Deep Dive: Artificial Intelligence on the Rise

Alan Weil: Um, thank you for joining us. I don't think it's just the rain that brought you in under the tent. Uh, this is a hot topic. Before we start, I do just want to ask how many of you were in the big data session this morning? So I want to that. Fewer than I thought. That's good. So we're, this is not going to be a repeat, although there was a lot of discussion of artificial intelligence this morning, but I want to assure you we're going to go in a slightly different direction. What we're doing is, uh, called the deep dive only in Aspen is 80 minutes enough time to go deep, but it's more than the 50 that you get in the other sessions. And what that gives us is a little opportunity to go deeper. And for us what that means is we're really going to talk about use cases.

Alan Weil: We're going to talk about applications, where we're going, what the opportunities are, not in some theoretical sense and how much those opportunities are continuous or discontinuous with what we were doing before we had artificial intelligence. So we have a sense of grounding of how real this is, how it's being used, and we'll get into a variety of other issues as well. We're going to start with a, a bit of a Primmer, uh, because the terminology is complex and we want to make sure we're working from the same place. We'll then get into some use cases from our various panelists at a little bit past the halfway mark. We're going to start talking about AI in low resource environments, which raises a different set of issues. And one of the great things about Aspen ideas is that we bring to on the health side is we bring together domestic and global.

Alan Weil: We're going to do that in this session as well. Uh, Mike Know Haley is to my immediate left the senior vice president of strategy at Amgen. Lizzy Dorfman, senior technical program manager at Google Ai, a p on and don the CEO of the [inaudible] Institute for artificial intelligence in Catherine Moore. President of the intuitive foundation and more information on them is on the APP. Uh, but uh, enough to give you a sense of where we're starting. So who would you ask to give a primmer on AI? Someone from Google and um, indeed I did a session with Lizzie last year, so I know she's up to the task. So Lizzie, if you could just kick us off, get us all grounded so that as we take this deep dive, we're starting from about the same spot.

Lizzie Dorfman: Okay, thank you. Uh, so I'm coming from Google, but I'm going to spend the next couple of minutes trying to spare you all from googling words as we enter into this deep dive. If you were in this morning's big data session, there will be a little bit of overlap, but the terms that I'm going to be talking about, big data, artificial intelligence, machine learning and deep learning, whether they're part of your daily lives and work or not, these are starting to become water cooler terms and having everybody have some sense of what they do, what they don't mean, I think is quite beneficial. And I hope that you'll use this as you move forward in some of your discussions. So let's start with big data. This is probably the least specific of all of those terms. Uh, but generally speaking, this is referring to data sets that are so large, so complex and so diverse that they've really forced us to rethink methods around how we collect them, how we store and transport them.

Lizzie Dorfman: And in particular how we analyze them. Artificial intelligence refers to computers that are able to perform tasks that are typically understood to require human levels of cognition or intelligence. This is a applications of, this would be things like understanding and using human speech. We'll give you some Google examples are translate products which allow you to do realtime translation from one language to the next or realtime transcription. This is also the type of technology that's used to do things like suggesting the end of an email that you're writing. Another example would be the ability to understand and identify um, content within an image. The reason that you can pull out your phone and for Google image search, look up a word and instantly see all of the photos that have whatever you are searching for within them. This is based on Ai. Another example would be motion planning for robots might be surgical in nature.

Lizzie Dorfman: Uh, might also be a vehicle. Obviously Google has an organization called ways that's focused on, excuse me. Well, yeah, Waymo and mixed up ways is also Google, but that's traffic navigation. Um, one of the important distinctions to be very aware of when it comes to AI is between artificial general intelligence and narrow AI. Artificial general intelligence would refer to a computer that can perform any possible human cognitive task and uh, at present and for the extended unforeseeable future. This is the stuff of science fiction. We are not remotely close to achieving that level of AI, but what we do have is an increasing number of applications where computers are able to meet or in some cases even exceed human level performance for very, very specific tasks. Uh, we have seen several public examples of this in the gaming space. Lots of news when you had a computer beat world champions at chess and it go, we've also seen a number of examples of this in tasks related to image classification, including several medical imaging.

Lizzie Dorfman: And I'll be talking about several of those examples of it later today. Um, so the question that results is how machines become intelligence and there's several approaches that they can take. One of which is a brute force method, sometimes known as good old fashioned AI where you have preprogrammed all possible scenarios and explicit rules for the computer to follow. A example of this would be IBM's chess playing computer deep blue. However, there are a set of methods known as machine learning methods that don't preprogram explicit rules and instead they are showing the machinery specific training examples and allowing the computer to develop specific methods used to make classifications are predictions. Deep Learning, which we're starting to hear used much more is a specific type of machine learning that's exceptionally powerful when you're working with images. And the deep part of deep learning refers to the fact that there is a series of transformations that are happening to those images where you go from pattern detection of very, very simple features.

Lizzie Dorfman: For example, an edge to much more complex features that are frequently things we would consider human recognizable like a face or some text. So to wrap all of these up together, deep learning is an, is an increasingly common and very powerful type of machine learning. Machine learning is one of the ways that computers achieve intelligence or artificial intelligence. Uh, and, uh, big data is a, is still a sort of a fairly vague and broad term, but really it's primarily focused around having to rethink the methods that we use to analyze increasingly large, diverse and complex data sets.

Alan Weil: Great. So with that terminology, of course, applications for AI abound, we're here at, uh, ideas, health and so even within the domain of health, the variety of applications is quite profound. And what I'm going to ask each of our panelists here to do is to just spend a couple of minutes describing a core application, maybe a little bit of a mapping onto the, to the taxonomy you gave us so that we have a sense of the incredible breadth of what is actually already occurring in the domain of artificial intelligence. Catherine, you want to kick us off?

Catherine Mohr: Sure. Um, so I, I think of the panelists up here represent procedural medicine, uh, surgeries, biopsies, endoscopy's, the, the sort of active things that surgeons or practitioners due to their patients. And we're using AI and we'll be using AI in the future in a variety of different ways. But to give you a little bit of motivation around why we are focusing on a certain sorts of things that we are, I have cheat sheet on some of the data that I wanted to give you because it was apparent, especially from Peggy Clark's opening that surgery is considered a little bit of a luxury. And a lot of the things that we've been talking about has been in the, you know, sort of in the areas of what we would think of as more general medicine, but just a few numbers on surgery. So globally are little over 300 million surgeries per year, but there are 5 billion people on the planet that lack any kind of access to basic surgery.

Catherine Mohr: And so what that means is that there are about 17 million potentially preventable deaths that could have been cured if there had been surgery available. To put that in perspective, about 10 million people a year die of cancer, so 17 million from preventable surgical access if they'd had access to surgery versus 10 million a year that die of cancer versus 3 million a year for HIV, TB, malaria combined. So surgery and lack of access to surgery is a real public health issue. It's not a luxury, it's actually lack of access is something that's killing more people yearly than things that we think of traditionally as public health. So, but surgery is not a panacea. We do 300 million surgeries a year globally and 4 million people who received surgery die within 30 days of receiving that surgery. So it is not without risk and that means that we need to be doing a lot better.

Catherine Mohr: So keep in mind that 3 million HIV, malaria, and TB, 4 million are dying within 30 days of having surgery, which meant if we did it better, there's probably a lot of preventable deaths in that as well. If we were to expand the opportunity for surgery to the entire globe, we'd probably be looking at about 450 million surgeries being done a year and that would mean 6 million people dying within 30 days of receiving their surgery. So this is why too, just to give you that context, why I'm focusing on the somewhat peculiar areas, but without that context, you might wonder why are we looking at that? And so the areas that I'm looking at, our clinical decision making, how do we help surgeons make better clinical decisions during their operations? And then surgical training, if you want to train, you can build a hospital in a year. If you want to staff that hospital, it will take you 10 years to train the people to staff that hospital.

Catherine Mohr: If you're trying to meet the public health crisis of not having enough surgery, you're going to have to figure out how you train those people better and faster and how you build that kind of infrastructure. So what do we use AI and machine learning for? Currently in surgery, a lot of the image recognition, we're looking at a CT scans were looking at preoperative images. We are figuring out where a tumor is. You can navigate, you can help the surgeon navigate to where that tumor is inter operatively. There are uh, new kinds of fluorescence imaging that we're putting on there that you can then do image analysis and you can help people. It's essentially how do I give the patient that I'm working on a better surgery because I am a able to get better information out of this data. Where we will go in the future is as we're looking at video based data, a lot of these procedures, if they're endoscopic or robotic or laparoscopic or your, your, your colonoscopy, there's a video component to the data and you can do a lot of parsing of that into being able to recognize the different sections of that procedure.

Catherine Mohr: And as you start to look at in the future, uh, what kinds of outcomes came out of that. You can start recognizing correlations between actions that happened during surgery and better outcomes associated with that. From this, we start to be able to build AI coaches some, a sort of a long with you looking over your shoulder, helping you do a better surgery, but also being part of accelerating the learning curve on these millions of surgeons and, uh, millions of new nurses that we're going to need in this kind of procedural healthcare world. So I said I was going to focus on things that were very different from [inaudible] was talking about because uh, in, so in the world of procedural medicine, a lot of the things that we are thinking about is AI and machine learning in decision making, in being able to avoid complications interoperatively and in the future, that's what we're doing now in the future, being able to bring in these AI coaches. Great. Thank you.

Padmanabhan A: 100 on your work. Of course it goes beyond health, but if you could focus on some of the AI applications in health that you're working. Thanks Helen. Yeah, I ran, but money is true for AI. Uh, AI for social good. It's a not for profit working in a low, low resource communities. Primarily we are looking at health as one of the areas they held a for the global south. So definitely 3 billion people who don't have adequate medical care. That is not enough doctors, enough nurses. Uh, there's a ratios out of darted of 30,000 patients to a doctor and such. So clearly there are a number of places where we want to augment the capacity of a frontline health workers and others using AI type of technology. Then that's the opportunity in our inner, we've been around for one year. One of the things we learned is in choosing use cases is that you have to talk to three different parties, the governments who ultimately on the caregiving to these large populations.

Padmanabhan A: For example, India via base, we are talking about 800 million people who are part of the public healthcare system. Uh, NGOs that are filled, NGOs who know how to provide delivered solutions, executive whatever technology or develop, who know where to go and deliver the solutions she ever talked to them. And third sometime that our tech apps that are part of the system and the makers of those apps, you have to talk to them. And what about you build needs to be in consistency with these organizations because they're the ones who can enable getting the solution. So we in our first year journey, uh, you know, came up with a couple of areas where we want to focus. There are many, but just to pick the do maternal and childhood, maternal and neonatal child health and TB. So let's talk a little bit about use cases on each of them in MNCH.

Padmanabhan A: For example, one problem we zoned in on, again, conversations with all the partners is a low book baby birth weight detection. It turns out there are something like 20 million low baby, uh, but low birth weight babies born and India, the number is of the order of four to 5 million. But 30% of those babies say for example, in countries like India or not weighted birth, there is an existing protocol to do so using frontline health workers carrying spring balances and weighing at clinics or at home at home, but often that are supply chain problems. The balances don't work or the cumbersome process of doing it or even in fact you cannot touch the baby because of the taboos and sites. So we realize that it's put, it may be possible to do that by simply taking a cell phone video or a few images and using computer vision techniques, namely AI to create a three dimensional virtual shape of the baby.

Padmanabhan A: And from that shape take a lot of body measurements and we know that the volume of the baby and the waiter closely related density is almost constant. So we could get the weight and this is something as we start doing, we realized, cause there's a whole data collection problem, which I can talk about later, but it seems like an exciting opportunity. The second one, as I mentioned TB, the entire cascade of TB from finding the cases to diagnosis and screening and uh, you know, treatment ideas all faced with challenges number, sort of the order of a say in India for example, that is believed to be about 4 million annual TB, new cases, new TB cases, but only one and a half million known cases. Those who go to public doctors and then notified. Can we use that information and correlate with other demographic data or population data or even personal data to triangulate and see where there might be more cases in small regions, wards and blogs and sites.

Padmanabhan A: So we can focus the caseworkers effort on finding new cases. Again, you know, there are other aspects of this as well where AI is amenable, but in all of these, one thing that you learn is that you have to work with the partners that I mentioned earlier and they have to be cocreators and they have to be willing to implement. And the governments are we willing to scale. So great. Um, Mike, I'm gonna, I'm gonna jump over Lizzy because she kicked us off and then we'll give her a shot to win. Micah when she gives us.

Mike Nohaile: Um, I think the fundamental challenge we face as a biotech company is that even if you advance really impressive innovation that can make a difference in a terrible disease, uh, it's not always used properly. All the patients aren't found. If the patients are identified, they're not always put on the therapy that's optimal for them. If they're put it on the therapy that's optimal for them, they're not always able to stick with it or take it and derive the benefit. So I'm going to focus on one area osteoporosis, uh, that we work in. And it's a fundamental issue at the beginning of that funnel, uh, only about one in four women who have osteoporosis are being adequately treated and many of them don't even know they have the disease. So I'm going to give two use cases, one that is directly based on the techniques that you spoke of, which is an image, a technique, and then one that's based on something else that I'm excited about, but I would still car artificial intelligence.

Mike Nohaile: So the first one, the image one, uh, the challenges that many, uh, women start to get what are called compression fractions in their spine. So they have these small fractures in the spine. Sometimes they can be quite painful, but sometimes they're painless and so they don't know. But if, boy, if you have a compression fracture, you're heading for hip fracture, and of course you get hip fracture, you may never get up again. In fact, you're likely to never get up again if you're old enough. So, um, what we've done, uh, in this case working with an outside company is, uh, develop an imaging technique along the lines that was discussed. And what it does is it crawls through Emr in a hospital and it just looks at all the incidental images. Maybe you came in, uh, and they looked at your lungs Namony suspicion, they just wanted to check or whatever.

Mike Nohaile: But in the background is your spine. And so the radiologist doesn't look because they're human and they're busy and they have many, many other things to do. However, the machine doesn't get tired. So it just continues to crawl through and look at all those images. And then all it does is fly the doctor if it finds a potential compression fracture and say, hey, you should take a look here. Maybe this is somebody that should be worked up for potential osteoporosis. And we actually have that working. It's actually in a hospital system now in a, in the middle of a test and it looks pretty good. It looks like it pulls out a significant number of people in that situation. Again, it's that making the final decision, but it's highlighting people for the doctor to look at. And I think that's going be an important way in the future to think about identifying patients along the same line even earlier, uh, or even with less information.

Mike Nohaile: Um, the second technique, we're excited about a predictive analytic techniques where you take streams of data and new predict how the, how that would go into the future. It's not an image technique. Um, but what in this case, what we're able to do is use a set of techniques that are actually related to what we use in natural language processing to predict the next word in a sentence and take an EMR record and predict what the next word or that sentence was. And sometimes the next word is bone break, right? So then you can take, create a predictive score. So the, the standard today is something called frax, which creates a predictive score that says in the next 10 year, this woman is likely to fracture to a certain level. Well that's great, but 10 years doesn't really motivate people. 10 years, doesn't motivate insurers.

Mike Nohaile: 10 years doesn't motivate people. So using these techniques, we were able to get a version that we think might be able to give you a two year score, right? And two years, much more motivated even to insurers who can see the payoff for preventing that fracture in two years. So that's still more invalidation. Uh, the first data set looks good, but always let you have to do is when you do it, that you'd have to put it on naive data sets that have never seen it to make sure you haven't fooled yourself in this possible we have. But again, I think that's the kind of technology that hopefully will allow doctors and other caregivers to be armed with much better information and to focus their attention on the right patients because, uh, if they have to spend all their time just finding those right patients, uh, that's the looting, the precious time and skill that they have.

Lizzie Dorfman: And Lizzie, um, so I'm gonna give, I'm gonna give two examples. My, my background and training is in genetics in public health. And the examples I'm going to give our relate to neither of those fields. Um, but they, I have had some involvement in one of the projects and, and I would like to sort of plant in all of your minds a sense of the breadth of applications to health and health care. Uh, so the first example is in radiology focused on lung cancer. Catherine, you mentioned 10 million cancer deaths per year. About 1.8 million of those cancer deaths are from lung cancer. In the U S it's about 180,000 deaths last year. It's the top cancer related causes of death for men and for women. If you screen people using low dose CT, you can lower the cancer death rate for lung cancer by about 20 to 35%.

Lizzie Dorfman: In the US preventative services task force recommends that you screed and high risk individuals. Mostly smokers. The challenges is that we have serious shortages of radiologists and many places in the u s the ratio of radiologists to population is about one to 10,000 in some parts of India. This looks closer to one to 100,000. I was reading an article last week. Uh, Nigeria has 190 million people and fewer than 60 radiologists. That puts the ratio of over 2.5 million to one clearly cannot scale to be able to handle the lumbo the that we don't have the amount of expertise that would be necessary. So, uh, some recent work published just in the last couple of weeks by the radiology team within, within Google AI and Google health, uh, in nature medicine showed that we have been able to train a deep learning model. This is again the image classification method that we've been talking about.

Lizzie Dorfman: Um, that actually is able to meet it in some cases, exceed radiologists level accuracy for identifying potentially cancerous nodules in low dose Ct. If at cs one Lodo CT doesn't have any historical data, it actually was better than radiologists. It had 11% fewer false positives. And 5% fewer false negatives. And if it did have longitudinal data than its was on par with radiologists. So this was just published. This has not yet been turned into a commercial product or available algorithm. Um, but I think that's showing us how we can map these capabilities to, to areas of tremendous need. I'll very briefly talk about a second and completely different project, which is not at all about clinical decision support or decision making. It's about helping make doctors' lives easier by trying to not require them to spend as much time sitting in front of computers and typing in information.

Lizzie Dorfman: So this work is focused around a, we call it scribing or if this is known as scribing. Um, but is it possible to capture audio data of what's being communicated between a doctor and a patient in a medical exam in to auto generate at least the starting point, if not the complete content of a note that they would need to enter into a medical record to allow them to spend more time during the interaction, actually making eye contact with the patient and to not have to spend after hours time doing mind numbing data entry. So this is work that we have a in a research stage and are working with several healthcare partners to develop.

Alan Weil: Okay. So if there's one thing that's clear, the applications are many and they're quite diverse. And so the challenge is to have a conversation about some themes in an environment where the applications are so variable. So I'm going to try a few and we'll see how they go. The first one gets at the question of how transformative this is if we keep a focus on sort of the diagnostic and treatment part of, of healthcare. Uh, the historical model is you have a very highly trained clinician who learns a lot and is, is tasked with keeping up with the literature and they apply their knowledge to the individual patient and we expect them to adhere to a certain standard. As we move to these more automated methods, uh, they can augment the information that the clinician has and that's just a better library, a better, uh, a recent set of resources.

Alan Weil: But at some point, it sounds like we start to flip to where the automated, uh, information generation may surpass the skill and the capacity of the clinician, which fundamentally alters the dynamics here. So my question to each of you is, where are we in that continuum and what are the issues that arise as we begin to move in that direction, uh, from, uh, from everything, from professionalism to a liability and accountability to resource use. Um, that just seems like a, uh, a major pivot. So where are we in that continuum and what are the implications? Who wants to start? I can start Chris. We build

Catherine Mohr: surgical robots and a surgeon sits at a console and performs the surgery with the robot as an intermediary. And when people hear about surgical robots, the first question they always ask is, well, when's the robot going to do the surgery all by itself? When are you going to automate that? And um, the, the, the short is short answers is going to be really long time. Um, we have an existence proof that if you take a natural intelligence, if you take a human being and you train them and you sit them at this console, they can do all of the motions of surgery. A complete surgery can be done. So if you could con could train a natural intelligence to run that robot, you'd be able to automate surgery. But as Lizzie said, we're so far away from that that we've got very, uh, narrow areas of intelligence that we're thinking about.

Catherine Mohr: The other aspects of this come into these fundamental questions that we're going to have to answer as a society, which is when an algorithm makes a mistake, who do we blame and whose fault is it? And how, how much are we going to tolerate there being mistakes made by automation, even if they are fewer mistakes that are being made by the humans, we're going to have to sort this out in cars first. So we're not even going to touch it and health care until, until we've already figured out, you know, who's at fault when there's been a self driving car accident. But, but we will learn a lot about both our tolerance for risk and um, our, our sort of, our tolerance for errors coming from algorithms, uh, as a society and who we do choose to blame when something goes wrong in that kind of a situation. And so, uh, we will continue. I've, you know, for very long time having our clinicians in the loop on this and the, but the, the support structures that they will be getting from their AI are going to be increasing in sophistication to the point where at some point they are teaching them, not just helping in a decision, but I'm actually performing it.

Catherine Mohr: I'm, I'm not sure, you know, it's not just a technology problem.

Alan Weil: Yeah. So the complexity of surgery, I even for a non clinician like me, that's pretty easy to imagine. But in some of these diagnostics it seems like we could get there earlier. Mike, I mean you mentioned the, the, the computer going through all the images. Um, but they then hand that off to a clinician.

Mike Nohaile: Yeah. I think that's going to be the model. Um, little less than 10 years ago, I ran a business that had 50 pathologists. We were doing diagnoses of blood cancers and they were all depressed because they thought they were going to reprice by computers very soon. Um, but I don't think that's how it's working out. I'm not saying there's not some point in the future, but I think it's the center model that people have talked a lot about, which is, you know, half human, half computer. So the computer will pick up more routine tasks. There may be some tasks that he can actually do better than the human deal, very specialized where it's doing an imaging thing and they can pick up more information. It certainly can piece together more disparate sources and present that in a way that a human just can't comprehend that much at once.

Mike Nohaile: But I still see us many years away from taking the human out of that loop. And that's not even what we're trying to do. We're trying to put tools in people's hands where they can do their job better. And again, I would also come back to this, a focus of care. There was a study a few years ago sign an AI thing, but it's a similar concept where, um, this was, I think that moral Sloan, they took a 700 or so late stage parents or patients and all they did was once a week ask them how they were doing on a clunky web browser. And then they use a relatively simple set of algorithms, not even Ai. And they fed that back to the care people and it focus their attention on the sickest patients and they got a four month life extension benefit on that. So I think we can't underestimate that getting humans to do the things humans can do is unbelievably valuable. And I think we have a long way to go before the machines are replacing the humans on mass or the end, the end at any of these processes.

Lizzie Dorfman: I think one topic that's quite important related to whether how and when will use AI systems is interpretability and explainability. So the understanding how an AI algorithm made a particular classification or image, and this is important for a number of reasons. We need to understand and have confidence that these things are consistently behaving the way that we hope that they will. We also need to be able to detect and prevent abuse and fraud. And third, optimally we would love to be able to actually learn from the situations where we achieved superhuman capability in the computer was actually able to do something that humans were not. Um, do you sometimes hear the term black box? Uh, what I actually have been very excited about is the fact that there's an increasing number of methods that allow us to understand how and why particular classifications or decisions are being made.

Lizzie Dorfman: Um, one of the ones that we use frequently at Google, it's called attention. What was the model attending to? What specific subset of the input data was most influential for the final decision that was made in the case of an image? So in the case of any to say we're back to talking about the lung CT example, what pixels in that image that this model believe influenced its, its, its prediction of whether this individual has a nodule or has a cancerous nodule. Um, so what it does is this allows us to generate these heat map overlays, specifically lighting up the pixels in the image that led to that classification. Um, this is also how we can sort of autogenerate a summary of what's in it, what's in a general picture. We could say, oh you thought that was a cow and that was a fence or you completely got it wrong.

Lizzie Dorfman: By generating these overlays in the case of text analysis, it's going to be able to surface, which were the words in which were the phrases that were used primarily to be generating these predictions. As an example of this, um, our team has done a lot of work looking at what potential signals are available from picture of the back of your eye in particular. Are there things that humans can't detect? And the answer is yes. We actually found that with fairly high accuracy, we can assess someone's cardiovascular risk from a picture of the back of their eye when we applied these, these attention model. So that problem to say, okay, well what specifically does the model think is informative about this question? The heat map overlay is of the blood vessels in the eye, which when you think about it makes biological sense. And this is what helps us to build more trust in these models and potentially gives us avenues to pursue, to understand why it was that they were able to determine something with accuracy that a human wasn't,

Padmanabhan A: I don't know. Let me again speak about Lotus has communities to set the context. Uh, in a country like India that I bought 800 million people who lie on the public health system. There are about 25, 27,000 what are called primary health centers that have a medically qualified doctor who has a basic medical degree, about 150,000 subset does one for every four villages roughly in India. And they are run by nurse midwives. They are not doctors, but that number is still not enough to, you know, handle the large number of patients. And many people simply don't go because of the time and effort it takes to get to the nearest no sub. The primary health center we'll do, they operate on the basis of using a frontline health workers. These are in India and sometimes called Asha workers and these are community women, typically not high school, graduated who are now trained by these nurse midwives and doctors and they are operating under a certain protocol and procedure to visit, uh, you know, families, community families regularly.

Padmanabhan A: And as you might expect, a lot of that effort really focuses on maternal and neonatal child health because that's the primary place where they can actually make a difference. Now we are talking about how AI can augment. Now clearly it is a case that we just don't have enough doctors. We have 30,000 patients per doctors and the doctors, uh, kid is not available to every patient knows midwives, our outreach workers do some amount of triaging. The more we can increase their capacity to better do that theology, the more they can measure things, the more they can actually have position making systems behind that will actually learn from doctors and help the health workers are the right questions, the better they'll be able to try out. So here's the case. But I don't think it's far fetched to think about the capacity of AI based technologies immediately increasing the way in which care can be given even while you know, the system is to find more doctors.

Alan Weil: So I want to, I want it.

Catherine Mohr: Yeah. Cause I could just an observation in building on some of the things that you guys were saying, one of the important things to understand is, uh, if you want an AI to be able to make good decisions based on a data set that you've given it, you actually need to be sure that you've given it a complete data set. Because you know, if, if you're saying, where is the attention being paid in an image for radiologists, we don't encounter a CT scans or images of the eye in our normal daily life. And so we're not bringing a lot of contextual information. We're not bringing a lot of information outside of the image itself. So you can train image recognition, AI eyes on just the image and anyone interpreting that image is only using the image but making decisions on midwifery and what kinds of things you're going to be doing in a village. You don't know all of the information that that person is bringing to bear in the decisions that they're making. And so being able to predict what is a complete data set that's necessary to train the Ai, you may be missing out the critical component that is actually determining whether they're making a correct decision or not. And so yeah, these are, these are the differences in do you have, yeah. How amenable is the problem you're solving to actually having a complete dataset?

Alan Weil: Yeah. So it's funny as I was about to ask a question in and you took a summer different angle on it. So we think so much about big data and AI is this innovation and you can't go to a conference without everyone. You know, if, if innovations in the title, then everyone has to be in that session. That's why you all are here. Um, but what I'm hearing is most of these applications are about improving capacity, inaccuracy, which is phenomenally important, particularly in low resource. But even in high resource environments, we're trying to be more efficient. But it makes me wonder, something that I hope you'll convince me that I'm wrong. And what it makes me wonder is, is AI actually a barrier to innovation? And the reason I come to this question is, Lizzy, as you said at the outset, we don't have general intelligence Ai. So what we're doing is we're throwing a technology at problems that we've already identified with approaches and solutions that we've already identified and we're improving capacity and accuracy. But innovation is about taking a very different approach, seeing a problem and thinking about it differently or, or coming up with a solution that no one else ever did. And Ai isn't going to do that. So I wonder, just again, talk me down, you know, convinced me that that this is not actually or reifying status quo healthcare, maybe making it more

Mike Nohaile: efficient, but swinging our attention to something that's less innovative than a dentist pocket. Please make me feel better, which is, so I don't know how to make you feel better about a practice of medicine. I'd have to think about that. But the first thing we do is try to discover new medicines and our scientists spend lots of time trolling through literature for ideas. Re Yeah, it's actually quite helpful. It's surfacing stuff that then they can put together into innovative ideas that no one's ever thought about. So you know, that's a natural language processing application. But again, the idea is I want the scientist spending time interpreting experimental data. I don't let them fiddling around with getting it. I want them thinking about what are the insights from the literature and the Ai. Sometimes fools literature they would never read. Right. And that's where you get the innovation cause they're, you know, they can only read so many journals and follow so many journals and then they do some Google searches or Google likes searches and that's helpful.

Mike Nohaile: But they don't hit the right key words or the right stuff in the AI is constantly puling out. Surprising. And most of them are not relevant. But a few of them are interesting and then a researcher says, Huh, I wonder why those things are related or I wonder why in this field it's not related to mind. Son was talking about x and then they get a really great idea. So it's still dependent on the human. I'm not saying they're doing it for them, but I do think, you know, when I see my researchers slide week for awhile we had, in fact there's someone in the audience who fix this. They would go to one machine and write down the date and then you go to another machine and type it in as just a wasted time. Right? You want that all to happen so that that's not what they're doing.

Padmanabhan A: So actually, um, Lizzy made a point about what does good old fashioned AI. Yeah. In Modern Ai, which is the machine learning component. See Ai is not a technology that's baked in one place and applied everywhere. It is one kind of method. But the whole point of machine learning is that you give it the data that is relevant for your context and your situation and you train it over there. I mean you need to know, you need to know from examples what the data, how to interpret the data. It's not going to just magically figure it out itself. You need to teach it. But if you can teach it, you can actually train it and customize it to different situations. So in a sense that AI I think is open technology that for example, as Catherine was saying, that even in interpreting the diagnostic methods that may be done, we can apply the context.

Padmanabhan A: And so how do I actually read it? Given this is the population, this is maybe the lifestyle, this is the dietary habits, this is a prevalence of disease and and bias the outcome according to that learning. Right? So in a sense it allows you to do innovation in place rather than being a barrier. I think it's potentially one that's capable of expanding the scope of innovation overall. I mean the one device by the, I should mention one technology that we find is universally applicable and has become fundamentally necessary for all this actually is a cell phone. This is one technology that's made in one place, applies everywhere, but that's the delivery mechanism for AI that makes it all possible. So

Catherine Mohr: I would, I would say at a fundamental level, every single one of us up here are trying to solve real human problems, clinical problems. Uh, we are trying to solve issues that are out there. Ai As a tool and we're applying AI as one of the possible solutions to the problems that we're ever going. Oh, it's a really cool tool and it can help us do a whole bunch of different things, but it's not going to drive the innovation. What's gonna Drive? The innovation is us looking at the real human problems we're trying to solve and then figuring out what in our set of tools in our tool box is going to help us solve that the best way we can. So thinking that AI is going to drive the innovation is the tail wagging the dog. It's the, it's really the fundamental problems we're trying to solve that are gonna drive the innovation as a tool.

Alan Weil: And I'll just say, I was never thinking AI would drive it. The question is, does it actually potentially a impeded, but I'm hearing I'm, you know, I'm a lot better than I was 10 minutes.

Lizzie Dorfman: Yeah. I'll give you, we had talked to down a little while ago. Okay. I'll say a couple of things. One is, uh, if you're interested in innovation, there is a panel on innovation in this tent that follows this session. Um, I, uh, I think a lot about the question of like, what is the possible version of the future that I would love to see and how does that inform the incremental steps that we're going to need to take to get there? And I would agree that a lot of how we're using AI in health and healthcare today is incremental. It's taking specific well-characterized tasks that humans do and seeing if we can help extend capacity or improve accuracy. And I think that that serves two very important functions. One is we're not as accurate as we can and want to be and we don't have capacity to support what we need to do.

Lizzie Dorfman: So it is objective Lee valuable, but arguably even more important as I think that it's a building confidence that we can, we can use this technology in ways that are rigorous and robust and helpful. That lets us be braver and bolder in terms of how we think bigger. But this isn't saying this isn't really directly related to the question of visit is it is an opportunity cost in terms of mind share for other ways to approach problems. Um, I'm not, I'm not personally worried about that, but, but that's how I see the, the specific sort of opportunities to improve, extend and automate current clinical tasks.

Alan Weil: And certainly in the driving analogy was given early or, or application that was given earlier. There were lots of, of, uh, supports for active driving that are automated, that have given us some confidence as we move towards autonomous so that that path makes sense. No, I think you each gave me some very concrete reasons to see how AI liberates human potential to focus on innovation. Um, and, and even though it's not the source of the innovation, I thought that was very persuasive. So what I want to do is I at we've had a number of to low resource settings. I do want to turn to on and on to talk a little bit more deeply, particularly because something you said this morning about just the data sources, uh, being a huge barrier. A Catherine, you have also done some work in this area. You alluded to it with uh, the supply of, of uh, or the existence of surgical capacity. And then we'll probably, depending on how long these take with maybe a follow up question, we'll bring the rest of you all into the conversation. But let's, uh, let's talk about low resource center.

Padmanabhan A: Yeah. That are basically three parts to, in our developing an AI on any other solution for that matter in a low resource country. One is of course, defining finding the right use cases. As I said earlier, that you do in consultation with the people in the domain who know what the priorities are. The choices are so many and you don't know which is more gentle. You have to pick. Now having done that, if you're thinking about, you know like all of us AI researchers do, I'll go collect the data set, I'll annotate them. I'll have, the way I worked by the way the machine learning is that you take a lot of input information and then you take it on, take a subset of it where you might have known outcomes and use that as a training example, so it's learning the mapping from the input data to the output and then train it and then applied to new new data.

Padmanabhan A: The problem is that there's is simply not enough input data in in a in public sector, public health domains that data is sparse, meaning that some people are covered, some people are not very little longitudinal data that data or incomplete meaning for the same person. You may not have a lot of information. There are often incorrect, noisy and they are often false. As I said in the baby weight, for example, it's not uncommon for people to write the weight of two and a half kilograms as the baby weight because that's the threshold below which the baby is considered critical. Why? Because they eyeball it and say, okay, two and a half it looks okay, but you'll see a series of two and a half kilograms next to each other. You know, something's wrong. And if this is the data on which we are going to actually train or for that matter, any of that thing about diagnosis and so on, we have a big challenge.

Padmanabhan A: So finding good data requires a certain amount of investment of effort and creating perhaps an ecosystem. But equally importantly, you have to start instrumenting other caregiving methods to, you know, start producing data. It's, it's a, it's remains a huge challenge. And the last part is that even if you do all of that, if you want to actually do some field trials, there's no way you can actually take the solution to the end user. Now, of course, they induce at being a poet. Uh, you know, mother in a, in a village is not going to be able to use your solution. You will go through the frontline health workers, but you cannot get to them. You need to work through the organization and the agency that's working with the particular government, usually a state level government, which may or may not be ready to take this and then find the right match and in order to be able to deliver the solution. So fundamentally, every one of these requires special effort. Uh, and you kind of have to start with the places where you can make progress so that hopefully those examples will help others come around.

Alan Weil: Um, so

Catherine Mohr: everything you hear from me is in the, through the lens of surgery, but it's the, one of the things that you can tell when watching a surgeon in the u s is often you can tell where they trained, whether it was the east coast or the west coast. They've got techniques, they've got a ways in which they use particular tools. And so if you train just on west coast surgeons, you're going at the east coast, surgeons are going to see the kind of, uh, techniques that are being used and they're going to look foreign to them. You go to another country where the same surgical products, you can't get a stapler or you can't get a particular instrument. And the techniques that evolved there are very different from what you're going to see in the datasets that you get out of the United States. If you go into an under resourced area, things that you assume you can use like a stapler, they may suit you or something instead.

Catherine Mohr: So you see regional accents, but you see globally dialects in the way that people practice certain procedural medicine. If your data sets don't include all of the people that you're trying to, uh, coach or analyze, you've got this, the, the tyranny of the Dataset that you've used. You or you average the entire dataset that you're using. And you say the best procedure is the average procedure of this entire Dataset. And that's not true. There's going to be lit, you know, sort of regional clusterings of your data around best practices that correlate to that region of the world that those pieces of data are coming from. And so when we look at being able to create tools that are going to be used globally, our datasets have to be global and we have to really be thinking about weight the way medicine and surgery is being practiced there. And so what I might think of as best practices in a certain environment is only best practices in the environment that that was, uh, you know, that those were developed. And so I think this is a really important part of the way we're thinking about using these tools globally.

Alan Weil: So let me just follow this thread a little bit further. Um, it's, it's relatively easy to contrast an application at Memorial Sloan Kettering and, uh, measuring, uh, uh, birth weight in rural India, but there's a whole lot in the middle and all of you are operating in environments that, that span, that, uh, that range. So reflect if you will, on not just the question of whose data are used, but then who owns the algorithm. Mike, you described, you know, screening for images. Well that's presumably the investment that your company made. Uh, how that's proprietary. Presumably just asking these questions, how do we think about spreading that knowledge? How do we think about equitable access, the transparency that creates trust? Uh, the flip side is if it's a trade, it's a trade secret. If it's to be valuable to your company, there's only so much you can share. So I just wonder if we could sort of use this contrast between high and low resource as a framing question around some issues of ownership and equity. I know, uh, Lizzie, you've thought about this a lot. Maybe I'll turn first to you.

Lizzie Dorfman: Yeah. A couple of things. This, this came up briefly this morning, but from my perspective, one of the ways that that Google can actually help them most is to be developing capabilities that lower the technical barrier to entry for people to be developing their own models that are customized to their own problems and their own local contexts. And this is tools like we have one called auto ml or you don't have to write a single line of code. It's, it's facilitating the process of bringing AI technologies to your problem, but it is not fundamentally one. We're Google takes possession of your data in Google teams, build your data and it introduces these questions of who is owning the problem. It also allows the type of customization that is typically quite critical for problems in low resource settings where you really need to bring that specialized understanding of that problem to bear.

Alan Weil: Mike, how do you all think about these issues?

Mike Nohaile: Yeah, so I think this gets to a really fundamental problem with all of these technologies which are incentives and how systems are designed because while there are data gaps and then there's data engineering problems, we can always have better algorithms. I still think that the single fundamental problem is this. So if you take my example of the imaging algorithm, why does that work? Well, we have a bone medicine and we are excited about it and we want more people to be diagnosed now because it will directly go to our medicine. But we're trying to expand that funnel. We have a little company so we didn't fight with them over the Ip much because we're happy if everybody uses it. Right. There's other situations which are competitive where that would not be so much the case. So I think in every one of these cases you actually have to work back and ask, can I align the incentives properly?

Mike Nohaile: And sometimes it's, it's a public private partnership. So in the UK they have the UK bio bank, they've collected 500,000 high quality samples. Right. And they do have private folks come in and work on it, right. There's a company that has done a little bit of work sequencing, I think 50,000 of the 500,000 they have, but they will then make that data public after a relatively short period. I can't remember the exact number. So I just, I think this is the single hardest part of getting these technologies from really cool ideas that work in there oh. Cases to actually changing dramatically how the healthcare system works see in, uh, in regular tech world, right. One of the things that has enabled a lot of growth is having open innovation. Ecosystem open has been

Padmanabhan A: a very important driver. In fact, Google is a great champion of this in Loretto situations. One of the big problems that there's not enough capacity to pay. So a lot of the innovation has to happen, uh, in an open fashion and in a shared fashion because no single innovator has the incentive to go deep and invest a Lord. People on the innovation, but going to market is equally difficult. So we have to create ecosystems that partners and come and play and also go through the process of, you know, creating the ecosystem of Bellotti. Uh, fortunately or unfortunately, currently the situation is that this has to be done largely through government and philanthropic initiatives without a lot of IP ownership because they're just not enough players kept. We'll have the strength to do that.

Catherine Mohr: Yeah. For you moved from company to foundation. So it's exactly that. Yeah. So tell us a little, I mean, this seems really important. Yeah. Um, so, uh, I've been with intuitive surgical for now I guess 12 years and uh, ran research and a global strategy. And the conclusion that I came to was that the data needed to be open and available for everybody to use. So this past year we put together the intuitive foundation and one of the big projects that we're working on is building a global database of procedural videos, um, with associated metadata that will be open for researchers. So we're still working out any foundations out there who want to come in and build some coalition building right now. But, um, being able to allow people to build products on top of that, but essentially keeping the data in an open comments, kind of like the Wikipedia, Wiki media sort of, um, structure.

Catherine Mohr: So that AI researchers who can't get access to the clinical data and clinicians who don't really have connections to AI and ml researchers that, that this is a place that they can come together. And so we've come to that exact same conclusion that the data need to be a commons and that all of the other things can get built up on top of that. And that there, there isn't there, there's a strong role for philanthropy in that because if you try to extract value out of that, out of the data you get in the way of the data aggregation and what you really want is to get that there and then allow the value to get built on top of that. So I want to put one, yeah. One very quick please. Yeah. And then I'm going to ask a follow up.

Lizzie Dorfman: So just draw a distinction between open for research and valuable for AI. Yeah. Um, before we started working with, with low dose CT, we were working with chest x rays. We were interested in finding nodules. Google does not have a large corpus of this type of data. And we had to rely on partnerships and public data sets. And one of the largest, uh, to date public data sets was about to be released. We knew it was coming, we were excited about working it. We ingested into our system and we started training models and we thought, wow, our models looks great. And as we normally would do, he pulled up a couple of examples of cases that we got wrong cases where we had predicted that there was a nodule on, there wasn't or that we predicted there wasn't in, there was, we were sitting in a conference room, we were simulating a reading room.

Lizzie Dorfman: We pulled down all the shades in the blinds. We had a couple of chest x rays up on large monitors and the room went really quiet because we realized that some of you are probably guessing that there were pen markings on these chest x rays that were in circling or pointing arrows at the findings of interest, which if you're trying to develop automated models to be able to detect these things is quite problematic so they could beer, you know, lesson blurred for QC. But um, there are, there are, there it is critical that we are getting larger, uh, larger sets of data available for use. But it also there does take some intention to have these data sets be valuable for AI applications.

Alan Weil: Okay. So I, I'm thrilled to hear the discussion of data access and we had a really interesting paper in the journal about all of the efforts around that or a variety of efforts around the world around data sharing. But I want to push to the next step because it's relatively easy. I'm not saying it's easy. I know it's not easy, but it's relatively easy to talk about data sharing because each individual actor benefits from the data pool being larger. And in most instances, no individual actor has sufficient numbers to, to do everything they would want to do. But I want to go further down the chain to algorithms, to, uh, applications where at the end of the day, if we're going to have equitable or, or some form of, of dissemination diffusion beyond, uh, an individual, a private enterprise, we have to be willing and able to share algorithms and, and workflows and all of the other, you know, surgical procedures, not just the data. So imagine a world, I know it takes some imagination where we've solved the data access confidentiality, but pooling, sharing PR problem, how do we solve the next order of problems so that we can all benefit from these, uh, developments. So you mean like an algorithm marketplace? Yeah, I mean what's what I dunno, uh, use that AI creativity to solve this problem.

Catherine Mohr: Well, so, so we, we have a few um, structures that have organically grown up around these kinds of things where, uh, researchers love to publish their algorithms and to share. And you get, you know what I mean? Google's a good example of the opensource tools, you know, so, so

Lizzie Dorfman: why people like

Catherine Mohr: to contribute into this comments is because they're saying if I do something and everyone else can see it and can benefit from it and they can build on it, it's an enormous sense of community. And the academic community is driven by this betterment. And, and so we spontaneously c groups growing up in sharing. And I mean the number of hours that people put into, I use the Wikipedia example just editing cause they're doing this because they want to improve the world at large. And I see that same kind of human nature playing out if you can, if you can provide these datasets, people will be sharing what they do with it. There'll be quite a few people who make their Ip and make a a an entrepreneurial company around it and, and sell things based on what they're doing. But, but the people that are open are actually often moving a lot faster than anyone that's branched something off. So yeah,

Mike Nohaile: I'm less concerned that this is a problem. Yeah. Data, because I, without naming names and talking Lizzie's book, there is a, there is a large company sitting to my left that has been very open that has now they've used their cloud in some ways to help, right? You get into their cloud and you can get access to an unbelievable workbench that they've assembled for you to use. The other hand, there's another large company, many of you have their products in their pocket that for a long time was very, very close and they have not done as well in this space because it's very hard to say we've got all the answers and you can just sit on everything. So I don't want to say it's perfect, but I'm much less exercised about that. Then the access to the underlying data. Yeah, the less personally, but you know, cause I think again, Google shows, you don't have to now they're not about published peg page rank I think for us. But um, but for many algorithms you guys were incredibly open and generous with, with what's out there and we benefit greatly from it. Yeah, I should,

Padmanabhan A: I mean see files like there's a boat in a particular specific personnel. Llangollen a global one. For example, when we were looking at this low birth weight babies, we looked at on the computer vision and found how people are doing things but not for babies. They were doing adult body modeling and that has a lot of very interesting applications you can imagine from movies to fashion do or all that. But that core approach actually translates over, I mean it has to be done at Oxford and all that and the particular people that are working on it or we're quite willing to collaborate and you know, we actively collaborating with the people that have been doing it. One of the things we found is that when you go with the idea that you do something for social good, lots of doors open, people are quite quite open and willing to collaborate.

Padmanabhan A: That said, however, these are just one offs. I think what has been done in the industry, and again Google is a leader in is to create a tech platform as bunch of aps bunch of um, you know, uh, reference architectures for doing deep learning insights that other people can pick up and use and build and leave behind at something similar to that along with the data ecosystem has to be created for health and global health. And that, I think that's a journey. But clearly I think there's a lot of interest to do that. That's great.

Alan Weil: Good. Well, uh, we have left some time for you all to ask your questions. So, um, thought there might be a few, I see a little cluster right up here. So, uh, the mic, the microphone carriers, these two in the front row. Is that Patty, we'll get a microphone back to Patty.

Audience Member: Okay. Uh, hello. My name Jason Helgerson hovers does solutions group like for a long time I was in New York State Medicaid director. So a lot of my question comes from, from that, from the context of vulnerable populations and that people that the healthcare system often leaves behind. And um, I kind of want to Echo Alan's earlier point and I feel a little bit, um, uh, more concerned even you Allen in the sense that it feels like AI in the use cases that we're given are more sort of like incremental improvements on a platform that in my view is burning. Um, you mentioned some stats, billions of people don't have access to basic services. That's just the reality of the system in this on our planet. And so my concern is, is that we can't train our way out of that. We can hire our way out of that.

Audience Member: Um, and so the IBM, my fear is that technology has to be the solution. And especially for you in the center here is Google the company wanted the most powerful companies in the history of the globe, has unique role responsibility in my view to actually create a vision for how we can fill that gap. And my fear a little bit with AI and maybe you guys are a can dissuade me of this fear, is that we're a little too deferential to the healthcare system itself. We're a little too deferential to the role of physicians, at least historically they plate. And my hope, especially for companies like Google is not to, uh, to, uh, acquiesced to that fear and be willing to say maybe in the future, machines will in fact replace human beings. Because I fear that feel that the only way we're going to solve these huge access problems and really achieve equity is if we can find a way to automate.

Lizzie Dorfman: So interesting. Your thoughts. Yes. Fantastic remarks. And I agree with, I think everything that I heard you say, um, vulnerable populations and healthcare was my self defined focus when I was in school undergrad. Uh, when I was in graduate school I studied obstetric pharmacology because it angered me that we were basically doing research on pregnant women by giving them medications that we had never studied to if they were safe or they were effective. And the equitable distribution of risks and benefits from AI technologies I think is one of the most important topics. And the reason that I am at Google is because I think it is uniquely positioned to have impact on order of billions of people on this planet. But it has to be done. Being mindful of those who are already winning and those who are already losing. I I think, I don't have answers to your questions and I don't think that there is a single answer to your question.

Lizzie Dorfman: The part of the reason that I took the job that I did is because we need to have people who are, who are thinking about these problems and are engaging with people like you to better understand where, where are we doing okay and where are we not and what are the ways that we can make sure we have one direct representation that was funded on a grant that studied tribal communities of the Pacific northwest. And understanding the extent to which there can be harms that are unintentional that result from people trying to do good things is arguably without limit understanding the extent to which we, we harm people by trying to protect them from risks and then exclude them from the benefits of our work. So I, part of the reason that I love conferences of this nature is the chance to hear in a very concentrated manner from people who are very, very focused on underserved communities, vulnerable populations and regional and localized needs. The, that's why I'm here. [inaudible] and I and I intend to be working very hard at Google to make sure that these, this, these issues that you raise are prioritized in our herd. And Insofar as I'm able to speak on their behalf, I believe that most, if not all of my colleagues feel similarly.

Padmanabhan A: I think to me your message was health system requires disruption. A technology slash AI might be the way to do it. And companies, technology companies like Google should not be afraid to go take it all the way. I totally agree. One thing that they look at is that though disruption can sometimes more easily begin from the bottom, but the cell system is actually not functioning. And in some ways that's where we see the opportunity to get in. Sorry.

Catherine Mohr: Yeah. I actually disagree with a couple of the premises that you have because you're saying the population that needs services is the problem and I actually think the population that needs services is the solution, which is why I am much more about education. There are brilliant people everywhere in the world and a lot of them are growing up without the benefits of an education. But you can get a population to be able to care for itself and care for one another. If you can figure out a way to get people the training that they need to be able to be those carers. And you don't think of them as passive recipients of a healthcare system that we're structuring in order to serve them, but a, but active in being able to learn to care for themselves and learn to care for one another. And so the, the why my focus is on the educational side of it is I believe that the doctors and the nurses necessary to take care of one another, art in those populations that they need to, to serve and we need to fix how we're teaching them.

Mike Nohaile: Okay. Can I take on a slightly different angle on it, which is I disagree with another premise, which is that, um, I, I'm a little skeptical yet, uh, very impressed with Lizzie and the work that Google does, but I'm a little skeptical that it's tech companies are the most obvious solution to the problems that you outlined. And the reason is we've had a move fast and break it ethos. I the that better not say this Google, but I'm a little concerned about that ethos being the ethos that we have for the healthcare system. I think it's got to be policy, right? Maybe that's boring, maybe it's not popular or whatever. But I think it starts with, with, with thoughtful policy to address these issues. I'm not saying that AI and things can't help. I think it can be tremendously helpful. I'm a little worried that we blow up the system and assume that what comes out of it is better than what we have.

Mike Nohaile: It's exactly what you were talking about. We may be trying to help, but you may get, you know, some people doing, you know, Facebook moderation in ways that you don't like and you didn't think about that ahead of time because it's a complicated system. So it's not that I don't think the, the, the ideas right are the pushers, right? I just don't leave out the policy element because a lot of the problems we have in the, not in the underdeveloped systems but in the developed systems, our policy and our locks that make it very hard to make logical changes.

Alan Weil: Oh wait, you had another question here, right here in the front row and then we're going to go to, Patty will take you and that's probably going to be about the end of it.

Audience Member: Try to keep long questions and answers. We do. All right. So I'll try to make this quick. Dave, for costume at the anthems Ai Division and uh, question is somewhat twofold. So the first one is about let's make it workable. So we're trying to automate the human process with Ai. How do we get past that and look at issues that go beyond human discernment? Because I think there's a lot of potential there and you know, that affects a medical education and so on as well. What makes for a good medical student in the future in this world?

Catherine Mohr: So I'll go back, I'll leap in on this one. Um, if you say, oh, I'm going to start with AI and I'm going to look at all the places that I can apply it, you're going to stunt the kinds of problems that you're solving. If you say, I am looking for a better way to detect cancer and I can light it up with a fluorescent molecule and I can see it with a camera and that goes beyond to humans. I mean that has got nothing to do with AI, but it's a really good solution to being able to solve that original problem. I would just keep coming back to, you know, is the tool we have to not track of

Mike Nohaile: what are the actual problems we're trying to solve and whether AI is a good solution to that or are there better solutions to the problems. So that's

Padmanabhan A: good. All right. We can get a few more in at this pace. That's what we're going to take. These two are here. Patty,

Audience Member: thank you. That was great. Uh, uh, Patrick ever. Um, I think we all know at, we have low or healthcare costs in this country, but we also know at Ari history is every time we introduce new technology in the healthcare it raises class. And uh, robotic surgery is a great example of, of uh, so how do you see this growth of new technology actually helping us lower healthcare costs, not more screen iknowmed board drug therapy. The more tools add add cops and I'm not sure I see how what you've said of art or his pops.

Mike Nohaile: So, um, I'll take a shot at that, which is what we're also seeing that is more evolutionary as you can use, uh, data techniques to actually understand what's happening in the system and run counterfactual experiments and try to really understand what would have been better. So for a long time in Medicare system, they couldn't even understand what they were paying for. Right? I was the diagnostic field for awhile and you coded things and called TPT codes and it said you ran a a particular aisle ago, nucleotide assay don't know for what and we're going to pay you for that. Right? So it's very hard to figure out what you should actually do if you don't even know what you're actually doing. So I do think these technologies can hopefully give us a, an actual picture of both of what's happening and maybe even a good picture of what would be better. I still come back to policy though. That's not going to be sufficient. I think it could inform it, but I think you're going to need policy changes to move to assist in that rewards innovation. But thinking clear out some of the older stuff in the older ways of paying for things and get that off the books to pay for the new innovation. And for me that's a political decision, not just a data issue. Although I think data will be critical to helping make the right decisions.

Padmanabhan A: I mean, ultimately it's a question of amortizing the cost of innovation. Right? Right. In other words, the reason why they broke up is because the innovation required a huge amount of vocabulary investment and somebody is going to try to make money out of it. I think, I don't know the of it as a panacea, but really having democratizing the process of innovation and making it open and available to a large number of people by sharing the key elements which is in this case will turn out to be data is probably the only way to do that. The more people out there, the less the cost will be. Well, and I also,

Mike Nohaile: I would say, um, we need to be careful what we think

Catherine Mohr: our costs versus value. Um, big pieces of surgical equipment like surgical robots are often pointed to and said, oh, they raise the cost of health care because the math is really easy. You spent $1 million on a piece of equipment. But what's harder on that calculation is figuring out what role the costs downstream that it avoided. And if you do a net value on a hospital system that has this equipment, is it serving more people for less cost or is it at or getting better outcomes for less cost? Or is it for more cost? And what we've seen in US hospitals, so I'm going to defend my robot for just a moment, is that hospitals will get one and they start looking at the bottom line in terms of costs of complications and costs of readmissions and things like that. And pretty soon they find they have seven because for every single one that they're bringing in, they're actually reducing the overall cost of care to treat. And so one of the nice things that's happening in the U S is we're moving towards, um, a, a bundled payments and being able to look at an episode of care and that will actually make a more focus on really looking at what it costs to, to do care. But, um, these technologies are being used in systems all around the world, not just the u s system where everybody's already looking at the holistic cost associated with that. And every new technology is going to need to earn its way into our health care systems.

Alan Weil: So, I'm sorry, I'm not going to be able to bring you in. I'm, I'm very sorry, but we, we got 80 minutes. I can't turn it into 90. Um, let me just close by saying, I hope from this deeper dive, you got a sense of the heterogeneity of the potential applications in Ai, uh, that it's immediate potential is largely focused on improving capacity and accuracy, but there's a lot of unknowns about how disruptive it will be, whether, uh, it will lead to the cost savings. I'm going to challenge you on those robots, by the way. Um, but, uh, but fundamentally, I think the remaining issues in addition to the technological ones, which we've learned a lot about today, they're issues of policy there, issues of infrastructure and that if we think about those that will help us use these applications to be disruptive, to be cost saving and not just a creative and additive to a system that has a lot of challenges. So please join me in thanking our [inaudible]

Catherine Mohr: [inaudible].