

What Happens Next – Sunday February 7, 2021

Artificial Intelligence, Caregiving, the Penitentiary, Religion and Capitalism

Michael Littman, Charles Isbell, and Seyed Sajjadi QA

Larry Bernstein:

I want to open it with a question for you, Seyed, about your first point on decision-making processes. And just out of my own ignorance, I know that Boeing had, I think it was the 737 MAX plane that got into some trouble, where I imagined there was some flight management system that was not run by humans that was failing and the pilots were unable to take over control of the plane in sufficient time to prevent a crash. Can you explain, is it really just a failure in design vs. how the systems have been operated? And for you, is that an ethical issue, is that a problem, or is it just a failure in the system to meet its own objective?

Seyed Sajjadi:

Yeah. That's a good question. So, I think in the example that you mentioned is the failure probably came from the human-computer interaction and the teaming there wasn't great. And that is something that Charles was mentioning earlier. One of the biggest things that we need to be thinking about is how to best pair up a human with the AI. And rather than replacing the AI, replacing the human with the AI that is being worked on, how could we augment the pilot in this case?

And also to your point, Boeing, the technology that they've been working on for contingency management, the approach that they've had to the best of my knowledge on 787, the Dreamliner has been at least, we want an AI that can be treated as a co-pilot. So then that AI can work with the pilots while being mindful of the work that the pilot is doing, and then help them make the best decision that they have to make. We've seen a couple of these failure examples of the flight management system, where there's a contingency and the pilot will have to go through literally a 700-page manual book to follow a decision tree on how to handle that specific contingency, and this has been a very big problem. So, the approach that they have been taking is we need to be mindful of the cognitive load that the pilot's taking, and then use the AI to not add to that load, but rather have a virtual co-pilot that can help the pilot in that case.

I think the technical challenges in this case are can we use advanced systems that can explain their decision-making processes when they do make those decisions? And that's been sort of the challenge that's been happening in sort of explainable AI that's been a big bottleneck in getting a lot of these algorithms certified.

Larry Bernstein:

Charles, I want to bring you into the conversation. You talking about how combinations of humans and machines together are better, but when you think about catastrophic results where the human kind of either sort of checks out. I'm thinking of those Tesla drivers that are either playing a video game or doing something else than driving and crashing the vehicle. Or even in that example of a plane crash, there was that Air France flight that went from Brazil to France that crashed, and it had gotten mixed signals

from instruments and the computer had been, I guess, forcing it to go up when it should have been going down to regain speed and it fell out of the sky. How do you think about the fact that we human beings when we think a computer is only going to call on us in a moment of need maybe 1% or a hundredth of 1% of the time where we're both mentally and physically checked out and can't do our role in the teamwork?

Charles Isbell:

Well, I think that's definitely a key issue in the way that we build these systems. Conceptually, is not any different from the co-pilot not paying attention because he thinks the pilot is going to be making some decisions, and then suddenly he's called upon to make a decision. If we think of these as partnerships and by the way, as human beings, we're very likely to anthropomorphize the machines we're working with and sort of treat them in the same way that we would humans. I don't think the problem's actually any different. And by the way, it goes both ways, right? Although we had a celebrated or even a movie around a pilot that got a plane to land in a lake outside of New York, there's a lot of evidence that suggests that the machine actually would have done a better and safer job than the human did. So, these things can go either way.

I think I would actually phrase your question and make a point slightly differently, which is that we tend to think about these AI, just to get on this notion of narrow and general AI, we tend to think about these AIs, whether they're in partnerships or alone, as if they're going to become these general things that will be intelligent. And then we worry about those things. But that's not really the worry. The worry, to me, is that we can use these kinds of systems to make terrible decisions more efficiently. I mean, that's really ... when we look at the things that we're suffering with now, the problems that are coming out of the increasing use of AI and machine learning, it's that we are just encoding our own terrible decision making and just making it more efficient to do the work computers are very good at doing.

Larry Bernstein:

Michael, sometimes when I think of AI, I imagine that the machine is learning stuff along the way. We've just used the example of flying a plane. How do you think AI can learn about mistakes that are made in the air to become a better program? Is that something that Boeing uses right now? Is it gathering this information real time to improve the efficacy of its flying?

Michael Littman:

I really hope not. The people who deploy these kinds of automatic flight systems are ... they're safety engineers. And so they ... there've been decades of AI researchers really wanting to have a positive impact on these kinds of systems, talking to the people who actually field them, and being totally shot down because what they want to know, what the engineers want to know, is can you guarantee that under no circumstance or under no foreseeable circumstance, will this system drop the plane out of the sky or basically cause instability to happen. And we say, "Well, probably not." And they say, "That's not good enough for us. We're not going to let you touch our systems."

So most of the learning that's happening, certainly in these mission critical settings, is happening in a very batch or offline setting, where data from the flight is being collected and stored and then it becomes big data that can be processed by machine learners offline. But it's definitely not the case that

people are fielding, yeah, systems that are adapting in the sky to decide what to do next because, "Hey, maybe this'll work."

Larry Bernstein:

Following up on Seyed's point about we don't understand the logic of the algorithm. And I imagine that's going to be true with almost every one of these programs, whether it's trading stocks or flying a plane, or driving your car. At some point the levels of complexity, the human logical system can make no sense of the code itself or what it's doing. But sometimes the stakes are incredibly high, a plane could fall out of the air, we could run over a child with a car. How do we as human beings give up the logical construct where we don't know what we're doing and don't know why it's doing what it's doing, but we have an expectation that it is safer? Is that something that we are going to be able to do easily or not? As a public policy matter or as an individual making those decisions?

Seyed Sajjadi:

Yeah, I think there are a couple of different ways that that's been sort of looked into. One is to run a lot of tests. So, to the point that Charles made earlier, a lot of these systems are actually safer than humans. And you're seeing this as the biggest arguments that autonomous car makers are making, and that is, it's not that it's cooler for you to drive an autonomous or be in an autonomous car, it is literally safer for you because of the millions of experiments that we run from getting from point A to point B. They've shown that statistically it's safer for you to be in an autonomous car than not. And that's one side of it. I think that's probably the most logical thing to do to make sure that these systems are safer. For planes, it's also the same thing. So, getting the safety is running a lot of experiments during simulation or in real life, is the biggest thing that's been happening.

The other one, which is also interesting is looking into, as Michael was mentioning earlier, what are the imaginable situations within the input space? Meaning what are the situations that we can put the plane in or the car and that are sort of edge cases, and then test the results of it. The challenge there, however, is we cannot exhaust the search space because you cannot test for every single weather, for every single lightning, for every single sensor failure. And that's been a challenge. But we can, I think, relatively do a good job at that. We cannot explore the whole search space, but that is the second thing that's been also done to assure that the algorithms are performing well under those conditions.

Larry Bernstein:

And just going back to the ethical questions, and I know this is sort of a classic, we've always been asked the question, should we change the path? If we go in one direction, we're going to kill five, if we go in another direction, we kill two. And then we hope the computer kills two. But there's also all sorts of other intermediate steps where we take the probabilistic odds of running over a dog versus running over a person, even if it's incredibly trivial. How do we tell the computer or what are the ethical implications of making those judgments, whether it be people or animals or whatnot, as to decisions that make part of a decision tree? Seyed, why don't we start with you on that one as well?

Seyed Sajjadi:

Yeah, that's a tough one. And I think those are the questions that I think about a lot and I know a lot of other people that are designing AI systems are thinking about all the time. There are obvious ones like,

which is a game of numbers, okay, if I have two people on one side and five people on the other side, okay. Both are horrible choices, but I guess I'm going to go with the two. So there are the obvious ones that we can look at, but the challenge is that certainty or lack of certainty and not knowing that, okay, if I had gone to the other side, there could have been a 5% probability that I wouldn't have hit anyone. And then I think that would be really a challenge on how we do it. Generally speaking, I think we humans handle that probably intuitively. And for sure, that's the big area that needs a lot of further discussion for designing AI systems.

Larry Bernstein:

Maybe switching subjects a little bit to you, Michael Littman. You mentioned the Reader's Digest definitions of different types of programs and those teachers that inspire. I'm wondering, do you think ... to take Reader's Digest literally, do you think you imagine a world where AI will be in the education business and will find ways to inspire learning relative to other programs?

Michael Littman:

Right. Well, so just to clarify, the Reader's Digest quote is from decades ago and is actually about teaching, like people teaching people. It wasn't meant to be about computers. I just thought it echoed really nicely the different ways we have of instructing computers. Because in many ways, it's not that different from how we instruct each other.

That being said, the question about the role of AI in education is a really exciting topic. I think that there's been a lot of work over the decades towards a vision of, well, boy, why can't we have a smart computer act like a tutor? Like a personal tutor. People do seem to do better with their learning when they've got one-on-one attention. If we can program the computer to be smart about this, it ought to be able to do better than a teacher lecturing to a group of people.

And that hasn't really been the case to date. It's actually really hard. The thing that people do to connect and to teach is something that we don't know how to automate particularly well at this point. So at the moment this isn't showing huge gains, I think there's a lot ... there continues to be a lot of interest in trying to get this right. And the opportunities have increased of late.

I think in the early days, the good old fashioned AI days, the idea was, okay, well, I'm a smart programmer and I'm a good teacher, so I'm just going to write down as a computer program, all the great things that I do as a teacher and then the program will implement it and it will be fantastic. It turns out that we're not very good at reflecting on our own teaching just like we're not very good at reflecting on our own ability to well solve the trolley problem, which we're probably not very good at or recognize objects. Like how is it that you know that that's a wolf and not a dog? I don't know. It just looks more wolfy, right? And so teaching is a little bit like that.

The way that we've been handling those kinds of problems of late is by using data, by having systems learn to basically program themselves. That's the direction that we're starting to go with the teaching as well. The idea being that we can now gather a tremendous amount of information about how different students online react to different kinds of educational methods. It's like how much did someone's

understanding increase once they watched this video, or once they read this chapter, or once we asked them this quiz?

And so now that we have the ability to collect this kind of data, lots of people are starting to look at turning it around and turning it into a learning problem so that we can actually create systems that will make smarter decisions based on exactly that kind of experience. But it's still quite early at this point.

Larry Bernstein:

Charles, this question is for you. Sometimes art or fictional representations of reality can inform us as to some of the problems we're going to have in the future. And I've seen one play and the movie, *Her*, where we have a robot, an artificial intelligence machine, interacting with humans to improve the quality of life for the human being. In the play that I saw, we had a woman with dementia and she misses her husband and they have ... I'll call it a robotic-like figure, obviously driven by AI, which would recall stories and give comfort to the demented woman and act as a constant caregiver. And in the movie, *Her*, this AI character, dealt with the loneliness of single man and provided love and compassion. How do you think about this fictional representation of providing emotional support to humans? Do you find that just too fictional or is that something that's going to be solvable? Will this allow humans to find more meaning in their lives?

Charles Isbell:

I don't see why not. I mean, in principle ... in both of those cases, you're talking about companionship and you're talking about someone or something that can sort of play a role of being a partner in this emotional connection that we're looking for. I mean, there's a question of whether building a machine to do that as the right thing to do for the machine if the machine ... but those are the usual questions we end up in these conversations. But I think that's entirely ... it's not just entirely possible. It's a natural thing that we'll want to do.

I think all of the questions you've asked us recently and all the conversation that we just had, they're really all versions of the same thing. I mean, intuition is just another way of saying, I don't know what I'm doing or at least I can't explain it. Explainability is about making you trust who you are. These kinds of ideas of, can I educate you? Who's responsible for these outcomes? They're all really sort of circling around the same basic questions of what is the right way to deploy intelligence? And they're more about that than they are about whether the intelligence is possible or not. I actually find it useful when asking all of these kinds of questions to try to pop up a level, imagine about sort of how human beings would work if we had sort of perfect understanding and see how we feel about it.

So to go to your early example about autonomous cars and driving, the problem of self-driving autonomous cars and safety is actually very easy, if you just require that everybody drive an autonomous car. If there are no human beings, all the cars are talking to one another, and the stoplights are talking to the cars then this problem becomes significantly easier than the one we have now. We don't even come up against the trolley problem because most of it goes away, most of the time. But I don't know that we're willing to do that. And I don't know that we're willing to take the question you just asked me to its logical conclusion.

But let me just put the question back to you. The companionship is perfectly fine. It makes a lot of sense. I think it's something we'll be able to do. But maybe we should be building machines that are going to satisfy people's sexual needs. Maybe we should be building machines that would stand in for what we would all consider possibly immoral and certainly illegal activities that you would otherwise engage in. And that would solve potentially some problems. Do we even want to walk down that path and ask that question? These are policy questions. These are high-level issues of fundamental notions of ethics and morality. And they're much less about questions of whether taking the derivative is going to maximize your objective function.

Larry Bernstein:

Charles, just to follow you up on the ... I'll call it the transitional period for driverless cars, where you have both a combination of driverless cars and regular cars working together compatibly in the environment. Can you see a world ... imagine you were the public policymaker? It might be difficult to tell people they can't drive their existing car, but it might be much easier to persuade a trucking industry that they had to use driverless trucks. Can you imagine a world where we would have a driverless truck lane on the highways and that we would see these trucks, nose to nose, going at 80, 90 miles an hour through our urban environments, or going through, I should say, on the highway environments? Are we going to ... because we could build two separate systems, I'll call one an AI system and one a non-AI system, one driverless and one non-driverless? And you'll see the driverless system flying by as fast as it can and the non-driverless stuck in traffic, looking at those cars going by at 80 miles an hour, going, "All right, I give up. I got to go buy one of those new cars." How do you think about that transitional period?

Charles Isbell:

Yes. And in fact, by the way, not just trucks, but public transportation period. By the way we have this now. I mean, we have HOV lanes. And some people prefer the HOV lanes. And certainly, in cities like the ones I live in, it makes a massive difference in your daily commute. I think that if my goal, this is not necessarily my goal, but if my goal were to get everyone to be in autonomous vehicles, that's exactly what I would do. I would create very visible parallel mechanisms by which if one has an autonomous vehicle, one can get from here to there with no problem whatsoever. I would take advantage of the kind of governmental infrastructure to do it with public transportation. And I would take that one HOV lane and I would make it two and then I would make it three. And then eventually the trade-off would just be obvious. And then you would end up with, in parts of the country, small places, people who still wanted to drive around trucks or whatever would do that. But you would just make it so that the outcome is inevitable. I don't know that that's the world that I want to see. But if I were trying to get to a world past the transition, that's exactly what I would do.

Larry Bernstein:

Question for Michael. My question comes from one of our other speakers, Lawrence Friedman. He's wondering what you think AI ... how it would affect the job market? We have a number of relatively low skilled workers in our economy, and we're going to have some very bright AI related machines that may be able to do that work at less cost. How do you think about that interaction and what that will do for relative wages for low skilled workers?

Michael Littman:

Yeah, so I'm not an economist. This is definitely an economist question more than an AI question. But I've spoken with some economists and the sense that I get from them is that historically automation, and this is just in many ways just another form of automation, tends to result in a short-term shock to the system that has the negative repercussions that you were just listing. But then the market equilibrates, that people find ways of earning a living, and in the longer term, things end up being better for more people at the cost of these really difficult situation in the shorter term.

The big concern, as I see it, is when these shifts, these switches to automation happen in quick succession over multiple industries, without really a chance to equilibrate, that's what I think we need to be trying to build some policy around, to smooth that out. Because if we don't smooth it out, we could easily be in a situation where just things keep getting automated out from under us and people are constantly out of work essentially, because the job market hasn't re-equilibrated from the previous shock. It's just shock, shock, shock, shock, shock, shock.

And so I think that the government can play potentially a very powerful role in smoothing these things out. The upside of this story is essentially a town that was based very much on a manual labor economy wanted to automate. And what they decided to do was basically continue to pay all the people who had the manual labor jobs and just pay them until they retired. But all the new jobs, everything, all ... the new system all went with the newly automated version. And so it's really trying to get the bumps out for the people who are directly affected by the automation that's really hard to do, but potentially can make a huge difference in helping to just smooth out the labor market.

Lawrence Friedman:

Yeah. I think that the issue, the job issue, of course it's an economic issue, but it's also a political issue. I mean, the people are being displaced and it's okay to say that in the long run there'll be more jobs. We don't live in the long run. In the long run, we're dead. And so the dislocations which could be caused might produce, and I think have already produced, an enormous amount of political turmoil. We're always told that extremism is partly the result of people who are frustrated because the world they were comfortable with has gone away. And that very often is the labor world that they were comfortable with, the world of high-paying factory jobs.

So I think it's really not satisfactory just to say either that, "Oh, in the long run, there'll be more jobs," even if that's true or, "The government should play a role." Because what role should the government play? And one last point is, suppose that you have unemployed workers, they've lost their job, and you said, "We're going to pay them for the rest of their lives." But people don't want that. They want to work. They want to feel that they're doing something productive. So I think that there's a real problem here that has to be addressed.

Michael Littman:

I don't know that we're disagreeing. So yeah, I think that this will be a problem and we got to figure how to deal with it.

Larry Bernstein:

All right. I'll move on. I have another question for Seyed, I know you're potentially in the ad insertion business, and I imagine how it will work is, I'll be on Google, looking up what play I should see, and I haven't bought tickets yet, and I can't decide if I should go see the show. And then sure enough, when I'm playing a game, there'll be some sort of an insert reminding me that I should buy some tickets for that show. Now, Mitch was saying before that somehow when we get some recommendations, it doesn't feel creepy. But what sort of challenges do we have when ... where there's this intermingling of, I guess, personal information as it relates to something that's maybe in the open?

I know there's a classic case where a daughter had just gotten pregnant and she'd done some Google searches for diapers. And then in some other form, diaper advertising started surging on the computer and the father was able to ascertain that she was pregnant. And so how do we sort of limit or protect certain private information as it relates to, I'll call it search behavior, and then how new advertisements get inserted into a public space?

Seyed Sajjadi:

Yeah. To the best of my knowledge, and the example that you actually provided, that is a classical one, and that was from Target. And I don't think it was even on a computer. So this type of recommendation was being done before. So when you go to a store, your purchase and all the things that you've bought would be analyzed, and the coupons and the brochure that you get was being personalized in the example that you mentioned. At least the version of it that I read was that the father found in their mailbox, some diapers, and then through their mailbox, he figured that his daughter was pregnant. Not the best way to break the news.

And on the note of personal data, I think the very good trend that's been happening, the GDPR and CCPA, is to provide people with more control over their own data. So understanding who's using your data and who's using it for what purposes. I personally am happy that some of the not so intrusive ads that I see on different platforms, they help me with choosing a better product, whether that's for a health or a fitness application or any type of diets. Recently, I've been just getting a lot of keto diet ads. And I would say probably they have positively contributed to my overall health and wellbeing. And it is important for us to keep that trend and not really provide just intrusive ads that are not relevant to the user whatsoever just to charge the brands for the eyeballs.

With that, it's important for the users also to give it a little bit more attention to the user agreements that they sign and also keep track of their own data and understand ... Google now provides this. You can just go to your Google dashboard and then actually disable the use of your data on Google, and they can't sell it to any third party. This is already happening in Europe. It's already happening in California. But I think the rest of the country also needs to catch up very soon.

Larry Bernstein:

Charles, just to follow that same argument with you for a second. So let's imagine that facial recognition gets to be very, very good. And so you walk into Target. Target has learned from Google that you've recently been searching for jeans. And just like Michael's phone a minute ago, your phone explodes,

there's a coupon, a two for one special on jeans in aisle four matching something very similar to the one that you searched for. This is communication between a whole variety of, as you would describe it as, machine systems and the human itself. Do you see that as something as a force for good? Is that something that's going to be in our future for sure?

And then maybe just to layer on top of it, let's imagine that I've been arrested for theft. And so facial recognition sees me, I enter the store, a security guard comes over and says, "I'm sorry. You got to leave." "Why? What have I done?" "I don't know. It's unclear to me, but I've just been given a notice that your face is recognizable and therefore we're going to ask you to leave." And I don't know if I'll ever be able to say, "No, no, that's a doppelganger. That's someone who sort of looks like me." Or maybe I have attributes that make me more risky than other people for maybe committing a crime and I've actually done nothing wrong. How do you think about some of these moral layers as it applies to making predictions from algorithms that may make my life uncomfortable?

Charles Isbell:

Well, I'll just point out that your second example, we do that now. I don't know how many times you've been followed around in the store, but I've been followed around in store multiple times in my life. I've been accused of theft. I've been in this situation more than once. So we're doing that now. Again, this will be a case of making the decisions that people already make significantly more efficient. For the first example ... so that's going to happen. So by the way, the answer to both of your questions is, yes, this is going to happen. It's going to be widely available. And whether you think it's a good idea or a bad idea, I suspect will depend a lot upon whether it makes your life more convenient or not.

In particular with the first case you described here are the jeans, here's a two for one coupon, well, if the machine is most of the time right, and gives you a coupon that makes it easier for you to get the jeans you're looking for, Then you're going to be extremely happy with it. And if it's bad at it, you're going to decide that it's spam. I mean, the biggest problem right now is that clearly somebody who is following all of my searches is they keep offering me coupons for things I've already bought. Now, if they had done that before I bought them, maybe I would be okay with it. But it's ... whether it's useful or not is something else. So the question of the kind of ethics and morals around it, I think it's an interesting question and sort of, I think, the question for all of us in the future. But I think that we're going to answer that question in a very practical way about whether we think as individuals it helps us or not.

And I will point out that ... two things. One, the ads are very visible, but the amount of information that Tesla, Apple, Google, and everyone else we know and care about has on all of us is enormous. They know what we're going to do before we do it. And they're very good at leveraging that. And a lot of the times it's invisible. So it's unclear to me that we have really thought very hard about, or even have the tools and language to decide how we want to deal with this as a society and as a culture. But I think that in the end, we're just ... the truth is we're going to go with whatever's most convenient.

Michael Littman:

Michael Littman here. If I could just jump in for a second. I think Charles is focusing on the impact on, let's say the consumer. My worry, and there's already plenty to worry there, but my worry is that the decision makers here are the companies, so Target, say. And they're not making decisions necessarily

that are in our best interests. They're making decisions that are in their best interests. And I was wondering if Seyed had any thoughts about what are the forces that are encouraging companies that are deploying AI to be ethical and to concern themselves with the rights of people? Or are we in the middle of a losing battle?

Seyed Sajjadi:

Yeah, that's a great question. I think one is for sure, the basic one, regulations. So now for every single European customer, there's GDPR and the customers that have their data in California, you have the California Consumer Privacy Act, the CCPA, that's in place. And following the regulation, because otherwise there's a \$25 million penalty, that's the very first thing that I think all businesses want to follow. And then the second one, is the communication that comes with the partners. So us as AI company, we work with other businesses. We're not a consumer product company, we're a sort of B2B company working with those partners that are mindful of their customer data and they themselves are also following the regulations and are also open with their customers.

For instance, if we are working with a customer in the security space that has access to tens of millions of cameras around the States, we want to make sure that that business customer is disclosing to their consumers that, "Hey, by the way, we're sharing your data with this AI platform and they're going to help you provide ... help us provide you this intelligent solution." And having that disclosure, not on the 43rd page of the user agreement, I think is very important and being able to disclose that upfront.

The user agreement I think is very important, and being able to disclose that upfront so that customers can make that decision. And understand the case, do I want this and am I okay with this third-party company using my data so I can get this and that feature. I think those are the two pieces that, realistically, incentivize companies. We want to follow the regulations and we want to follow the tastes that the consumers have and we want to be transparent with them.

Larry Bernstein: Charles, question for you. I imagine that as this thing begins to evolve, there's going to be sort of like an opt in or opt out. And if those people who are very concerned about privacy, they can opt out of the coupons at Target. They can opt out of getting recommendations from Amazon on what books to read, or they can opt out or pay extra in their gaming applications not to have advertisements. But that said, it seems to me that there's a... My question really relates to generational preferences. It seems to me that my father's generation were very concerned about aspects of their privacy, almost religiously so. But my children seem almost completely indifferent or don't care about their loss of privacy. How do you think about this generational shift of behavior? Do you think that the older generations were more concerned by 1984 like concerns and the younger generation, that concerns have gone away as they haven't seen or been abused yet?

Charles Isbell:

I do completely agree with the observation that this is changing generationally, and certainly true with my kids. At first, I thought they just weren't aware of how much information they're giving out, but they absolutely are. They just don't care. And I think I care a lot less than my mother does about certain things. And that's just maybe the way it's going.

Michael Littman:

But I wanted just to say that I don't know that it's a clear, like not care. I think there's a cost benefit analysis of sorts that says, sure, you can opt out and you can kind of opt out of a lot of this stuff now. You can always use Google in that mode, that sort of dark spy mode where it can't track you, but nobody does that, partly because you lose out. There are things that you miss out on by not participating in this, I don't know, information economy. And I think, for many people it's like, well, I could opt out, but I can't afford to. And they're kind of forced it. It's not as free a choice as it feels like.

Charles Isbell:

I think that's right, Michael. but I'll also point out there's the other end of it is, what do I want the privacy for? There was a point in time that if people knew certain things about you then maybe you can never be a Supreme Court Justice. We're moving into a future where people can know those things about you when you can still be a Supreme Court Justice. If the consequences for large swaths of things that you used to think of as secrets are no longer secrets, then maybe it just matters a lot less. And maybe that's a good thing. Maybe it's a good thing that we don't care as much about the things that we used to care about before.

Larry Bernstein:

Why does our younger generation who face similar trade-offs seem to not care about their privacy as much, or have they bought into Google's belief that first do no evil? Do they trust the companies not to do evil or do you think it reflects a difference in their own utility function as it relates to privacy?

Charles Isbell:

I think it's their utility function. That's what I'm getting at least listening to my children. So that's anecdotes, not data. Although I do think there is some evidence that the younger generation thinks about these things very differently from older generations and the things that they think are important to keep to themselves or not. But again, you could ask yourself the question, if we lived in a world where there were no secrets, would that be a better world? And some people very much think that's true, and some people very much think it's not. My guess is the line, the percentages and the proportions have definitely changed over the last decade or so. They're probably going to continue to change. And we'll just see where it comes out at. I don't have a strong opinion about it because I just don't have a strong opinion on it.

Larry Bernstein:

Seyed, you're in the business of applying AI technology to solve business problems. And when you look out there in the field of opportunity, where do you think the richest and most fertile ground for AI applications is with regards to business applications?

Seyed Sajjadi:

Well, generally speaking, there's this thing I think in the advertising business that, if your revenue's not growing fast enough, just turn into the ad business. And I think that's exactly what a lot of AI companies have done. So, ad technology definitely has been an area of growth for a lot of ad companies. Security has always been one. Using cameras to understand what are the activities that are happening in front of

the camera, that there's someone breaking in. Is there a fight on the street, so on and so forth. The ones that VR are also passionate about and this is the very first customer that we had is the space industry using AI for space, very similar to how we're using it for cars and planes, getting autonomous system on the spacecraft. However, I'm not sure about the size of that business. And probably the number of customers are very limited to three or four space companies, like the NASA's and the SpaceX's of the world.

I think one that I'm particularly very excited about is seeing a lot of the AI algorithms and reinforcement learning breakthroughs that we've seen in StarCraft to be applied wider to games. So rather than having just easy, medium, hard, predefined AI rules for FIFA or NBA or any of these other games, having world champion AI to be integrated into the games that can help it visit players. I think those are a couple of the areas that I can think of.

Larry Bernstein:

In your talk, you mentioned cameras on the street. We had a discussion with Robert Vargas who's a sociologist at the University of Chicago, and he mentioned that minority communities in Chicago are demanding more cameras on the streets to reduce local crime. Do you think that if facial recognition improves that it will be revolutionary in crime reduction in terms of informing the police that there's trouble going on in the streets, making predictions of where to extend resources? We also had another discussion with Peter Moskos who told us that in Baltimore, when he was a police officer, they would be receiving millions of 911 calls a year, and the police had really no ordering of how to deal with problems. Could AI be helpful for police enforcement to send the cop cars into certain areas to minimize crime? How do you think of it in the terms of criminology and crime reduction?

Seyed Sajjadi:

Yeah. There's been a couple of experiments, especially on the East Coast side and New York is the one that I am remembering. For sure it can be used within the police departments. However, it is important to not arrest someone before they commit a crime.

Larry Bernstein:

Is that a Spielberg movie?

Seyed Sajjadi:

Yeah. So it's important, I think, if they have resource allocation problems, which is always the case, and they want to prioritize AI to understand, okay, if I'm receiving these many 911 calls, which one should I really prioritize and which one will have the higher chance, which neighborhood or zip code at this point in time, given the things that I've heard and received historically are going to have the highest chance of an actual crime? Then yes, for all means in terms of resource allocation, but not arresting people before they do things. The obvious ethical boundaries, I think, as long as they're being met, they can be quite helpful in criminology.

Another one is the body cam. There are a couple of companies that have started analyzing the body cameras that cops put on to help to go through millions of hours of content, and then see if there's been any incidents that they need to provide attention to.

Larry Bernstein:

Before we wrap up the panel on AI, I thought I would go around within this one panel and ask, what are you most optimistic about in this machine learning, AI world? And Michael, let me start with you, as you think about your field, what are you most optimistic about how it would improve human life?

Michael Littman:

Well, my hope is that at some point we're going to start a push for something more like universal programming literacy. The idea that we all should know how to tell our machines what to do. Just like it turns out to be a really valuable thing for citizens to be able to read and write. I think it's a really valuable thing for people to be able to tell machines what to do. At the moment, the companies are much better than we are. And I think that's part of the manipulations that we're seeing. If we can get to the point where more people have more control over the computers, I think then that's a good thing. That's more empowerment for individuals. And I feel like AI technology, as it exists, could help us get there. It could help bridge the gap between the complexity of current programming and regular people. All of us being able to tell machines what it is that we want them to do on our behalf.

Larry Bernstein:

Thank you. Charles, what are you optimistic about in AI?

Charles Isbell:, I think both in the sort of educational space, but also in the ability to stand up to larger organizational forces and institutional forces, I'm hoping that as we move along this path, we will be able to allow individuals, enhance individuals ability to make reasonable decisions and to understand what's going on. So even though there's an enormous amount of power for, I don't know, Amazon to know everything about what I'm buying, there's actually an enormous amount of data that's out there that also tells the individual about Amazon.

I'm not picking on Amazon. I'm just staring at a screen that has Amazon on it right now. And I think that, there's a good chance that if we can connect everybody to these tools, we can actually democratize this sort of power, at least when it comes to analysis and making somewhat better decisions. That's my hope anyway, and I think it can happen.

Larry Bernstein:

Seyed.

Seyed Sajjadi:

I think two things that have made me really excited and hopeful for the future of AI are the things that happened over the past two or three years. One was using AI to recognize a black hole in the middle of our own galaxy. That was amazing. And the second one was AlphaFold that was released a few months ago. Generally speaking, I'm most hopeful to see more applications of AI in other sciences, specifically life sciences and biology in a way that it can help with longevity and life expectancy and, can help with solving and curing a lot of diseases. I think the power of information processing that AI has and pattern recognition can really help us get an edge there and accelerate our learning life sciences as well as other sciences.