

# Future of Health Technology Symposium

Presentation by:

Rosalind Picard, ScD. FIEEE

Author, *Affective Computing*, Director, Affective Computing Research,  
Co-director, Things That Think Consortium

*This podcast is presented by The Centers for Disease Control and Prevention.  
CDC – safer, healthier, people.*

## **Rosalind Picard**

It's a real honor for me to be here. I grew up in Atlanta and I've always had this sense of awe about the CDC so to finally be here is very exciting for me. Thanks for inviting me. I'm going to give an overview of a bunch of things kind of fast and I just want to let you know it may leave you with more questions and I will give a URL at the end and be happy to hang around through break and all to try to answer or point you to more information on those. To start I'm going to cover very quickly three areas and the first will have to do with the basic idea of sensing affect and particularly some states that we think correspond to risk situations in decision making. Second, the importance of and one illustration of delivering an empathetic interaction via technology, and then third some of our new work in trying to help people with autism and perhaps also those who are vision impaired. So very briefly I think this is common sense but you know in psychology you have to do 20 years of study on it before people believe it. But basically we know that certain emotions strongly influence the behavioral choices we make as well as judgment, cognition, perception, a whole host of cognitive phenomena these days. And one of the biggest concerns I share with you is how to prevent a lot of these harmful behaviors and at the same time foster the ones that would be more beneficial to us.

In particular we wish we could detect, and we're working toward this goal, those emotions that are triggers, whether it's a craving state, whether it's well I was doing good with my diet and I got all stressed out and I just had to eat that whole thing of ice cream, or seeking drugs are much more dangerous situations, unhealthy sex and other behaviors. So how can we get technology to actually help people when they're say, all stressed out at a moment of risk. Now this is particularly challenging because if you're like me, most of your experience with technology is it just increases stress, right? This is not the thing you would usually seek for a solution, right, you'd seek solace in a good friend, not going to your computer which is probably in many people's experience going to aggravate them more.

For years we've been working on new ways our lab is, the pioneers and measuring emotion while people are actually doing things such as filling out a health related form and seeing not only how to improve the usability by detecting very subtle changes and frustration and stress, but also how to take that ability measure and use it to make the technology less frustrating. Let me give one illustration. This is a health feedback form that's used as input information by Jack Dennerlein, at the Harvard School of Public

Health, who does a lot of research on ergonomics and risk factors. We coupled filling out that form with use of a mouse that Carson Reynolds in my lab built with pressure sensors on the mouse and what you'll see here as I play this little movie is you'll see the trail of where the person moves the mouse and if they squeeze the mouse a little harder which might be associated with frustration and stress, you'll see the line get thicker and redder. I don't know how good the red will be here so it will also get thicker. So watch this. We're heightening stress because we're asking them a lot of personal information, asking about smoking cigarettes – oh, we're also deliberately using bad usability here to try to heighten stress a little bit. What is your current marital status? And here we're going to tell them that the children's ages should be separated by commas and you'll see, this is a little annoying here, ages should be separated by commas, please re-enter information. So we find that we can present the same form, the same mousing behaviors with and without subtle stress, and in this particular study we ranked people's experience of using this interface from most frustrated to least frustrated and in addition to seeing increased mouse pressure, significantly higher in the most frustrated group compared to the least frustrated group. There was also significant increase in the specific muscles that are associated with wrist injuries and so that was good and bad to understand that a little bit of irritation on the computer contributes to the stress that many people experience in their wrists, and so little timers that tell you to just take a break every half hour to help your wrists may not be as effective as a pressure mouse that starts squeaking when you've been squeezing it too hard after a while. Something more adaptive to you has been our aim.

We would like to go mobile with this and we have in a couple of cases now, and the study I'll describe is kind of small but it builds on a number of principles we validated in larger studies and Tim Bickmore who follows me is validating the much larger ones, too. So I think it has merit beyond the limits of small numbers. And here we didn't have as snazzy a device this time as Astro just described. We customized a heart rate monitor that was on the chest with accelerometer information from your foot plus context beacons that we could drop into the environment so you could privately choose to put these where you thought you wanted to have it mark your location. This is not somebody spying on you, this is me saying I need one in my kitchen, I need one in my office, I need one in my car, that's where I want to know where I am when these symptoms happen. And our challenge was to do a kind of experience sampling that wasn't aggravating people. Usually people drop out of experience sampling studies; the experts say you can't pay them enough to label how stressed they are 12 times a day because it just increases their stress so much that they want out of there. So our challenge was could we design a way to get this labeled information from people while getting their physiological information so we could start to build a model of what happens in their body that we should be measuring that predicts their stress or anxiety or whatever, depression. Here we focused on stress. And then do that without irritating them even more.

Our question in a lot of this work is to ask what is the equivalent human-human situation and this is inspired by findings by Reeves out at Stanford who say that it's often helpful when designing technology to think about what would people do if you take out the

computer and put in the person. So we thought, could a person enter up to 12 or 15 times a day without increasing stress and we thought, well there's some people who could and there's some who couldn't; now what's the difference? And one of the key differences that we think really matters is the ability to use empathy. And I'm jumping ahead of myself. There's lots of sensors, Astro showed one, I'm wearing another one I can tell you more about later. But how do you enter up people frequently without being annoying and one of the key things is not just a sense what's going on with them but to sense and respond to what's going on with them with something like what this dog does. So here, I like this cartoon. He's the first one who convinced me that a machine could deliver empathy and have it just not drive people crazy. So here we have a master coming home at the end of the day and he looks like he's had a pretty bad day, and we know that by looking at him. We don't know how the dog knows it. We know that the dog is always happy to see you no matter what, and the dog's tail is wagging. The master sits down and he looks upset. Now we know the dog recognizes this even though we don't know how. We don't know if the dog looks at you, listens to you, hears you slam the door differently, your footsteps, you smell different, or what. But somehow the dog recognizes it because the dog puts his ears down, his tail down, and mirrors your state. A very simple but very effective mirror because instantly master feels a lot better. Dog makes master feel better and now dog's happy again, too. Wag, wag, wag. So all of this happens, of course, much faster than I just explained it and we know this effect that a pet or person can have on you if they help you understand that your feelings are being understood and then somehow you're able to get past those bad feelings. This is not just something for cartoons, it's been measured in medical situations and I've been told is the driving impetus behind why physicians now have to be tested on their bedside manner before they get fully certified because it has an impact directly on the bottom line in terms of malpractice lawsuits. Physicians who commit the same error, one is empathetic, one is not, which one is likely to get sued and cost the insurance company the most, the one who's not empathetic. So there's lots of other things that suggest that this is a really important factor and it has been missing from technology. How many of you have had a computer apologize to you or say I'm sorry I had to keep you waiting, or gosh it looks like I was really frustrating. So we have been implementing that in our technology and finding it makes a measurable difference in people's behavior and interest in using the frustrating technology.

So here's just a quick example of that to show you just how simple it is. This is our mobile system; people took around a PDA. They could either interrupt it to give it information or it could interrupt them. If it interrupted them, we thought it was extremely important that it be empathetic. So we compared, everybody used the system for four days, everybody used each of these two systems for four days. System one and system two, we didn't tell the subjects this was a study about the role of empathy until afterwards when we debriefed them. And for example, without empathy they have this user friendly system. I want to emphasize empathy is not the same as being friendly. So it's friendly, it's social, it tries to build a relationship with them, morning Jane, actually doesn't succeed in trying to build as much as Tim's stuff that I'll show you afterwards, but it tries to get at least the social stuff you would have at the beginning. You know the drill, feeling stressed?\*\*\*\* thanks so much for your input. The bottom dialogue in the

other system, system two, is identical except for inserting this red line that customizes a response based upon how they say they're feeling and these are all pre-scripted so you can deliver it instantly in technology. It does limit the choices of the user but at least it seems to still be very effective. What we find is things like people do, again a small group but in other studies that are larger we find similar affects and of many, many measures they were all in the same direction. People consistently seem to prefer this empathetic dialogue over the one that just ignores their feelings. There's some other things that change the system. Again, I'm flying by a lot of details here so I'll be happy to address those later if you'd like. Again, there are lots of measures that suggest this here.

Now our next challenge is to take something like this mobile system and to sense not just things like are you stressed now, but are you having a craving now. We're teaming up with some physicians at Children's Hospital Boston who are experts in addiction treatment and trying to see if we can deliver a system that can help addicts figure out what it is in their life that's triggering these episodes that cause them to use, whether it's a particular GPS location that they're in or some physiological changes in their body that we can measure. In this particular system, I can't show you details here but they're basically including heart rate, change measurements, skin conductance measurements and accelerometer measurements as well as location information all in a wristwatch form factor and possibly something on the ankle right now. So the idea is to give them this small sensor, have it hooked to their cell phone. The cell phone is a GPS source for location and also a lifeline source delivering interventions like you'll see from Tim or interventions from a real human being, and it's going to involve machine learning which our lab also does a lot of to learn which patterns might predict your cravings so that before you actually get a strong craving, it can say, hmmm looks like Joe is headed for trouble, gosh he's headed towards his dealer, I'm going to interrupt him now and empathize and try to be understanding. And then it could hopefully be there at the moment of high risk to either calm you down if that's what you need, or remind you of other personalized preventative interventions perhaps that you learned in cognitive behavioral therapy or other situations whether it's playing with your strong attachment to a child or another loved one or something else that motivates you; try to deliver that at the moment of weakness so that we can help people really overcome these things and prevent relapse in behavior.

I want to emphasize in all our work that this is not stuff that's being put on people and forced on them to spy on them. Our systems take a very different approach which is very user centered, giving people technology as Astro has .... that they can choose to wear or not. They can take off, they can put on. But trying to make the technology so cool, so calming, so wonderful that they actually choose to wear it and learn more about themselves and change their behavior as a consequence. So we're constantly looking for ways to make the technology more caring and more empowering of people's feeling that they are in control.

So I've quickly run through two things there that very briefly mentioned, we've building technologies to sense different affect of states in a mobile way on the body. We couple that with machine learning and pattern recognition to try to do things like predict cravings that you've labeled previously by this device that interrupts you. We are able to just change the wording in some cases so that it appears more empathetic and caring and that makes a significant measurable difference downstream in whether people choose to use the device longer or not. And finally I'd like to attempt a live demo here if I can of our new autism work. So a lot of things here. Let me first introduce my colleague on this work. This is a picture of Rana el Kaliouby. This software I'll be showing you began as her doctoral dissertation at the University of Cambridge and then she's post-doc with me at MIT where we've been building the version you'll see here. The basic problem, gosh you guys have been all in the news lately because of your recent report about increase in incidents of autism estimated now at about one in 150 through the survey of 14 states. We can't prevent autism yet, we don't know how to do that but we're trying to help people who are young adults and older adults who are fairly high functioning, they are able to get about and try to interact socially with people but they have serious social deficits. For example, they have a very hard time reading, some of them even looking at another person's face and really are often clueless about what's going on. For example, if you're really concentrating on something and this is a bad time to interrupt you, they wouldn't really notice that and they would interrupt you anyhow and they would be perceived as rude, they would often inadvertently be very irritating and frustrating much like our technology any many ways. Our technology is also autistic. We've been trying to develop things that improve the experience for people with autism jointly with the experience of people using autistic technology.

The current states that are in the system that can be read in real time from a person are these very complex cognitive affective states. They include things like agreeing and disagreeing and as sub categories all these things in white. These sub categories were given to actors in Britain that aid in autism research gathered to develop a mind reading DVD to help people with autism learn to read these mental states of other people and so this is a DVD. The current interventions, you sit down, you play games on this DVD, watch lots of examples and you try to improve your ability. The challenge though is that what a person with autism learns using the computer and the DVD may not generalize to what they do in real life. It's very hard to get it to go to real life. So we thought, well we should be able to build a mobile system so that I can wear it around and learn what these things look like for the people I interact with on a daily basis. So we've now taken this off of a DVD and put it into a real time mobile system so you can just be talking to somebody; you can't read their face but your camera can start to do that.

I won't go through technical details here; it's already running.... But just to let you know there's a whole lot of pattern recognition, signal processing, computer vision, machine learning; you may hear those terms, networks to recognize things which are relatively robust to low level features that may not be so robust. Let's see if I can do this live for you here. I have a little camera attached to my computer right now. The little green dots are a tracker made by Nevin Vision that we use at the front end of the system that hopefully we find my face there, yeah. This camera is really tiny. We often just stick it on

a hat or on eyeglasses or the ear or you can put it on your chest. Right now I've just got it stuck to the top of my laptop so you see the little green points are finding my face. I haven't just glued little green points on my face. People sometimes get confused about that. And this is kind of a complicated display; this is not one we're presenting to the people with autism except when they want to study the system like in the forum factor right here but it's pulling out lots of low level features like head nod, head shake, head tilt, head turn, head forward, head back, lip pull like when I smile, lip pucker, these kids really want to craft relationships in their teens, so brow raise, actually lip pucker is good for confusion and thinking. And then here we have six states, disagreeing, thinking, confused, concentrating, interested, agreeing, color coded and what you see here is the probability of each of these six states as a function of time. It's kind of complicated to watch all this at once, but while I'm nodding and smiling it says I look like I'm agreeing. I want to emphasize, too, this is not telling anybody what I'm truly feeling, it's simply telling them what I look like I'm displaying, so you can fake it just like you can fake out other people socially.

Accuracy of these things is a challenge to test. The hardest test we've done right now is – it's trained on these British actors and we want to test it under a completely different situation with completely different people, different lighting, people who aren't trained and so that's the results here. The basic idea is to look at these six states when people are trying to communicate them and see what a group of 18 people received. Okay, a bunch of people are communicating them on video, 18 people rate what they think the person was communicating, we take that as our gold standard and that's what's shown here. It would be perfect if it were 100% across the diagonal here but instead what we see is disagreeing is most reliably recognized followed by agreeing followed by some of these others, concentrated, interested and so forth. The computer algorithm tested on the same data right now that it hadn't seen before is better than 17 out of 18 panelists. The thing that might strike you as surprising is how low both of these numbers are. We were surprised, too, to see how poorly people did. Some of this may be due to our presenters not sending the signals as clearly. In fact, when you train and test on different subsets of the British actors, the rates are up around 80-something percent, mid 80's. So it does a lot better with people who better communicate this information. This is closer to the real world situation.

We are in design trials right now trying to involve people autism in the design of the system. These people often have incredible insights into details that those of us who just do social interaction without thinking about it don't really realize are there. They have to systematize what's going on and pull it out in a way that is really what we have to do to give it to the computer, so we stand to learn and benefit a lot from involving these individuals early on in the design process. Maybe I have time for one more quick video. We are also trying to help them with their number one request right now which is monologue detection. I figure this is a good one to put at the end of the talk. The problem a lot of these individuals have is they're fascinated by what they're telling you about but they can't read your face. This is not live, this is pre-recorded. I'm having a dialogue here with a friend of mine. He admits he has a real hard time reading faces. He's a very interesting guy and throughout this dialogue I am listening to him and I'm

very interested in what he has to say. And the blue line is highest for concentrating and the red line is next highest because I nod a lot, I'm agreeing. But after a while you'll see that while he's sort of consistently concentrating and interested, he's not looking at me. And though I try and try to understand and follow what he's saying, he never pauses to give me a break. So while the press says, you know, it's going to detect if people are boring, he actually wasn't boring me, it's just that I needed to stop and think and he wasn't reading that signal. So you'll see after a while that you can listen to part of this..... See what happens to me after awhile, but he keeps going. Now he goes on and on and finally I stopped him and I said, you didn't see what I did, did you? No. So actually our number one request right now especially from adults with Asberger's and high functioning, is could you just give me a back off signal or something that tells me when to pause, or when the other person is not interested. Or actually in this case it wasn't not interested, I just needed to think and process so I needed him to pause and let me go, oh, okay. Something like that. We are, as I said, putting this into mobile form so we have it in a hat, a self cam where people can look at themselves if they're not good at looking at other people's faces because we think maybe if they could see the association between their own face and their own feelings it might inspire them to look for that information in other people's faces. There's a basic fundamental question that's unanswered in this area of why these individuals don't find it intrinsically rewarding to look at people's faces, why is it that they avoid faces. And then we're building a number of different form factors to give information back, and we are also just recently talking to some experts in helping the blind and possibly putting this into future systems for the blind. One of my colleagues who works a lot with the blind was talking with one of his unsighted friends about what mattered most that he missed in life and he said, you know, it's not that I can't see movies or the Grand Canyon, stuff like that; he said, the thing I miss most is that I can't read the face of the person in front of me. I just can't see how they're receiving and feeling about what I'm trying to share with them. So we can actually do that in real time right now with a little bit of audio feedback for a limited number of states. Again, it's far from perfect but it is performing well in this current evaluation as 17 out of 18 people.

I'll just summarize. We've quickly gone through a whole bunch of things, probably raised a lot of questions so I want to just give you a URL and invite you to come and learn more. There's papers and a lot more detail on each of these topics. Thank you.

*To access the most accurate and relevant health information that affects you, your family and your community, please visit [www.cdc.gov](http://www.cdc.gov).*