

The
Economic
Club of
New York

ESTABLISHED 1907

The Economic Club of New York

117th Year
744th Meeting

Dr. Kai-Fu Lee
Chairman and Chief Executive Officer
Sinovation Ventures

February 13, 2024

In-Person/Hybrid Event

Moderator: Caroline Hyde
Anchor, Bloomberg TV

Introduction

President Barbara Van Allen

Good morning everyone. I think we'll go ahead and get started. We do have some stragglers. The weather, as you can see out the window, has not made this morning's commute easy, but we must carry on and try to keep to our schedule. And we also have folks joining us virtually. Usually we have at least as many in the room, we're twice that in our virtual audience.

So welcome. This is the 744th meeting of The Economic Club of New York. I'm Barbara Van Allen, President and CEO of the Club. Known as the premier nonpartisan forum in the nation for social, economic, and political issues, The Economic Club of New York has been in existence now for more than a century, and we've hosted over 1,000 prominent guests during that time with great excellence frankly that continues up to today.

I'd like to extend a warm welcome to students who are joining us virtually from Rutgers University, Zicklin School of Business at Baruch College and Mercy University.

Congratulations also, and welcome to members of the 2024 Class of Fellows. We have the largest class ever, and equally diverse. On behalf of myself, the board, and our membership, we look forward to being with you and enjoying this year's special

programming.

For today's program, we're honored to welcome Dr. Kai Fu-Lee. Dr. Lee is Chair and CEO of Sinovation Ventures and President of Sinovation Venture's Artificial Intelligence Institute. Sinovation Ventures is a leading venture capital firm focusing on developing the next generation of deep tech companies. With a unique VC+AI model, Dr. Lee also leads the venture-building efforts of seven AI startups.

Before founding Sinovation, Dr. Lee was the President of Google China, and a senior executive at Microsoft, SGI, and Apple. He received his bachelor's degree in computer science from Columbia University, a Ph.D. from Carnegie Mellon University, and Honorary Doctorate Degrees from both Carnegie Mellon and the City University of Hong Kong. He co-chaired the Artificial Intelligence Council for the World Economic Forum Center for the Fourth Industrial Revolution and is recognized as "TIME100 AI" Top 25 AI Leaders in the World, 100 Most Influential People in the World, by *Time* magazine in 2013, *Wired* 25 Icons, and followed by over 50 million on social media.

In the field of artificial intelligence, Dr. Lee built one of the first game-playing programs to defeat a world champion, as well as the world's first large-vocabulary, speaker-independent continuous speech recognition system. His 2021 co-authored book, *AI 2041*, predicts how artificial intelligence will change our world in the next twenty-year

horizon.

The format today will begin with prepared remarks from Dr. Lee followed by a conversation in which we're honored to have Caroline Hyde as our moderator. Caroline, for those of you that don't know, is Anchor at Bloomberg TV, where she hosts their flagship daily technology show from New York. We will end promptly at 9:00 a.m. As a reminder, the conversation is on the record. We do have media. Without further ado, please join me in welcoming Dr. Lee to the stage.

Opening Remarks by Dr. Kai-Fu Lee

Thank you, Barbara. It's such an honor to have a chance to present to such an esteemed group of people here. I want to just give an overview so we have some background on the amazing technological revolution we're experiencing. I often talk about generative AI as the most important technology revolution in human history, and here is why.

AI has been around for a long time, but it's really plagued with the issue that for every domain you have to retrain the AI. In my book, *AI Superpowers*, I see a copy there, I talked about data is the new oil. Meaning that the more data, the better AI performance. That comment has transcends historic, basically classical AI, and generative AI. So the

question that's thrown upon us is, well, how do we give the most possible training data ever so AI can be super-smart. And the answer is generative AI, which is trained in an unsupervised kind of way.

If you've read a little bit about it, it's actually a fairly simple algorithm where you feed all the text in the whole world and basically say predict the next word. And that has forced AI to build up its own intelligence, which is a black box of neural nets and weights, hundreds of billions of weights that basically comprise intelligence. So if you present five billion words, you only allow it to retain a few hundred billion, then it's forced not to memorize, but to learn to categorize and inference and to be able to predict.

So that's the power of generative AI is that you build the middle part, which is called the foundation model from all the data in the world. And you don't have to train it. You don't have to tell it, this is John, this is Mary, that's a table, that's a building. You just give all the data and it trains itself. That's the amazing thing. So every six months it has a ton more data and it gets smarter.

And how fast does it get smarter? Well, this graph shows you pretty much all the important AI systems in generative AI. And you see a graph from left to right, and it looks like a mild curve, but it is not mild at all. The bottom horizontal line is the average human intelligence. Think of that as IQ 100, if you will. Now the best AI system is now at

roughly IQ 300, if you will. So basically it's outperforms any single human by a mile. And one of us would be lucky to get 55 on this IQ test. So, you know, multiply by 3, for normal IQ. This is not IQ exactly. But AI is now at 90.

The only way you can match AI performance is if you get, you know, a few hundred Ph.D.'s in a room and for any question that's given, find the expert and have that person compete with AI in history and philosophy and logic, math, AI, and so on. So that's how good AI has become. So it has gone from human level intelligence to super, super, superhuman intelligence in merely two years. So imagine what will happen in the next year and the year after that.

And so I think of it, so let's back up to business. What's the business implication? I think forgetting all about the AGI and the science fiction elements of how brilliant AI will become and the existential threats that it may have, all those are legitimate problems, but this is the Economic Club, so what is the business implication?

I think the business implication is simply that we've had two technology revolutions – PC-mobile, and now we're in the AI revolution. And if the PC revolution is characterized by the phrase "PC on every desktop" which was brilliant when Bill Gates formed it back, you know, 40 years ago, mobile revolution is anywhere, anytime computing that knows who you are and where you are. Something much bigger than PC on every desktop. AI

is super intelligence injected into every app. So it should be self-evident. That is ten times bigger than everything that's happened in the past.

So what are the killer apps? Well, it already has two killer apps. This graph shows you the number of months required for an app to reach 100 million users. Previously the record was held by TikTok at nine months. And before that, WeChat, Instagram, WhatsApp, YouTube and Facebook and so on. But now these records have been blown away not only by ChatGPT at two months, but Character.ai at six months. So these are examples that the killer app is already out there. We just don't think about it that way. And there's no doubt in my mind that in another two years, this whole picture will be all orange, all AI apps. There will be no more mobile non-AI App that remains on the Top 10 or 20 lists in terms of what we use.

So what's important is to build these apps in an AI-first kind of way. What does it mean to be AI-first? It means the app thinks about AI in first principles. In fact, if you remove AI, the app stops working. To give you an analogy, a mobile first app would be Uber. That would be a great mobile-first app because if you remove the mobility element out of it, Uber stops working. Can you use Uber on your desktop PC? You cannot. Right? So Uber was built as mobile-first.

What's an example of an app that wasn't built as mobile-first? Yahoo wasn't. It was a

website that, with a fine app, but it could work on the phone or PC. It doesn't use the fundamental element of mobile, which is it knows where you are and who you are. Now, AI is the same. So a future of successful apps will not be simply glue AI onto an existing app, such as Microsoft Copilot. It will be fundamentally a breakthrough app that's built with all of AI's power built in. And if you remove the AI, the app stops working.

So I'll leave a little bit of something to your thoughts, and we can discuss what those apps might be in the conversation that follows. And this is a quote from Andreessen Horowitz, which I agree with completely that the generative AI, the value is somewhere between all of software and all human endeavors. Thank you.

Conversation with Dr. Kai-Fu Lee

CAROLINE HYDE: Fascinating jumping off point. And what I worry about is that we have about 45 minutes or so and we're going to want a lot longer. Thank you all for braving the weather.

It's a joy to be sat here next to you, Kai-Fu. I want to drill into that element that you say AI first. I know that you're busy. Not only have you been an innovator throughout, but you've been an investor, but you're back to innovating. You're building your own AI apps, the killer ones. Can it be retrofitted? We associate the big winners as Microsoft,

because of their ability to see the opportunity in Open AI. We see that Google is still really front and center when it comes to a race, if we think about it here in the U.S. But can they be AI-first?

DR. KAI-FU LEE: Theoretically, they're the best equipped to be AI-first because they have the most AI experts in those two companies, if you count Open AI in Microsoft's count. But what they face is a fundamental business issue of cannibalizing your own business. So for a large company, imagine, if you will, you're a Microsoft or Google, the first thing you're going to think to want to do once you invent this technology is how do you make money? And, of course, Google will make the most money from putting it in Search. Microsoft will make the most money by putting it in Office. So that will be the low-hanging fruit.

But where Microsoft becomes a disruptor that says I'm going to build a new engine that's AI-first that causes more people to no longer need or want to use Office or pay for it. Few companies, if any, have ever done that. Right? When Kodak invented the digital camera, they shelved it because they think it would kill their film business. It did, but it wasn't because of them, but other companies. Right? So the same thing happens with Intel not embracing mobile, Qualcomm not embracing AI. And the list goes on.

So it's fundamentally really hard for a large company not to take this technology and put

it in their most, bestselling product, to make that easy low-hanging fruit. And it's incredibly hard for them to build a destructive product that will challenge, cannibalize, and potentially kill the revenue that is currently a cash cow for them.

CAROLINE HYDE: When I think of what you are doing now with 01.ai and not being 01.everything is how it would be interpreted in the U.S. You're doing things in an open-source manner. And what was so interesting is you sort of referenced the two killer apps thus far that developed AI-first and you're thinking about Open AI's ChatGPT, you're thinking about Character.ai. You also have the curve in the graph that showed sort of proprietary LLMs and open source LLMs.

When I was reading *AI Superpowers*, you were thinking about ultimately when the next iteration of AI would happen, how it would happen. And you sort of hoped – dare I read between the lines – that it would come from academia and it would be open. And it wasn't. It came from Open AI, but that's a proprietary large language model. How did you feel when it was a proprietary model?

DR. KAI-FU LEE: Well, I think Open AI should change their name to Closed AI to avoid misleading those who don't know. It did start as an open company, which I think was the right direction. But as greed is always the enemy of goodwill, once they saw the opportunity, they thought closing it down was the best way to make the most money,

which I don't disagree with. But the problem, when you have one or two top closed models, then everybody becomes beholden to those models. If you think this is the way we get to the, you know, IQ 500, IQ 800 system, that's so amazingly smart, but also so amazingly dangerous to let one or two companies control that is very dangerous.

Furthermore, in a closed model, who loses? Right? It's really those who need access to technology. Professors, researchers, hobbyists, entrepreneurs, students. They would not be able to learn how to train and tune and align these models because underneath is a black box that can't be trained or aligned or tuned. So I feel open source is a way to provide that accessibility to everybody. And it's not a country-by-country, but really a segment by segment. The next innovation is much more likely to come from the group of researchers, professors, students, and entrepreneurs, than it is from any large company. So to remove those tools, I think would be the wrong thing to do for humanity.

So we decided, even though we're a much smaller company at 01.ai, we will partake in the open-source revolution. We have made, everything we've done are open source so far. As the other companies have done that, Meta, I think, has done the same. And, of course, it doesn't mean in an open-source company, everything is forever open source. We still do need to make money and satisfy our investors. So we will figure out ways to make money either with larger models or APIs or applications. We think providing that access to everyone is an important and a responsible thing to do.

CAROLINE HYDE: So when you are thinking from an open-source manner but also applying it to some of the killer apps that you want to build, which industry do you target first? I know this was an audience question, really wanting to understand what are the first ways and means of adopting generative AI, in which industry groups?

DR. KAI-FU LEE: Okay, so first, you have to also acknowledge generative AI has some problems. The biggest problem in the near term is hallucination. That will get fixed over time. So I can think of applications in the near term that can deal with hallucination. One is content creation where the human remains the author because then the fact-checking remains with the human. So a Microsoft Office Copilot or a disruptor version of that, I think would be reasonable.

Another direction, inspired by Character.ai, is you can align these AI to be interesting AI people. They can be empathetic, cute, funny, romantic, knowledgeable about food, tourism, or your business. And they make wonderful friends because they don't lose their patience. They don't get angry. So many of us, and certainly many of our children, are going to end up having many AI friends. So I think that direction of moving towards entertainment, social, short-form video type, I think is another big one.

Now, back to your question about industry, I think the former, productivity content application, will find the biggest usage in basically routine white-collar work. So think

about your departments. Many of you run successful businesses. Your HR Department, Legal, Finance, back office, people who do assistant work, people who copy, paste, and do repetitive things, although each task is a little bit different, well, AI can easily learn to do that.

The latter part of social, think about parts of your organization that does social type of tasks that doesn't require deep human to human connection. So customer service, a sales organization that is doing the first stage of pipeline generation, not the part that's relationship building. Those kinds of social interactions can also be automated. Those would be the first areas for companies to think about adopting Gen AI.

CAROLINE HYDE: It feels as though, and this is an audience question, so 2023 was definitely all about ChatGPT and Sam Altman all in one. 2024, therefore, with some of the applications that you're talking about, does it become about AI agents? Are we there yet? Or is there something more that we only focus in on this year for its application?

DR. KAI-FU LEE: I think there are two big things that will be the theme of 2024, one of which is agent. So think about the content creation that AI does. Right? You can tell it or give me recommendations for the best New York restaurant. I love Italian cheesecake. Or, you know, tell me what I should buy for my wife for her birthday. And it gives you excellent recommendations. But why recommendations? Why not just get it done? Why

not just say book me the best Italian cheesecake place near my hotel or buy, my wife's birthday is coming up, buy her flowers, cake, and a present. And then it's done. So that task of taking recommendations, content, into action and it gets done, is the work of a brilliant personal assistant. And AI will be able to take that on and we'll see it begin to happen this year.

I would add a caveat that it's not trivial to take that on because, think, what will 1-800-Flowers think about my agent getting the flowers for my wife? They would lose their brand and attraction for their customer. What would Amazon think? So effectively, the ultimate agent will be the greatest convenience for us. It will save, you know, hours per day for all of us. But it will relegate an e-commerce job into a warehouse. So they will be fighting very hard against it. So I don't predict agents will be powerful this year yet because there will be many companies who will not want it to succeed, or who will want the Amazon agent to succeed, but not the Google agent. So that, I think, is the first part.

The second, I'll touch on briefly, but it's really important also, is what we call multi-modal. That is AI is starting to understand not just text, but images and video and sound and audio and begin to understand the world, with a world model in which we live. So this is already beginning to happen and we've actually launched a multi-modal model as open source as well, as Google and Open AI. They've launched their products, but not open source. So we've been kind of trying to push on that side. Currently, it's mostly

with images, which are relatively easier, then video, then audio. And then eventually what's called embodied AI, where you're going to see robots roaming around, touching, measuring, creating data and feedback.

Because unlike text, where we have thousands of years of graded texts you can feed, a world model doesn't exist. So robots have to feel it out. Kind of like how Google's cars that drove around and measured the streets was quite a novelty to see ten-plus years ago. But now it has transformed into value in Google Maps. You can see anywhere in Google Maps. The same will happen with these robots that feel the environment. But that's probably, the embodied AI is definitely not for this year. So this year will be agents in relatively smaller demonstrations and domains, getting people excited, but not working full-scale. And it will be multi-modal with image and audio and video, but not quite embodied. That will be later.

CAROLINE HYDE: Will it also be a year of job losses?

DR. KAI-FU LEE: This year will be the beginning of job losses that are measurable. Previously, this was not easy to track. And a lot of CEOs don't want people to know this so they give excuses like, well, we'll just elevate our employees to do more complex things. Well, how many more complex things are there? And would you really train a routine worker to do complex things? Or would you hire a trained person to do the

complex things? So those arguments don't hold water. They've already happened last year. But I think this year, numbers will come out significantly.

Also, I think last year's numbers were masked with issues related to immigration, related to Covid, related to remote work. So that makes the job, the unemployment rate not as comparable to previous years. And then this year, I think, we'll have a reasonable 2023 to compare against. The same basis for Covid, the same basis for remote work, presumably similar in immigration, and then that difference will be very pronounced, I think, this year.

CAROLINE HYDE: Having written a book that was thinking about 20 years' time and envisaging 2041, you go into the ethics. You go into the challenges, the worries. Are we capable of dealing with this level of disruption, of job losses, of Amazon and e-commerce fighting back and not wanting to see its own agent not winning? Are we ready for it, first from an American perspective, but a global perspective?

DR. KAI-FU LEE: No. I think we're definitely not ready. I'm usually an optimist. In both of my books I showed a highly optimistic future despite recognizing the risks. But it's just happening too fast. And all the issues you mentioned plus issues like suitability for one model for the world, and if each country wants its own model, then what happens? Who will build that? Can it be built? Will that be good or bad for the world? These are all

unanswered questions on top of what you just asked. And governments don't usually move at the fastest pace to address these. So that's why I think forums like these, people like yourselves, are probably the most likely solution to the problem, is that it may come from the private sector.

CAROLINE HYDE: Can you paint a picture of what a solution could possibly look like? Is it that each country does have one model that wins them all, and you hope that it at least it might be open source?

DR. KAI-FU LEE: Yes, I think the end game isn't clear. I think what is clear is the open-source approach is creating adoption and putting, becoming serious competition to closed source. So it hopefully will, at a very minimum, force the closed companies to be more open. At maximum, there are some people who think the open-source models will win because it will align the massive number of people, just like Linux has in some sense won the OS revolution at least on the server layer. So I think those are possible. So I think that closed source, open source, I feel, we're on a reasonable trajectory.

You know, Cathie Wood from Ark, has projected that open source may become the best model, at least will give the closed model a run for its money, which I share. So I think that is the least of my worries. The bigger worries will come from people not understanding what's happening, the loss of jobs, the further polarization among

classes of people, types of companies, and countries. Those all seem likely. And not to mention the bad people who could use this to invent the next biological weapon, figure out how to get away from airport security, use it to manipulate elections, and mislead people, create deep fakes, the list goes on.

CAROLINE HYDE: Do you think that any of that can be regulated upon in the here and now? I mean if you're thinking of an open-source future, much of the pushback on that is that bioweapons suddenly fall into the wrong hands or the capability of building them. What is your answer when people push back on open source for that reason?

DR. KAI-FU LEE: Well, closed source is potentially even more dangerous. Right? Because fewer people understand it and you're relying on one company or two companies to build, to fight against bad people. With open source, you're aligning the great majority of good people who want to do good things to do that. I mean Linux is not less secure because it's built by open source. And there are closed-open-source collaborations so I wouldn't worry too much about that.

CAROLINE HYDE: There's often a narrative, particularly in my industry, the media, of who is winning the race? And when they say who, they just mean which company or which open-source large language model is winning in the U.S. We mean which country is winning. And I know it's a question from the audience of with your deep knowledge of

knowing, back in *AI Superpowers* you were sort of trying to measure ultimately Chinese innovation versus U.S. innovation. And at that time, China was behind in AI, but you foresaw that it would leapfrog. When you're building, as you are, in China at the moment, how do you, do you think it can be compared? Or there is actually a lot more internationalization and global coordination from an academic perspective that we shouldn't think like that?

DR. KAI-FU LEE: Well, first I think, I have to answer your question, which is that I stand by my comments in *AI Superpowers* that U.S. is better at innovation, China is better at execution. So to the extent you believe we're in the execution phase, China will surely catch up. To the extent you believe there are more innovations coming, the U.S. will hold its lead. And to the extent these innovations are – this is where there's a paradoxical – to the extent that these innovations are held behind corporate walls, the U.S. will hold its lead longer. But that would argue against open source, but that's another dimension of the question.

So I think we're seeing that play out in Chinese AI. But I think I could spend much more of my time thinking about closed versus open, how to alleviate the dangers, because I think we first have an obligation to the world, then to our countries. And I think the world needs more open models. The world needs to solve the bigger issues that face all countries. And, in fact, I think both U.S. and China are likely to end up just fine no

matter how the geopolitics play out because both have strong AI companies and vibrant markets.

It is the countries that don't have the resources that become a problem. Right? Because if all the routine jobs are going to be gone, many countries that have mostly routine jobs and no AI technology and no giant to create the income that can be taxed to redistribute and help people get back on their feet, I think this is a disaster in the making in the countries that don't have the means to profit but have large numbers of routine workers and those are likely also the countries that aren't aware how fast this tidal wave will come.

CAROLINE HYDE: You travel a lot. You go to see usually relatively early adopter kind of countries, but which countries are starting to realize and asking you to go visit? And where are you worried about?

DR. KAI-FU LEE: Yes, I think, you know, clearly the western world is very excited about this. France is all excited about Gen AI. The UK is very excited about this area. I also think some Asian countries, Singapore, Korea are usually the ones that are thinking ahead. And then some emerging countries like UAE and Saudi are also thinking about this. But these are not the four countries I was worried about in the previous question.

CAROLINE HYDE: Where are you worried about?

DR. KAI-FU LEE: I'm worried about, you know, countries in Africa, poorer countries in Southeast Asia, in Latin America and so on.

CAROLINE HYDE: So, we put it to the room that some of the problems and challenges we face can be fought or can be tackled almost by corporates in and of themselves. What responsibility is it upon business leaders to think about these countries, these people, or indeed these workers that they employ in the here and the now?

DR. KAI-FU LEE: Well, to the extent you have a labor force in those countries, I would say finding ways in which they can find the path at least to providing for their families. Right? For example, people who do outsource work to Asia may want to see what part of that population can take on the responsibility of managing AI to do the outsourcing. And also when AI takes over, you still need some routine work. We just have to be on the watch as to what they are. It turns out the biggest area is labeling data right now. And so if you have an outsource center somewhere in Asia that do routine work but these are intelligent people, even well-educated, which does happen in Asia, I think they can do interesting labeling work.

You might say, well, labeling, that's boring, it's marking table, chair, humans. It's not. It's

actually labeling of here are two different large language model answers to this complex question. Which one is better and why? So we hire labelers and what we ask for the labelers are those who do well on the SATs, equivalently. So these are highly intelligent people who are caught in maybe some routine work, perhaps in a country that doesn't provide easy retraining and upward mobility, and it turns out labeling is an area. I'm not claiming labeling will solve, you know, 100 million job problems, but there are probably new emergent work that you can help the groups take on.

I think the more, you know, earlier Amazon had a brilliant policy of providing four-year training to all the employees, which I thought at the time was brilliant and very generous. But I think those are still nice if you're willing to offer it, but they're so altruistic that I don't think most companies can afford it, and I think they take too long.

So I think it's time, because how fast the tidal wave is coming, you know, now it's labeling, and then you watch out for other things that may help a more rapid transition of the people who are inevitably going to lose their jobs working for you in this country or more seriously in poorer countries. What can we transition them into within the next six months? I think that's more urgent because I can't imagine most business people would say, well, we'll lose our profitability and offer 5,000 people a four-year retraining. It's just too altruistic and not profitable.

CAROLINE HYDE: You think it's six months?

DR. KAI-FU LEE: I think, well, I think business leaders can take six months to redirect and help and retrain their employees before they're laid off because six months is not much larger than a severance package. I'm just speaking pragmatically now.

CAROLINE HYDE: Sam Altman, for example, has always countered the demons, should we call them, with the idea that, you know, just the more data, the more we'll do with it. The more iterations and innovations we have, the more compute power we'll get, the more productivity we can have. We're not suddenly going to go to a universal basic income and everyone can just have a three-day working week and we will create more, do more, be more productive. Do you agree with that? Or do you think ultimately we can have a more, sort of productive and idealistic society where people can rest more, focus more on creativity?

DR. KAI-FU LEE: I think if you look at a 30 to 50-year horizon, I'm open to that train of thought. That, you know, people are incredibly, humanity is incredibly resilient. We've overcome a lot of job displacements in the past. It's just that this round is too fast, too sudden. We're too unprepared. And the jobs are going to go quickly. And the new jobs are going to arrive more slowly so there will be at least a misalignment of the disappearing jobs and the appearing jobs.

CAROLINE HYDE: What sort of strikes me is if you followed Kai-Fu Lee's work, you have been at the forefront of AI for two decades. You were creating these standout labs with Microsoft in China, you were working...how are you caught off guard by the pace of innovation? It amazes me.

DR. KAI-FU LEE: Yes. So to defend myself a little bit, if you have *AI 2041*, you will see I did predict the future. The power of generative AI is all Chapter 3 was about. The scenarios are exactly as happening. The technological breakthroughs are happening. The interfaces are proper and correct. It's just that it's just, I didn't see it coming this fast. I didn't think the particular algorithms would just work. Right?

If you think about what I describe, take the whole world's data and teach an AI to teach itself to predict the next word, that would be the emergence of AGI. It did seem inconceivable to me. I thought five or six more huge breakthroughs were needed. I think it's understandable. I think many people thought the same. But Open AI was proven right. I mean despite what I said earlier, I hold great admiration for the persistence and the belief and the brilliance of those people who engineered the system that worked and surprised, I would say 99% of the AI community.

CAROLINE HYDE: Many at a time six months ago said we should have a six-month pause. And everyone said, well, that's totally impractical to actually effect. But does it

feel like that if we've just got, as you said this tidal wave coming towards us?

DR. KAI-FU LEE: If it could be implemented and reliably detected, I would actually be in favor of that, but it's just too impractical. I'm a pragmatist. This is not like nuclear weapons. There are no fish in a bowl materials that can be detectable. And people talked about putting some kind of fixed code somewhere in the GPUs so they would shut off upon something. But this is science fiction. I think, you know, people's curiosity, intellectual, technological curiosity, just simply cannot be suppressed. And people will feel so strongly about the innovation, they will find resources, find GPUs to pursue their life's dream. And we can't stop it unless there's some way the resources with which it takes to build these large language models can actually be controlled, and they cannot.

CAROLINE HYDE: Is that in any way the blocker at the moment, access to compute? How have you managed that at your own startups?

DR. KAI-FU LEE: Yes, I think the GPU shortage is a problem worldwide. And I think the American policies about not shipping GPUs, high-end GPUs to some countries does cause a short-term inconvenience. But I wouldn't call it a stifling effect because this is all software and hardware. If you have good enough GPUs that are not great, that are hard to program, if you have a team of hardworking, persistent people, you will overcome it. You may need to pay a higher price. You will need to take longer. You'll need more

people. And you may not be able to build the absolute best large language model, but you can build a pretty good one as we've demonstrated.

CAROLINE HYDE: Do you need \$7 trillion worth of investment in chips worldwide, which is about 8% of...

DR. KAI-FU LEE: I think anyone in investing would love to get \$7 trillion. We'll figure out what to do with it. But if I had \$7 trillion, I would use it to figure out the solution to the problem we discussed rather than building some science fiction future which compounds the problems which we already cannot solve.

CAROLINE HYDE: So what is the next innovation? Large language models, generative AI felt like it's, as you said, a tidal wave towards us. Is it now iterations? Is there going to be another major shift that you see coming next?

DR. KAI-FU LEE: There are people proposing new things. I think a lot of what I talked about is not trivial. Agent technology is not trivial. Embodied AI is definitely not trivial. We don't know how to do it. World model is not trivial. So I would say those are the ones that are pushing the frontier for the AI people. Beyond AI, I think the other, even bigger, thing that takes a little bit longer is quantum computing. We're not sure, 100% sure it would work. And if it does work, it will likely take about ten years. So I think most

of us are putting that aside and worrying about AI for now. But in five or six years we'll need to start to, people will start to worry about that.

CAROLINE HYDE: So take us back to how you got, you got it right, it was just the time frame which was slightly off. What will 2041 look like? What will quantum computing change for us in the way that AI is currently changing?

DR. KAI-FU LEE: Yes, it's a double-edged sword. So the quantum computing in ten years, for sure, almost for sure in 20 years, will deliver basically a transformative underpinning of our computing. All the existing software including generative AI will be thrown away, and they have to be reinvented and rebuilt let's say. And it will solve problems that couldn't be solved before. So on the positive side, you know, we'll be able to model our body and what happens with a new drug in terms of curing a disease for humanity or for a certain segment of our people or for myself. So that's the power it has. But it will also have the even more amazing power to build weapons of mass destruction. And for every drug you can invent with the technology, you can also use the same technology to invent toxins, poison. So everything is going to be double-edged in the future.

CAROLINE HYDE: You said briefly that you were always an optimist and now perhaps you're becoming a little bit more aware of the pace of change and disruption. Would you

call yourself still an optimist? Are you optimistic for your children, for our children, our children's children?

DR. KAI-FU LEE: I think I'm cautiously optimistic because if we look at the history of humanity, we've faced a lot of crises and eventually our resilience, resourcefulness, and brilliance have overcome them. And at the time they were overcome, we didn't know how they would be overcome. So I'd like to hold out that belief. So I don't have the answers to the challenges we discussed, but I think history should give us faith that we will eventually overcome.

CAROLINE HYDE: What onus do you feel to question, to talk more about the ethical implications and questions that you have? You're doing this. You're talking to us now. Is this almost as big a role as the innovator and the investor that you have? You're going around and talking to governments and policymakers as to how we can prevent an out and out disruption?

DR. KAI-FU LEE: Yes, I think that's one. And then, as an author, I propose a lot of solutions. These are not for sure solutions that would work, but they hopefully will begin the imagination of people who care.

CAROLINE HYDE: Talk us through one that you're really focused on at the moment.

DR. KAI-FU LEE: Well, I'm not focused on it, but in the book I talked about a few ideas. These are not things I could focus on. So things I could focus on for my company is open source and basically outreach and let more people be aware. That's the limit of what I can do in my position. But in the book I talk about what's the replacement for money. How do we keep people happy even when they don't have jobs? And they're really out-there ideas, which probably will not work, but I wanted to sort them out as the straw man so more people can come up with their own ideas.

Not to governments. If you give *AI 2041* to government people, they'll generally say, well, this is some dreamy nonsense, which it is, but at least we need to get the dreamy nonsense started because the ultimate solution to these problems aren't just going to be how do we regulate AI, how do we catch the bad people, how do we use existing laws and have new laws. I think those cast much to conservative and slow-moving an approach. While my ideas may be wild, I want to get more people thinking, and maybe someone else will have an idea that works.

CAROLINE HYDE: Do you have optimism in self-regulation?

DR. KAI-FU LEE: No. Because the greed will drive companies to do things and rationalize them and that has happened in social media. That will surely happen in AI. I do think some government-required, mandated self-regulations that is measurable can

be a very useful mechanism. But I would say the government, I think the first priority would be to use existing laws to prevent catastrophes. There's already laws against falsehood in advertising, against fraud. So apply those to AI, gen AI-type of outcomes. That's the lowest hanging fruit.

The other is, well, you, Gen AI companies say you will self-regulate, here's how. Here's how you will be measured. And here are the consequences. I think that will carry meaning in the self-regulation. So I'm not against self-regulation. I think it's actually quite important. When I was with Google, at that time we were very much self-regulating in terms of caring deeply about the relevance of search results. There were third-party watchdogs that watched, even though it wasn't the government. So I think such mechanisms can work.

And the governments should not try to invent a new way to regulate Gen AI because it probably can't really understand them to the depth to manage them. And if there were an easy way to manage Gen AI that's brand new, one of us in AI would have come up with it. It's not trivial. So I would say the two approaches I outline are more practical.

CAROLINE HYDE: What do you advise your own children to be doing at the moment to be learning, to be adopting?

DR. KAI-FU LEE: Well, first, embrace the tool, embrace AI. Use it to advance yourself in the line of work that you do. Don't fight it. Second is that it's all the more important to follow your heart and do what you love to do as opposed to taking on a position that you don't like. When you don't like what you do, then you have no hope of competing against AI. If you love what you do, again I have hope in human resilience and brilliance to find a way out of this challenge.

CAROLINE HYDE: Too many people in this room are not adopting AI quickly enough. Why are you seeing companies being too slow?

DR. KAI-FU LEE: I think almost all traditional companies are too slow. I haven't seen any traditional company that's fast enough. And I think you should get some help, you know, management consultants and your IT departments, hire a chief AI officer.

CAROLINE HYDE: At great expense.

DR. KAI-FU LEE: Bring some college kids, some new grads. I think we have some online. Those are not going to be at great expense. And some of my companies' best people are in their mid-20s. And these people actually grew up not with all the handicap of using old AI technologies, so they come in thinking large language model, generative AI for everything. And I love working with them. Maybe that's a reasonably priced

solution for you.

I'll give you one example. I know a company that builds very ordinary games in China. And they have 1,300 people. They do understand a bit about technology. They're not super-technologists. But, you know, in one month they reduced that from 1,300 to 800 by just making AI, using AI to generate all the images that were being used, that alone. They replaced the graphic artists. So that kind of number and that kind of speed should appeal to you.

And even if you are not ready to embrace Gen AI in a big way, find someone to go through your company, find the easiest area to adopt it. Either to reduce costs or to create value. And let that be the catalyst that will cause your management team and leadership team to embrace the technology. You'll need a relevant real win in your company to convince the other people in the leadership team who might feel threatened or choose not to believe.

CAROLINE HYDE: You turn to your 20-somethings that you've hired, who else do you look to, do you read, do you absorb, do you listen to, do you call for the cutting edge of AI expertise at this moment?

DR. KAI-FU LEE: I think actually there are many brilliant people on Twitter, I should call

it X now. So if you follow some of the smartest people in AI on X, and also there are two types of people you could follow to help yourselves. One is the true AI experts. For example, you could look at the Davos Panel I was on. These were super experts discussing the future. Those types of discussions and the people who know deeply can really help you understand the technology. One caveat with these super Ph.D. technologists is that they're not always practical and they usually have something they want to prove. Maybe their technology or their company. So be a little bit careful but they're the experts.

There's another group of people who make money by proposing tools you can use. And these people are not technically deep and they are paid to promote tools. But among the tools they promote are some good ones. And start to use them, and I think they will give you ideas on how your company might embrace Gen AI.

CAROLINE HYDE: Give us a day in the life of Kai-Fu Lee using tools. What do you now totally depend upon as an app?

DR. KAI-FU LEE: Right. I don't want to really do advertising here, but I'll give you an idea. If you follow me on X, you'll know what tool I'm talking about. We have a tool that is intended, we've invested in a tool that's really ambitiously intended as a disruptor to Microsoft Office. Think for a moment, if you have an IQ 300, 100 Ph.D.'s, expert sitting

next to you, would you really type anything ever again? You would not.

What are the main two things you want to do with such a super intelligent AI assistant? One is read for you. Because we're very slow in reading. Read 5,000 pages for any problem you have. Pick the most credible sources which surely you and your team can do. Feed them all to this AI engine and then have it learn and then tell you back what you should know. Answer your questions or teach you what are the questions to answer. That's the first step.

The second step is generate content in a way that suits our minimal bandwidth. I mean, I think our brains are pretty good. They're not as good as super AI, but we have some strengths. The one area in which we are super limited is our I/O bandwidth. We read very, very slowly. And we're not very good at reading thick things so we lose our attention. So we like things like flow charts and the McKenzie 2x2s and the PowerPoints and SWOT analysis. Right? These things are definitely way over-generalized and not the most suitable way for, that is the best in explaining complex concepts. But they're a very easy way that everybody can grasp.

So the AI should produce those things – PowerPoints, line graphs, flow charts, McKenzie 2x2s and SWOT analyses. So that is a tool that I would really be able to use, that maximizes the power of AI, that it can absorb, learn, and teach us back, and makes

up for our problems, liabilities, which is minimal bandwidth. So have it tell us in ways that we can quickly absorb and understand and also potentially use it for our employees, customers, bosses and so on.

CAROLINE HYDE: I think I might know the answer, but do you mind that AIs have been, that large language models have been trained on your books? Do you want royalties for it?

DR. KAI-FU LEE: No. No. I believe LLMs learning from text is something we have the privilege of providing for, so I'm happy to have AI learn from all of my books. I don't want any royalties. I mean the only case I will have is if they start printing books and selling Kindle, then I would be unhappy with that. But they're using it to learn. This is all of us contributing our collective knowledge. I would have no problem at all. And, by the way, I am training a special AI with all of my writings and all of my posts on X and others that could potentially give the next talk in my place. We're also training a version of me with a video and voice, and so it will talk like me, look like me, and think like me. And with Barbara's permission, maybe I can offer another talk in nine months that is not me, although I can appear at the end of the talk.

CAROLINE HYDE: Yes, please be here. Kai-Fu Lee, it is such a joy to be with you in real life, IRL, as they say. And I hope that we still get the privilege of being able to do it

in real life even if there's going to be an AI version of you pretty swiftly.

CAROLINE HYDE: I'll make one for you too.

CAROLINE HYDE: Oh, thank you.

PRESIDENT BARBARA VAN ALLEN: Well, I have to say, I think that was one of the greatest conversations that we've ever had at the Club. So thank you. Really, really exciting. So looking forward, which we always like to do as we wrap up, on the 21st of February, we have Melissa Kearney, the Neil Moskowitz Professor of Economics at the University of Maryland. She'll be discussing her new book, *The Two-Parent Privilege*, with Club Trustee Greg Mankiw, the economist up at Harvard. The book makes a data-driven case for marriage by showing how the institution's decline has led to a host of economic woes.

Seats are still available for our first-ever event in Florida on February 22nd featuring Tom Brady, Seven-Time Super Bowl Champion. So if you're not able to attend but you have guests that would like to come, they are welcome, and seats, again, are still available. On the 26th of February, we'll have Charlie Cook join us looking forward into the election cycle that we're in. And to finish out the month, on the 27th of February, we'll have Eric Holder, Senior Counsel at Covington & Burling, and the 82nd Attorney General of the

United States.

In the spring, we'll be hosting David Ricks, the CEO of Eli Lilly. Christopher Waller, Member of the Board of Governors of the Federal Reserve. Professor Jeremy Siegel at Wharton School over at Penn in a conversation with our Chair, John Williams.

Chairwoman Jessica Rosenworcel of the Federal Communications Commission will be joining us. Susan Collins, the President and CEO of the Boston Fed. And Jamie Dimon, and that event is April, I want to say, the 23rd. And it's filling up, so please do, if you're interested, make reservations for that. And then our Club Chair, John Williams, will be doing an event for us as well this spring. And there's more actually that we're still in the process of confirming. So all the events are listed on the website. Please be sure to review the dates, add them to your calendar.

And, as always, we'd like to take a moment to thank our members of the Centennial Society joining us today as their contributions help to make our programming possible. Again, thank you all for attending. We'll see you again soon. For everybody digital, thank you for joining us. You are lucky you missed the snow. Thank you. Bye bye.