community effort has worked on. The other is, we have mug shots or controlled frontal illuminations. There's been fantastic progress in the last 17 years, as shown by missed evaluations. But as we now move and solve that problem, we're now moving upon to the other cases. What happens if we do point-and-shoot? What happens with video and a lot more constrained? In many ways, this is the cutting-edge of research we're just moving into, and it's probably -- probably what I've seen more models of the southeast social-media concerns being addressed. And I think this is now an active area of research, and time will tell over the next five years if we see similar limits in performance or maybe the different capabilities. Thank you. [Applause]

>> Mark Eichorn: Thank you both very much. Let me, I guess, start off the questions with one sort of taking off on the "Minority Report" theme -- and you both addressed this in different ways -- but with the increase in images that you talked about and the, you know, vast increase, the acceleration of the accuracy rate in the last years of study -- Is that "Minority Report" kind of situation feasible now or feasible in, you know, 10 years, with sort of predictable advances in, you know, the number of images, the increase in computer speeds, the reduced costs of computer storage and developments such as that, as well as the developments in the facial technology itself?

>> Dr. Jonathon Phillips: So, you -- Oh, this is -- Oh. Oh. So, one of the things about "Minority Report," it just knew who that person was moving forward, which is a huge database. I think, you know, these being frontal -- I think for the moving forward, I think that'll be always a challenging problem. I think the limits of where you can push face recognition in unconstrained situations is an open question. In other words, there's some fundamental limits -- that we don't have "X" quality of images, where I mean quality, not -- or maybe quantity over enough viewing conditions. I think it's an open research question of what the limits of recognition are in the unconstrained environment.

>> Mark Eichorn: Ralph, do you have any comment?

older ones I've shown taken with a digital single-lens reflex are maybe six megapixels. And the conclusion I've come to after looking at these is, it's not the number of megapixels you have. It's the quality of the camera and the lens, the optics, and the ability to F-stop down itself. So, the issue is not how many more megapixels you can stick on your smartphone. It's, can you stick a better lens and focusing mechanisms on the smartphone that dictate the performance over a wider range of the quality of images for recognition.

>> Mark Eichorn: We have a question from the audience, and the question is to Dr. Phillips. Has NIST issued any standards on performance quality or integrity of facial recognition technologies?

>> Dr. Jonathon Phillips: So, one of the research projects we're starting out with is to measure quality -- to start measuring quality of images for recognition. This is not a solved problem. We have, for example -- I should say we're in the process of formulating a challenge problem to be able to measure and assess what are effective quality measures for face recognition for matching.

>> Mark Eichorn: So, another question is, as we get more and more images of each of us and images that are associated with us, does it then become easier to identify an unknown person when you're comparing to sort of a known data set? If you have 100 pictures of me that are known to be pictures of me, is it easier to identify that 101st picture of me?

>> Jonathon Phillip: So, the ones that are cooperative, you actually go up and interact, like a foot away from the scanning to get a high quality. I would say acquiring irises at a distance is currently a research topic.

>> Mark Eichorn: And you said that retinal scanning was a different type of process?

>> Dr. Jonathon Phillips: Yeah. Basically, I think, if I remember correctly, retinal looks at the blood vessels in the back of the eye, and it's much more intrusive than retinal scanning. I don't know of any serious applications with it.

>> Mark Eichorn: Okay. Could you talk about the differences between facial detection and facial recognition? You discussed it, Dr. Gross.

>> Dr. Ralph Gross: In terms of performance?

>> Mark Eichorn: Yes.

>> Dr. Ralph Gross: Yeah, generally, face detection performs better. You see it quite frequently in use in your consumer-grade cameras, that you have a little green box hovering over faces. And across a large set of conditions,b EMC   nsumice642rA1.725 Td   [( )Tj   0.ovetssh Gri e7ge Td   (-).725 xID 3 >
ieli?41(Td   (-18.7co)u(m)8(orbu(y)5(g)-2)4(?)t)-2(e)-1(c)4(t)-2(i)-(n yol)-2(y)5(p)-2?ou se-23(t)- 1(. A)2(nc

>> Mark Eichorn: And what kind of criteria -- We've heard of age and gender being used or detected, some type of categorization based on that being made. Are you aware of any other types of sort of identifications that are made by facial detection?

>> Dr. Ralph Gross: I'd say these are the most popular one, gender and age. You know, ethnicity is certainly one that I've seen, although not as frequently. So, yeah.

>> Mark Eichorn: Okay. Well, with that, we'll wrap up the first panel, but I really appreciate both