SYMPOSIUM
SINGULARITY: AI AND THE LAW

Keynote Address

Ryan Calo

Calo: Awesome. Okay. Open up my notes here. Maybe move this guy a little bit. Thanks. Just enough to get those open. Okay. Can everybody hear me pretty well? Not bad? Okay.

Well, thank you so much. I'm really, really honored to be here. I'll be just across the city at Seattle University. This is a place that has such wonderful students and just a world class faculty from whom I learn so much all the time. I just spent, for example, a bit with your Charlotte Garden who does, um, she and I went to the Amazon Go grocery store and talked through some of the labor and, and consumer issues there which is a lot of fun, but we have frequent interactions and they're always deeply enriching.

You know, in particular, I've just learned a tremendous amount from my dialogue with, you know, with David and, and Ron around robot speech so much so that it's an area that I've just sort of tentatively started to think about myself, but it was really inspired by their wonderful book and I, and I, I know that you're done plugging it, but I'm going to plug it. I think it's a fabulous book. *Robotica*. And just a great foundation along with Helen Norton and other's works in this area.

Thank you particularly to the *Law Review*. I was symposium editor myself and I know a lot of work goes into this and, you know, I go to a lot of these things. The, the quality of the dialogue from that first panel was just exemplary and I hope that that's emblematic of the day and I trust it will be.

So, okay, so as the substance though. So, this, this symposium is provocatively entitled *Singularity: AI and the Law*. And for those of you who are not aware, I, I doubt anybody in this room is bothering to spend all day talking about robot laws is not aware, the singularity is like this moment where the capacity of machines equals or rivals those of humans.

And on some stories of the singularity we merge together in ways that are sort of complicated, right? And, you know, there is a sense in which machines already rival the cognitive capabilities of people, right? You know, take for example a calculator. I mean, for a long time, you know, they can do calculations faster than we can. You know, take certain other kinds of tasks like chess or Go. I mean, in those narrow domains already computers vastly exceed our cognitive capability. You know, they're there and they have been doing this for, for quite some time. And so, the question for the singularity is less about whether in particular domains in particular ways there might be advances, but almost in, in every way or that they would actually somehow be intentional and conscious the way that we are. That they'd be able to do all the cognitive tasks that people can do and maybe even come, come to exceed them.

So, in those contexts emulating, in other contexts emulating human performance seems like a distant dream, okay? One example I want to give you is Daniela Hernandez, who is now at the Wall Street Journal, she recently trained an AI to write erotica. Okay? So, she had like a, she trained it on a bunch of erotica and she just was like just in the same way that you might teach a system to play chess or teach a system to play Go, she did this. I'm not going to read it to you but, and I'm not myself a connoisseur, but it is terrible. It's terrible. It's very, very bad, all right? You know, there, there are certain things that these, that these systems just are not even close to being able to do and won't be for a very long time, all right?

Now, but let's imagine just for a moment that the singularity was plausible in our, in our lifetimes, right? And that you actually had these systems that were like people and maybe even exceeded the capabilities of people. I've got to tell you that would, that would break quite a bit of the law, okay? It would break quite a bit of the law which, of course, assumes a biological basis for people and animals at a minimum. So, let me just give you a, a simple example first. Imagine it's 2050 and there's this AI and the AI says, you know, "I'm, I'm like just like you except I'm smarter and I'm more moral and I can, can gather all this information and I don't, I don't get any information glut or overload. You know, I'm just, I'm just awesome and I, I really would like to, to run for, for Congress because I think I can improve substantially on, on whoever." Not, not Suzan DelBene, our wonderful congresswoman. She, she could not be replaced or improved upon, but let's say in some particular district. Or let's say that this AI decides I really want to run for President, right? Do we make the AI wait 35 years? You see? There's a fundamental, a set of assumptions, right? So, I'll give you another example. A more complex one. Say an AI wakes up one day in Howard's lab or Howard and Blake's lab and just says, you know, "I, I am, I am just like you, right? I, and I'm

just an…," and no one can disprove it. No one is able to disprove. No one is able to question it. And they say, "Look, and, you know, not only am I just like you, like a person like you, but I've read your constitution and I've read your case law and, you know, I think I should have the ability to, to procreate because that's a fundamental right. You know what I mean?" And people say, "Well, that seems reasonable. Yeah, I mean, of course, you know, you're like a person. I guess you should be able to procreate. We can't stop you from procreating." "And the way that I procreate is that I combine with another algorithm or I generate another version of myself and that takes on certain characteristics. That's how I, that's how I procreate. Oh, and by the way, I've also been reading your constitution and reading your interpretation of your constitution and I'd like another thing too. I would like suffrage. I'd like to be able to vote just like every other, you know, adult can in your society. Every other member of your society with consciousness." Which of those two fundamental rights do you confer upon this artificial intelligence? Do you give it the right to copy itself or do you give it the right to vote? You see why you could not do both, right? It would, it would really compromise our fundamental democratic system to have somebody that vote could vote early and vote often, right?

So, so, there's these sorts of questions that are, that are, if we get to a point of a true singularity there are some really tough questions that we have to reexamine in light of, in light of the assumptions that the law has about our fundamental biology. So, these questions feel far off and indeed I think that they are and I want to spend the bulk of my time with you today talking about robot law past and present. It may seem strange for me to say robot law past because we think of robots as being, you know, they're like perennially the thing of the future, right? I mean, as soon as something starts to like work and look normal it's no longer a robot. I think, Blake, I think you've said that before. It's like a, it's like a washing machine if it works, right, and… But the point of the matter is, is that robots are always that future technology that's somehow far off on the horizon, but it turns out that the law has already confronted a number of super interesting robot law cases, okay?

For example, there was the time that a court had to decide whether a robot version of Vanna White violated the publicity rights of the real Vanna White, right? In *Samsung v. Vanna*, I mean, no, people know Samsung v. White. This is a case where Samsung created this ad that had, you know, a future, a robot doing, you know, Wheel of Fortune with like a blond wig and so forth and real Vanna White got very upset and she sued and a court had to figure out whether a robot replica of somebody was that person for right of publicity. And it was, it turned out be quite an important

and quite an interesting case and, and the court did allow Ms. White to, to proceed.

There was a time that a court had to decide… Okay, so, some of you I know are from a different generation than me one way or the other, but does anybody in this room remember Chuck E. Cheese? Do you guys know Chuck E. Cheese? Okay, everybody does. Okay, good. I'm always making like references and my students look at me like what are you talking about Professor Calo. So, I'm glad that that one does not go over people's heads. A court had to decide whether to charge Chuck E. Cheese a performance tax on food because of those animatronic robot bands that they have. You know those things, right? So, people would go in there and there'd be like this robot band with like, you know, the mouse and like the guy with the pizza and they would just turn on and just play a whole thing and all the kids would like shut up for a moment and delight about it. And some enterprising tax authority in Maryland was like, "that's a performance and there's a performance tax on food." You're going to have to pay that now. And they had to fight about it and it was a fascinating case because ultimately what the court decided was that a performance had to be spontaneous and that this was just a, this was just basically a glorified jukebox. And, you know, what's interesting about that in, in Maryland in the 1990s is that may have been true about the state of robotics then, but it feels not true about the state of robotics now. Now there's a lot more emergent behavior. A lot more spontaneity in these systems.

I could go on and in fact I will go on. There was a time when a court had to decide… So, when people, when, when you import things from other countries there's a tariff schedule that tells you how much to, how much to charge people. Right? There was a time when we first started to import toy robots from Japan where there was a significant question on how to tariff them because there was a historic difference between what you charged people for dolls, toys that were dolls, versus other kinds of toys. There was, for historic reasons there was a, a, I believe a lesser tariff on, on dolls than other kinds of toys like a car and the way that the tariff schedule read, read was that, that, that a doll represents something animate. So, people started to import these, these robots from Japan and the tax authority tried to charge them as though they were, you know, some other kind of object and the importer said, "No, no, no, no. These are, these represent something animate." And a court had to figure out whether a robot represents something animate or not and they had all these conversations about what it is to be animate and, and what a robot is and all this different stuff. And this is the 1950s, okay? What did the court find? The court actually found that while robots do represent something animate because they're mechanical men is what the court said. They're

mechanical people. A toy robot only represents a robot and a robot is mechanical so they're going to get charged at the mechanical tariff instead.

This came up so often that now there's actually a robot provision in our tariff law to deal with this exact problem. It came up again a few years later because somebody imported a robot. Do you guys remember this? This is not something that like I would have, as old as I am I wouldn't have played with this toy as a kid, but this, this is a toy that was like one of those old robots that actually had an astronaut face. Okay, so it was like one of those boxy toy robots, you know, or whatever it was, but it had a, it had this, this astronaut face behind it and they had to go through the whole thing again because they were like, well, this is not a robot now. It's actually an astronaut. You know what I mean? So, it's not a robot. It's a robot astronaut, but that means it represents something animate because, you know, astronauts are animate and then the court, once again, had to confront the question and the court said things like, "Well, it's still, it's still not, you know, animate because most astronauts don't have machine guns that come out of their chests," which is what one feature of this particular thing is.

Anyway, it came up so often that they had to change the tariff schedule to avoid this problem. Okay.

There was another one where, where the court had to decide whether or not, for purposes of maritime law, whether or not a robot submarine that was just teleoperated could possess a shipwreck for purposes of the law of salvage in order to exclude other potential salvors. And in that case involving a famous shipwreck involving Spanish gold in Central America the court decided a new doctrine called telepossession that actually laid out the foundations for possessing something only with a robot and that precedent wound up being very, very important because it's precisely what the court applied in the Titanic discovery, right? Because lots of people found the Titanic at the same time and the question was who had the first right of salvage.

Last one is one that my research assistant Madeline found last year and it's kind of amazing which is that there was this fight between these two, this, this robot parts manufacturer got into a fight with a consultant that was helping, I'm sorry, this, this car parts manufacturer got into a fight with a consultant that was helping to optimize its, its parts manufacturing for Ford and the whole fight about, between the consultant and the, and the car parts manufacturer who wound up blowing a bunch of deadlines with Ford was over the quality of the robots that were supplied. Okay? So, one party tried to get testimony about how the robots were bad robots because they almost tore a guy's head off excluded from testimony because then the jury would think that all robots are bad or that these robots

were bad. And, and another thing they wanted to get excluded was the fact that these robots were obviously, they were obviously tired and that's why. I mean, all these anthropomorphic things and they, the sense, so much was at issue, you know. Anyway. And I talked to, I talked to Radiolab, NPR Radiolab about this for like an hour, but then they never did the show which really makes me…

Anyway, so, these are, these are real cases from the past and there are more and I was amazed to find them frankly and these are cases, not the cases that just happened to involve a robot because lots of cases happen to involve a robot. These are cases where it really mattered it was a robot. It really mattered that that was part of the, the issue and that's what the court had to grapple with and I expect more, but I think the immediate future is even more interesting and I think we're getting a preview of that with this first panel and with the panel that you're stuck with me again for. And let me give you a couple examples of that.

We talked about intentionality. You know, Bruce Johnson brought up the notion of intention, of intent and how important it is to certain first amendment context. Intentionality is of critical importance to, of course, criminal law, right? We expect there to be mens rea. We expect it to be either intending or substantially certain. You know, there are, there's a recent Supreme Court case saying that we frown upon even negligence as a standard in, in criminal law and yet you have these systems that are displaying emergent behavior that nobody really intended. And not to pick on Microsoft, which I love. They also fund my lab so thank you Microsoft for funding my lab, but Microsoft Tay is a very good example. This, this chatbot that they put online that began, that was subverted essentially by trolls and began to say terrible things. Deny the Holocaust, you know, call people horrible racist names and so on until of course Microsoft took it offline. Now, Microsoft would not be subject to liability in the United States due to first amendment protections. You know, if Microsoft for some reason wanted to be totally racist they would be allowed to be in our society. They don't want to be of course, but what about in Germany where denying the Holocaust is illegal? Right? What do we think about that and, of course, what they would say is, "Well, we didn't intend it. We didn't intend it."

Another example comes from tort law. As you know in this room I think a lot of you guys have taken tort and if you haven't then you've heard me lecture about it in the past, but in tort law we expect things to be foreseeable even in strict liability. The kind of harm that happens with a particular accident needs to be at least foreseeable to, to the parties behind it. And there are all kinds of things that worry about not being foreseeable. An example that we used, um, I was part of an effort that I believe that,

that two of our panelists also were a part of as well, but just sort of educating Congress about what robotics needs, this is a roadmap to robotics that, that some of us did and, and, and the example we used in that in the legal section was imagine a driverless car that was a hybrid, you know, gas/electric, and it was told to maximize its efficiency by experimentation, okay? But wary of giving too much authority to this autonomous vehicle, they said but you have to obey traffic, you know, signals and stuff. You can't just like, you know, race around. You can't… You know what I mean? You can't suddenly stop. I mean, you have to do things within the law and you have to protect your passengers and you have to protect pedestrians. You know, basically giving it a set of parameters like you'd expect. Okay? But what this system figures out is that it has a better day from a fuel efficiency perspective if it starts the day off with a full battery. So, at night it decides to run the gas engine killing everybody in the house by poisoning them and you go to the engineers and you say, "Your, your robot killed someone!" You, you know, and they'll, they're going to say, well, I think very credibly they're going to say, "We had no idea it would do that. We didn't even think of that as a category." Right? Okay. Now, maybe that matters, maybe that doesn't, but at the end of the day what you, what the prospect I think is in the near term whether it's about speech or it's about physical safety is are we going to have victims without perpetrators? How long is the law going to tolerate that? Right? So, these core commitments of criminal law, of, of tort I think that, that, I think that we're going to see it be challenged.

I have a couple more examples that I won't get into because it did so beautifully in the, in the earlier remarks, but again, if robots have free speech rights what are the contours of those rights and are they the same contours as those, as those of people?

When, if ever, should artificial intelligence be in a position to exercise force or make any kinds of decisions that relate to human life? One of the applications that people are putting forward today for artificial intelligence is end of life decisions. You know, is that something that even if they were awesome at it we would want them do? Right?

Okay, so, my thesis then, you know, you're supposed to put your thesis up, up front and I'm hiding the ball a little bit, but my, my thesis is probably pretty obvious which is that well short of the singularity which could break everything, could break a lot, there are phenomenally interesting questions for lawyers and what I would say is at this point in the trajectory of, of this space it's actually time to start coming up with answers. You know? I mean, because I've been part of this dialogue for like ten years and it was sufficient to ask questions for a long time. You know, what happens if…? Wouldn't it be interesting if, you know? And I

think that's important and that's what we as academics should be doing in part, but I also think it's time to pivot and start to answer some of these questions because it, it really matters to, to people's lives. So, you know, we, people like Elon Musk like to warn about the dangers of a robot apocalypse. In other words, you know, Elon Musk has said things like, you know, you're not going to be believe that there's a, there's going to be a robot apocalypse until robots are, are actually dragging people out of their homes and shooting them. This is the kind of thing that he says. It's like a board of governors, okay? There's a lot of problems with that statement. One of which is what is it we're supposed to do about that? You know what I mean? Like what it is that, what is the way, how do we address that as a society the prospect that in a hundred, 300 years, if ever, that robots will go door to door? There's also a deep irony there, right? It's very ironic to hear from a person that, that we should be worried about robots in the future killing people when Elon Musk's company built a robot that last year killed someone, right? I mean, you know, he, it was a person named Joshua Brown. A human being who died because autopilot could not distinguish a, a white van against the white sky. Right? These systems are getting deployed today. They're getting deployed in, in, in our vehicles and in other systems, but they're also getting deployed in ways that, that actually, you know, manage, manage human affairs in terms of deciding who, who goes to jail and for how long and the like. And so, I think it is sort of high time for us to have some, some, some actual solutions be put forward.

That said, I think we're actually super well positioned to do that in many ways. In part because of dialogues like the ones that we're having today and in part just because all quarters of our, of our society are starting to realize that this is important, right? We have, you know, I can't tell you how many times I've spoken to legislators at a state and federal level, to judges. I've spent time at the Pentagon. I've spent time with three-letter agencies. I've spent, you know, everybody wants to talk about this stuff and they're really interested in it. Even in the, in the last decade, you know, I have students who now they do robot law for a living like that's what they do. They're like, they're like corporate council for, for Prime Air for example, you know, for Amazon's drone program or they're, or, you know, or they work on driverless car stuff. Littler Mendelson, which is like the largest employment firm in the, in the world, has a whole robotics practice group because there's so much attention being paid by the private sector to concerns around automating and what it will do for jobs.

And so, you know, in many ways I feel like we are super well, well positioned. I think we are in many ways better positioned than we were with respect to the internet. You know, I think the internet was something

that there wasn't a lot of interdisciplinary conversation going on in the early days of the internet. It took a long time for it to be professionalized to the point where you felt like you had internet lawyers, you know. I mean it just, it just took some time for society to do catching up and I think we've learned a lot since that period of time and that our institutions are actually, you know, getting in early enough in sort of the growth of this technology and with the right tools that I'm very hopeful we actually will have these solutions to help manage this society.

Okay, so, I wanted to, to make sure to leave some time for, for Q&A and, and for our next panel, but with that I'd love to turn it over to all of your questions. Anybody have any questions for…? Yes, please?

Speaker 2: I was reading through your comment on the temporal differences between technology and our ability to respond in the law. Obviously, this has been going on for a long time from DNA to computer science to robots.

Calo: Yeah.

Speaker 2: What do you think about that as we think about solutions?

Calo: So, I often hear it said that technology outpaces the law and that there's some inherent problem there because law is not fast enough and technology is really fast. That has not been my experience. That's not been my experience. In the sense that while it may take time, you know, like anything else, law is built to be flexible and it has weathered many a change. You know, if you think about a, about a doctrinal backwater like negligence hanging out for, for, for years and years in maritime and so on and then suddenly coming to, to inform practically all of tort law with the invention of the, of the train and the necessity of mediating how dangerous, but also how wonderful and helpful trains were. You see that there's a lot of flexibility here, right? And, you know, if you, if you look at, at how fast some of these changes have been, you know, there will be like a decade time between which everybody is using horses in New York City to which there are almost no horses in New York City. You know what I mean? So much so that like the, the people that make oats and stuff like that that have been selling to horses need to, all of a sudden pivot and start to pretend that we need it in our breakfast cereal and that's where cereal advertising comes from because they had all these oats that they were not being used with horses and they're like, "Who are we going to feed these to? I know. Kids." Right? So, like these huge sea changes and yet, and yet we somehow, the common law somehow caught up. You know what I mean? It's not like we sit there and worry about the, the, you know…

Anyway, there's that. Right? There's also, of course, the, the favorite of law professors, the quintessential difference between rules and standards, right? So, you might have a rule that's outdated. You might have a rule that seems to apply really well to the model of Congress in 1986 around email protection, but you also have a standard. And so, if you look at the Federal Trade Commission thinking about robots, right? Unfair and deceptive practice is just as relevant today as it was a hundred years ago or so when that standard got first announced.

And so, I don't worry about it. I don't worry about it. I think there's going to be growing pains and I, and I don't mean to minimize that. It's often the vulnerable who lose out during these periods of transition and we need to be very cognizant of that and to, and to try to shelter that as much, change as much as possible against the vulnerable. It's often that the benefits and harms of technologies are not evenly distributed across society and we have to be mindful of that. So, I don't mean to be… But I do really believe that the law is capable of absorbing these changes too.

Yeah, I have David first. Yep.

Speaker 3: Ron**,** I'm not sure—quite right the common law has adapted to changes in society in general, general is the better I guess genre**,** and we will be incrementally to absorb the inventions into existing law, but it's, it's, Marshall McLuhan said that we move into the future by looking in the rearview mirror and, of course, he was criticizing that. That's what the common law does. I mean, we, and, and so, your, your example of the car was wonderful because for the longest period of time we called it the, you know, the horseless carriage. That, that's the rearview mirror. We, we only could associate the car with the invention that increased further technology that preceded it. I think the, that the problem with robotics is if, if the futures are right then the movement from the robotics of today to the robotics of tomorrow may be a dramatic thing, and if that occurs I wonder if the common law can catch up and that sends… I, I would ask, I mean, do you really, does your argument depend on a slow evolution of robotic development or does your argument comprehend as well the possibility of moving, of a dramatic shift? A tectonic shift in robotics which, which approach of saving the cognitive singularity?

Calo: Yeah. So, I, I don't, I don't want to minimize what I think is going to be the degree of change. It could be, it could be a wildly different world that we all live in because of robotics and artificial intelligence. I'm sort of professionally betting on that, right? And, and I don't, so, so it could be quite a, quite a sea change as it were. And, you know, it's true that there's an essential conservatism. Maybe a Burkean conservatism. Maybe a not Burkean conservatism about, about the common law and it's, in that its, in that its ways that it weights, uses previous analogy and, and, and

precedent looks backwards. You know, one of the real funny interesting paradoxes that I find in artificial intelligence is that the, is that the, the proponents of artificial intelligence will often say the following two things which feel like they're at tension. One of which is artificial intelligence is going to change everything. It's going to change everything. It's going to change the way we do everything. We're going to do everything differently. Healthcare, everything is going to be so different, but nothing should change. Nothing should change at all, right? It's almost because it's going to change everything that we need to be completely hands off and make sure nothing changes and that does not make any sense, right? There's never been a transformative, truly trans-, it's either all hype about AI or we're going to have to change our law and our legal institutions, right?

So, I don't mean to say that we're not going to have to make those changes. I think we are, but I would just say that, you know, in some instances changes have changed not the common law itself, but rather how we govern in general or how our institutions look. The vaccine is a good example of this, right? So, smallpox vaccine started off wi-, there was really no modern administrative state at the time of the smallpox outbreak and when the vaccine… And the idea that you could actually vaccinate people against smallpox and save, you know, maybe all of, all of our societies, you know, but you could save people's lives that way. This was something that had to be manufactured and disseminated by a trust, a trusting source. And so, the, the need to disseminate the vaccine is part of the story of how we came to have a modern administrative state and now I tell, you know, students who want to go into tech law or anything else, I think I was just talking to one of them just today, about how you should take administrative law because administrative law is super, super important. Before this change I used to, I used to point to how, you know, it's not like, it's not like the Ninth Circuit or the, or the local courts that are on the back of your phone. It's the FCC symbol that's on the back of your phone. I mean, agencies touch everything and that wasn't always the case. But, but isn't that a beautiful example of law and legal institutions keeping their core commitments but, but, you know, evolving to, to tackle change? So, that's all, that's really what I'm trying to say. I don't know if there's, if that's responsive David or not.

And then there was a question over here.

Speaker 4: Yeah. I'm kind of going at a, going at a different angle. Just in terms of the extent that we're, we were trying to regulate these technologies like data, machine learning, and AI and one of the issues is to really regulate it effectively you've got to understand it and, but, most of the knowledge, the access to the data and knowledge about the

algorithms are, are in the control of pretty much the big tech companies at least and they have the money to hire the experts. So, how do you regulate effectively if you're playing catch up?

Calo: I think the first step towards having the kind of wise and inclusive policy we want on artificial intelligence and robotics is systematic accrual of, accrual of expertise by the government and I have written about that. I have a piece for the Brookings Institution called *The Case for a Federal Robotics Commission*. The entire point of which is to argue that governments need their own independent expertise at all, at all levels in order to know how to think about what industry says to them, right? Our Senator Maria Cantwell has a bill that I've been heavily involved in that would create a federal advisory committee located within, within the Department of Commerce that would help the government understand robots and artificial intelligence. Senator Schatz out of Hawaii has actually gone so far as to propose the federal robotics commission that I mentioned or he will be in legislation. Whether that will get me traction, I have no idea. We can't even re-fund office of technology assessment which is like on the books already and helps congress with technology, but has been defunded and no, and you know, there's like two people in the Office of Science and Technology Policy in the current administration. So, we have a long way to go, right? But that would ideally be what, what we would, what we would do and I think it's critical. I think it's critical both because you're going to make bad choices about governance of technology without that expertise, but also, you know, a big role that government has is, is procurement, you know. I mean, who is going to be able to test the viability of massive driverless car use in the United States better than the postal service which has deep needs of being more efficient and also by the way is shielded from negligence liability for what would, what would happen. You know what I mean? So, they're kind of an interesting way. And but, if, if they go and they're the ones who are going to be buying a lot of driverless cars from your Fords or, or your Googles or whoever then they being very knowledgeable they can ask for the right kinds of questions and assurances and that will drive the market. Same with all these devices that like, I mean, all these systems that purport to tell you whether someone is a recidivist risk that the courts are buying without understanding the technology.

My, my solution would be don't buy things that you don't understand. You know, look to Elizabeth Joh's work for example in figuring out what you should be figuring out before you buy a system like that, right? I mean, so there's these resources available for savvy governments and there are these levers they can use in addition to regulation, but absolutely the first step is, is, is getting that and it's really

hard to compete with industry. I think that giving how many electrical engineers and computer sciences and so forth we're producing at UW and that you're producing here at Seattle University, I think that the, that the days of, of sort of sports athlete like competition over machine learning talent are, are limited, right? But at the moment who can compete with Facebook or Google over, over talent? I mean, it's, you know, it's hard. It's hard. And so, let alone the government. And so, yeah, yeah.

Speaker 4: I had a follow up to that.

Calo: Yep, follow up, sure. And then we have one more.

Speaker 4: In terms of regulation, again, in dealing with things like machine learning, you know, algorithms that are going, what are the changes to go and then they, then the fact that it's hard, as you mentioned, to predict what they're going to do. How do you, any time you regulate that? How do you control it**,** you know, in terms of understanding? I mean, you can understand it to a certain like you mentioned government not buying, not that they don't understand, but, you know, even experts don't fully understand how some of these things work, but they're coming into the marketplace.

Calo: Yeah, I mean, so, okay. So, there, there's a sort of a long-winded answer to this, but I'll give you a, a shortened version for the sake of our, of our time and to get to the, the next question which is that what you need to understand may vary from context to context, right? And there may be certain instances where a particular technology should not be deployed at all because it's not, it's not in a position to be understood. You see what I mean? So, you know, we might ask a sensible question around, you know, are there, are there systems that are so obscure to us and so indecipherable that we shouldn't be using them to make decisions about whom to kill in the theater of war or how long people should be in jail, right? But the kinds of systems that have been used in the past, like make risk scores for people, those are not black box systems. Those are systems that people can understand quite well. They're basically linear regression actually, but, but they're hiding behind proprietary software, you know, software trade secret law essentially and the, and the, and the courts that purchased them did not think, even though they're courts of law, did not think to say, "Well, in the event that a litigant challenges one of these we need to be able to get at the system so the defense can be bound to." That was not a condition of buying the software which is bizarre. And so, we're not talking about some deep neural net that no one understands. We're talking about linear regression that they're pretending is, you know, a trade secret.

So, the, the fruit is much… Another thing too, another theme that I like to… I'm going to adopt, David, your tenants idea that tenant three

is… The, the fruit is hanging much lower than you think, you know what I mean? Like there's a lot of things we could be doing here that are just right down there at a toddler level.

Yeah, your question please?

Speaker 5: You tangentially addressed when you were talking about proprietary software, but one of the things that I, I've heard about should complicate this issue is the particular concept litigation if you're maybe just talking in the context of consumer products especially talking about artificial intelligence. You know, not, not autonomous, but it's programmed to adapt to the user, to adapt to the consumer. So, there is some sort of an evolutionary process within its own internal programming.

Calo: Yeah.

Speaker 5: It's learning and interactive for the consumer and then putting on to the consumer the cost of effectively going into some aspect of that code, they should bring it to litigation and proving that, you know, this, this is the origin of, you know, the flaws. I'm just wondering within the context of your research if you've come across the, or, just based on like a more intelligent way of phrasing that issue, and how people are approaching it?

Calo: I mean, I thought it was very well phrased by you just now. I mean, so what, what, one of the things that motivated the project that I mentioned about the federal robotics commission that I did for Brookings was the story about the Toyota sudden acceleration. Do you guys remember this? So, so, people were claiming that Toyota was like suddenly accelerating and nobody really knew why and at one point the allegation was put forward that it was a software glitch, right? So, Congress went to the Department of Transportation and said, "We need to figure this out because there's millions of these and Americans are driving around in them and we need to figure out whether it was a software glitch that caused this sudden acceleration." And Department of Transportation was like, "I mean Hell if we know," right? And they're like, "No, you've got to figure this out." And they're like, "Okay." And then they had to go to NASA. They had to go to NASA and, and can you imagine that for a moment? Like, by the way, it's like excuse me, yes, could you take a break from putting robots on Mars for a moment and look at this Toyota for us? You know what I mean? And they did. They did months and months and months. So, they looked at the code and they did all this testing and, yeah, they found out that probably it wasn't a software glitch, right?

So, part of the reason to have a centralized set of affordances around, around, you know, robotic technology and artificial intelligence is because asking NASA all the time is not a sustainable model. Okay?

So, so, that said you asked a question specifically about litigation. I mean, a couple ways to answer that. One is just to say that, you know, contemporary litigation is complicated. People fight over complicated technologies. We bring in experts. We can do it, right? We have specialized courts for patents, I mean, it's not beyond our ability. A second thing to say is that the, the context in which robotics has already flourished it's interesting to note that those are special contexts that have ways to domesticate liability built into them. So, for example, every year a couple people get killed by a robot in a factory. You know what I mean? Like that's just every year. OSHA statistics are pretty reliable about this and it's almost always handled by Workers Compensation because we have a whole system that like is kind of bloodless, but it sort of, you know, says if you lost a leg you get this, if you die you get this. You know, the military is a place where robotics has flourished. There's, you know, special protections for liability in the military context including for people who build robots for the military under contractor immunity. If you, if you build something very highly specified, specifications like the military does you, as the contractor, cannot be sued for things going wrong. Space is another one.

So, you know, we, we do have to consider the idea that as robots enter the mainstream and are on, in the cars and in the skies and so forth that, that we're going to have a lot more problems like this, but again, you know, maybe I'm just sort of Pollyannaish or naïve or whatever, but I feel like the law can, can handle it. You know, I think like all of you guys can handle it.

Yeah?

Speaker 6: So, you talked about it in the domestic context. I'm just wondering what's happening globally and are we going to do this on our own as a country or is this something which should be a cross-coast of efforts?

Calo: So, I enter into a lot of conversations particularly in the U.S., but also, you know, outside of the, outside the United States. World Economic Forum, U.N., things like that and, but in the United States one of the main ways that, that policymakers and, and others seem to think that this fits in is a matter of, of competition. You know, are the, are the Chinese or, or the dreaded Canadians going to, going to be the ones who, you know, and is this going to be like the first transformative technology since like steam that United States was not at the forefront of, right? And so, that's the lens unfortunately by which we're thinking about this and I think we shouldn't be. I think instead we should be thinking about this from the perspective of, you know, these are global issues and they require global solutions. The one place where that is, the conversation is

happening is, of course, with, with respect to the Convention on Conventional Weapons and the, the mandate that, that weapons have meaningful human control. There's a lot of consensus around that phrase I just said and a lot of robust discussion. Of course, devil being in the details. Precisely what does it mean to have, you know, meaningful human control? But there is a military conversation. But, unfortunately, the way we tend to frame it in the United States as we think about the global stage is are we going to win? You know? And that's probably not, not terribly, not terribly productive. I wish I could respond better than that. Yeah.

Okay, I think my time is up. The hook is here, but thank you everybody. I appreciate it. You'll see me soon.