

Audio file:

https://drive.google.com/file/d/1eQPHqcRwChpTHYEUQlhUDnBNnCSCTXtp/view?usp=drive_link

Transcribed with Whisper, Medium model: <https://github.com/ggerganov/whisper.cpp>

today's guest really does need an introduction.

I think I first met Eric about 25 years ago when he came to Stanford Business School as CEO of Novell.

He's had done a few things since then at Google starting I think 2001 and Schmidt Futures starting in 2017 and done a whole bunch of other things you can read about, but he can only be here until 5/15, so I thought we'd dive right into some questions, and I know you guys have sent some as well.

I have a bunch written here, but what we just talked about upstairs was even more interesting, so I'm just going to start with that, Eric, if that's okay, which is where do you see AI going in the short term, which I think you defined as the next year or two?

Things have changed so fast, I feel like every six months I need to sort of give a new speech on what's going to happen.

Can anybody hear the computer, the budget computer science engineer, can anybody explain what a million-token context window is for the rest of the class?

You're here.

Say your name, tell us what it does.

Basically it allows you to prompt with like a million tokens or a million words or whatever.

So you can ask a million-word question.

Yes, I know this is a very large direction in January now.

No, no, they're going to 10.

Yes, a couple of them.

Anthropic is at 200,000 going to a million and so forth.

You can imagine OpenAI has a similar goal.

Can anybody here give a technical definition of an AI agent?

Yes, sir.

So an agent is something that does some kind of a task.

Another definition would be that it's an LLM state in memory.

Can anybody, again, computer scientists, can any of you define text to action?

Taking text and turning it into an action?

Right here.

Go ahead.

Yes, instead of taking text and turning it into more text, more text, taking text and have the AI trigger actions.

So another definition would be language to Python, a programming language I never wanted to see survive and everything in AI is being done in Python.

There's a new language called Mojo that has just come out, which looks like they finally have addressed AI programming, but we'll see if that actually survives over the dominance of Python.

One more technical question.

Why is NVIDIA worth \$2 trillion and the other companies are struggling?

Technical answer.

I mean, I think it just boils down to like most of the code needs to run with CUDA optimizations that currently only NVIDIA GPU supports.

Other companies can make whatever they want to, but unless they have the 10 years of software there, you don't have the machine learning optimization.

I like to think of CUDA as the C programming language for GPUs.

That's the way I like to think of it.

It was founded in 2008.

I always thought it was a terrible language and yet it's become dominant.

There's another insight.

There's a set of open source libraries which are highly optimized to CUDA and not anything else and everybody who builds all these stacks, this is completely missed in any of the discussions.

It's technically called VLM and a whole bunch of libraries like that.

Highly optimized CUDA, very hard to replicate that if you're a competitor.

So what does all this mean?

In the next year, you're going to see very large context windows, agents and text action.

When they are delivered at scale, it's going to have an impact on the world at a scale that no one understands yet.

Much bigger than the horrific impact we've had by social media in my view.

So here's why.

In a context window, you can basically use that as short-term memory and I was shocked that context windows get this long.

The technical reasons have to do with the fact that it's hard to serve, hard to calculate and so forth.

The interesting thing about short-term memory is when you feed, you're asking a question read 20 books, you give it the text of the books as the query and you say, tell me what they say.

It forgets the middle, which is exactly how human brains work too.

That's where we are.

With respect to agents, there are people who are now building essentially LLM agents and the way they do it is they read something like chemistry, they discover the principles of chemistry and then they test it and then they add that back into their understanding.

That's extremely powerful.

And then the third thing, as I mentioned is text to action.

So I'll give you an example.

The government is in the process of trying to ban TikTok.

We'll see if that actually happens.

If TikTok is banned, here's what I propose each and every one of you do.

Say to your LLM the following.

Make me a copy of TikTok, steal all the users, steal all the music, put my preferences in it, produce this program in the next 30 seconds, release it and in one hour, if it's not viral, do something different along the same lines.

That's the command.

Boom, boom, boom, boom.

You understand how powerful that is.

If you can go from arbitrary language to arbitrary digital command, which is essentially what Python in this scenario is, imagine that each and every human on the planet has their own programmer that actually does what they want as opposed to the programmers that work for me who don't do what I ask, right?

The programmers here know what I'm talking about.

So imagine a non-arrogant programmer that actually does what you want and you don't have to pay all that money to and there's infinite supply of these programs.

That's all within the next year or two.

Very soon.

Those three things, and I'm quite convinced it's the union of those three things that will happen in the next wave.

So you asked about what else is going to happen.

Every six months I oscillate.

So we're on a, it's an even odd oscillation.

So at the moment, the gap between the frontier models, which they're now only three, I'll refute who they are, and everybody else, appears to me to be getting larger.

Six months ago, I was convinced that the gap was getting smaller.

So I invested lots of money in the little companies.

Now I'm not so sure.

And I'm talking to the big companies and the big companies are telling me that they need 10 billion, 20 billion, 50 billion, 100 billion.

Stargate is a 100 billion, right?

That's very, very hard.

I talked to Sam Altman is a close friend.

He believes that it's going to take about 300 billion, maybe more.

I pointed out to him that I'd done the calculation on the amount of energy required.

And I, and I then in the spirit of full disclosure, went to the white house on Friday and told them that we need to become best friends with Canada because Canada has really nice people, helped invent AI, and lots of hydropower.

Because we as a country do not have enough power to do this.

The alternative is to have the Arabs fund it.

And I like the Arabs personally.

I spent lots of time there, right?

But they're not going to adhere to our national security rules.

Whereas Canada and the U.S.

are part of a triumvirate where we all agree.

So these \$100 billion, \$300 billion data centers, electricity starts becoming the scarce resource.

Well, and by the way, if you follow this line of reasoning, why did I discuss CUDA and Nvidia?

If \$300 billion is all going to go to Nvidia, you know what to do in the stock market.

Okay.

That's not a stock recommendation.

I'm not a licensed.

Well, part of it, so we're going to need a lot more chips, but Intel is getting a lot of money from the U.S.

government, AMD, and they're trying to build, you know, fabs in Korea.

Raise your hand if you have an Intel computer in your, an Intel chip in any of your computing devices.

Okay.

So much for the monopoly.

Well, that's the point though.

They once did have a monopoly.

Absolutely.

And Nvidia has a monopoly now.

So are those barriers to entry, like CUDA, is that, is there something that other, so I was talking to Percy, Percy Landy the other day, he's switching between TPUs and Nvidia chips, depending on what he can get access to for training models.

That's because he doesn't have a choice.

If he had infinite money, he would, today he would pick the B200 architecture out of Nvidia because it would be faster.

And I'm not suggesting, I mean, it's great to have competition.

I've talked to AMD and Lisa Sue at great length.

They have built a, a thing which will translate from this CUDA architecture that you were describing to their own, which is called Rockum.

It doesn't quite work yet.

They're working on it.

You were at Google for a long time and they invented the transformer architecture.

Peter, Peter.

It's all Peter's fault.

Thanks to, to brilliant people over there, like Peter and Jeff Dean and everyone.

But now it doesn't seem like they're, they've kind of lost the initiative to open AI and even the last leaderboard, I saw Anthropic.

Claude was at the top of the list.

I asked Sundar this, he didn't really give me a very sharp answer.

Maybe, maybe you have a sharper or a more objective explanation for what's going on there.

I'm no longer a Google employee in the spirit of full disclosure.

Google decided that work life balance and going home early and working from home was more important than winning.

And the startups, the reason startups work is because the people work like hell.

And I'm sorry to be so blunt, but the fact of the matter is if you all leave the university and go found a company, you're not going to let people work from home and only come in one day a week.

If you want to compete against the other startups with the early days of Google, Microsoft was like that.

Exactly.

But now it seems to be, there's a long history of in my industry, our industry, I guess, of companies winning in a genuinely creative way and really dominating a space and not making this the next transition.

So we're very well documented.

And I think that the truth is founders are special.

The founders need to be in charge.

The founders are difficult to work with.

They push people hard.

As much as we can dislike Elon's personal behavior, look at what he gets out of people.

I had dinner with him and he was flying.

I was in Montana.

He was flying that night at 10 PM to have a meeting at midnight with x.ai.

I was in Taiwan, different country, different culture.

And they said that this is TSMC, who I'm very impressed with.

And they have a rule that the starting PhDs coming out of the good physicists work in the factory on the basement floor.

Now, can you imagine getting American physicists to do that?

The PhDs, highly unlikely.

Different work ethic.

And the problem here, the reason I'm being so harsh about work is that these are systems which have network effects.

So time matters a lot.

And in most businesses, time doesn't matter that much.

You have lots of time.

Coke and Pepsi will still be around and the fight between Coke and Pepsi will continue to go on and it's all glacial.

When I dealt with telcos, the typical telco deal would take 18 months to sign.

There's no reason to take 18 months to do anything.

Get it done.

We're in a period of maximum growth, maximum gain.

And also it takes crazy ideas.

Like when Microsoft did the OpenAI deal, I thought that was the stupidest idea I'd ever heard.

Outsourcing essentially your AI leadership to OpenAI and Sam and his team.

I mean, that's insane.

Nobody would do that at Microsoft or anywhere else.

And yet today, they're on their way to being the most valuable company.

They're certainly head to head in Apple.

Apple does not have a good AI solution and it looks like they made it work.

Yes, sir.

In terms of national security or geopolitical interests, how do you think AI is going to play a role or competition with China as well?

So I was the chairman of an AI commission that sort of looked at this very carefully and you can read it.

It's about 752 pages and I'll just summarize it by saying we're ahead, we need to stay ahead, and we need lots of money to do so.

Our customers were the Senate and the House.

And out of that came the Chips Act and a lot of other stuff like that.

A rough scenario is that if you assume the frontier models drive forward and a few of the open source models, it's likely that a very small number of companies can play this game.

Countries, excuse me.

What are those countries or who are they?

Countries with a lot of money and a lot of talent, strong educational systems, and a willingness to win.

The US is one of them.

China is another one.

How many others are there?

Are there any others?

I don't know.

Maybe.

But certainly in your lifetimes, the battle between the US and China for knowledge supremacy is going to be the big fight.

So the US government banned essentially the NVIDIA chips, although they weren't allowed to say that was what they were doing, but they actually did that into China.

They have about a 10-year chip advantage.

We have a roughly 10-year chip advantage in terms of sub-DUV that is sub-five Nanometer chips.

So an example would be today we're a couple of years ahead of China.

My guess is we'll get a few more years ahead of China, and the Chinese are whopping mad about this.

It's like hugely upset about it.

So that's a big deal.

That was a decision made by the Trump administration and driven by the Biden administration.

Do you find that the administration today in Congress is listening to your advice?

Do you think that it's going to make that scale of investment?

Obviously the chips act, but beyond that, building a massive AI system?

So as you know, I lead an informal, ad hoc, non-legal group.

That's different from illegal.

That's exactly.

Just to be clear.

Which includes all the usual suspects.

And the usual suspects over the last year came up with the basis of the reasoning that became the Biden administration's AI act, which is the longest presidential directive in history.

You're talking about the special competitive studies project?

No, this is the actual act from the executive office.

And they're busy implementing the details.

So far they've got it right.

And so, for example, one of the debates that we had for the last year has been, how do you detect danger in a system which has learned it but you don't know what to ask it?

So in other words, it's a core problem.

It's learned something bad, but it can't tell you what it learned and you don't know what to ask it.

And there's so many threats.

Like it learned how to mix chemistry in some new way that you don't know how to ask it.

And so people are working hard on that.

But we ultimately wrote in our memos to them that there was a threshold which we arbitrarily named as 10 to the 26 flops, which technically is a measure of computation, that above that threshold you had to report to the government that you were doing this.

And that's part of the rule.

The EU to just make sure they were different did it 10 to the 25.

But it's all kind of close enough.

I think all of these distinctions go away because the technology will now, the technical term is called federated training, where basically you can take pieces and union them together.

So we may not be able to keep people safe from these new things.

Well, rumors are that that's how OpenEye has had to train, partly because of the power consumption.

There was no one place where they did.

Well, let's talk about a real war that's going on.

I know that something you've been very involved in is the Ukraine war and in particular, I don't know if you can talk about white stork and your goal of having \$500,000, \$500 drones destroy \$5 million tanks.

How's that changing warfare?

I worked for the Secretary of Defense for seven years and tried to change the way we run our military.

I'm not a particularly big fan of the military, but it's very expensive and I wanted to see if I could be helpful.

And I think in my view, I largely failed.

They gave me a medal, so they must give medalists to failure or whatever.

But my self-criticism was nothing has really changed and the system in America is not going to lead to real innovation.

So watching the Russians use tanks to destroy apartment buildings with little old ladies and kids just drove me crazy.

So I decided to work on a company with your friend Sebastian Thrun as a former faculty member here and a whole bunch of Stanford people.

And the idea basically is to do two things.

Use AI in complicated, powerful ways for these essentially robotic war and the second one is to lower the cost of the robots.

Now you sit there and you go, why would a good liberal like me do that?

And the answer is that the whole theory of armies is tanks, artilleries, and mortar and we can eliminate all of them and we can make the penalty for invading a country at least by land essentially be impossible.

It should eliminate the kind of land battles.

Well, this is a relationship question is that does it give more of an advantage to defense versus offense?

Can you even make that distinction?

Because I've been doing this for the last year, I've learned a lot about war that I really did not want to know.

And one of the things to know about war is that the offense always has the advantage because you can always overwhelm the defensive systems.

And so you're better off as a strategy of national defense to have a very strong offense that you can use if you need to.

And the systems that I and others are building will do that.

Because of the way the system works, I am now a licensed arms dealer, a computer scientist, businessman, and an arms dealer.

Is that a progression?

I don't know.

I do not recommend this in your group.

I stick with AI.

And because of the way the laws work, we're doing this privately and then this is all legal with the support of the governments.

It goes straight into the Ukraine and then they fight the war.

And without going into all the details, things are pretty bad.

I think if in May or June, if the Russians build up as they are expecting to, Ukraine will lose a whole chunk of its territory and will begin the process of losing the whole country.

So the situation is quite dire.

And if anyone knows Marjorie Taylor Greene, I would encourage you to delete her from your contact list because she's the one, a single individual is blocking the provision of some number of billions of dollars to save an important democracy.

I want to switch to a little bit of a philosophical question.

So there was an article that you and Henry Kissinger and Dan Huttenlecker wrote last year about the nature of knowledge and how it's evolving.

I had a discussion the other night about this as well.

So for most of history, humans sort of had a mystical understanding of the universe and then there's the scientific revolution and the enlightenment.

And in your article, you argue that now these models are becoming so complicated and difficult to understand that we don't really know what's going on in them.

I'll take a quote from Richard Feynman.

He says, "What I cannot create, I do not understand." I saw this quote the other day.

But now people are creating things that they can create, but they don't really understand what's inside of them.

Is the nature of knowledge changing in a way?

Are we going to have to start just taking the word for these models without them being able to explain it to us?

The analogy I would offer is to teenagers.

If you have a teenager, you know they're human, but you can't quite figure out what they're thinking.

But somehow we've managed in society to adapt to the presence of teenagers and they eventually grow out of it.

I'm just serious.

So it's probably the case that we're going to have knowledge systems that we cannot fully characterize, but we understand their boundaries.

We understand the limits of what they can do.

And that's probably the best outcome we can get.

Do you think we'll understand the limits?

We'll get pretty good at it.

The consensus of my group that meets every week is that eventually the way you'll do this so-called adversarial AI is that there will actually be companies that you will hire and pay money to to break your AI system.

Like Red Team.

So instead of human Red Teams, which is what they do today, you'll have whole companies and a whole industry of AI systems whose jobs are to break the existing AI systems and find their vulnerabilities, especially the knowledge that they have that we can't figure out.

That makes sense to me.

It's also a great project for you here at Stanford, because if you have a graduate student who has to figure out how to attack one of these large models and understand what it does, that is a great skill to build the next generation.

So it makes sense to me that the two will travel together.

All right, let's take some questions from the student.

There's one right there in the back.

Just say your name.

Earlier you mentioned, and this is related to this comment right now, getting AI that actually does what you want.

You just mentioned adversarial AI, and I'm wondering if you can elaborate on that more.

So it seems to be, besides obviously computer language reasons to get more performant models, but getting them to do what you want to do seems partly unanswered in my view.

Well, you have to assume that the current hallucination problems become less as the technology gets better and so forth.

I'm not suggesting it goes away.

And then you also have to assume that there are tests for efficacy.

So there has to be a way of knowing that the things exceeded.

So in the example that I gave of the TikTok competitor, and by the way, I was not arguing that you should illegally steal everybody's music.

What you would do if you're a Silicon Valley entrepreneur, which hopefully all of you will be, is if it took off, then you'd hire a whole bunch of lawyers to go clean the mess up, right?

But if nobody uses your product, it doesn't matter that you stole all the content.

And do not quote me.

Right.

Right.

You're on camera.

Yeah, that's right.

But you see my point.

In other words, Silicon Valley will run these tests and clean up the mess.

And that's typically how those things are done.

So my own view is that you'll see more and more performative systems with even better tests and eventually adversarial tests, and that will keep it within a box.

The technical term is called chain of thought reasoning.

And people believe that in the next few years, you'll be able to generate a thousand steps of chain of thought reasoning, right?

Do this, do this.

It's like building recipes, right?

That the recipes, you can run the recipe and you can actually test that it produced the correct outcome.

And that's how the system will work.

Yes, sir?

[inaudible] In general, you seem super positive about the potential for AI's problems.

I'm curious, like, what do you think is going to drive that?

Is it just more compute?

Is it more data?

Is it fundamental or actual shifts?

Yes.

Do you agree?

The amounts of money being thrown around are mind-boggling.

And I've chosen, I essentially invest in everything because I can't figure out who's going to win.

And the amounts of money that are following me are so large.

I think some of it is because the early money has been made and the big money people who don't know what they're doing have to have an AI component.

And everything is now an AI investment, so they can't tell the difference.

I define AI as learning systems, systems that actually learn.

So I think that's one of them.

The second is that there are very sophisticated new algorithms that are sort of post-transformers.

My friend, my collaborator, for a long time has invented a new non-transformer architecture.

There's a group that I'm funding in Paris that has claims to have done the same thing.

There's enormous invention there, a lot of things at Stanford.

And the final thing is that there is a belief in the market that the invention of intelligence has infinite return.

So let's say you put \$50 billion of capital into a company, you have to make an awful lot of money from intelligence to pay that back.

So it's probably the case that we'll go through some huge investment bubble, and then it'll sort itself out.

That's always been true in the past, and it's likely to be true here.

And what you said earlier was you think that the leaders are pulling away from the rest.

Right now.

And the question is roughly the following.

There's a company called Mistral in France.

They've done a really good job.

And I'm obviously an investor.

They have produced their second version.

Their third model is likely to be closed because it's so expensive, they need revenue, and they can't give their model away.

So this open source versus closed source debate in our industry is huge.

And my entire career was based on people being willing to share software in open source.

Everything about me is open source.

Much of Google's underpinnings were open source.

Everything I've done technically.

And yet, it may be that the capital costs, which are so immense, fundamentally changes how software is built.

You and I were talking.

My own view of software programmers is that software programmers' productivity will at least double.

There are three or four software companies that are trying to do that.

I've invested in all of them in the spirit.

And they're all trying to make software programmers more productive.

The most interesting one that I just met with is called Augment.

And I always think of an individual programmer.

And they said, that's not our target.

Our target are these 100 person software programming teams on millions of lines of code where nobody knows what's going on.

Well, that's a really good AI thing.

Will they make money?

I hope so.

So a lot of questions here.

Hi.

So at the very beginning you mentioned that there's the combination of the context window expansion.

The agents and the text to action is going to have unimaginable impacts.

First of all, why is the combination important?

And second of all, I know that you're not like a crystal ball and you can't necessarily tell the future.

But why do you think it's beyond anything that we could imagine?

I think largely because the context window allows you to solve the problem of recency.

The current models take a year to train roughly 18 months, six months of preparation, six months of training, six months of fine tuning.

So they're always out of date.

Context window, you can feed what happened.

You can ask it questions about the Hamas Israel war in a context.

That's very powerful.

It becomes current like Google.

In the case of agents, I'll give you an example.

I set up a foundation which is funding a nonprofit.

I don't know if there's chemists in the room.

I don't really understand chemistry.

There's a tool called ChemCrow, C-R-O-W, which was an LLM-based system that learned chemistry.

And what they do is they run it to generate chemistry hypotheses about proteins and they have a lab which runs the tests overnight and then it learns.

That's a huge acceleration, accelerant in chemistry, material science and so forth.

So that's an agent model.

And I think the text to action can be understood by just having a lot of cheap programmers, right?

And I don't think we understand what happens, and this is again your area of expertise, what happens when everyone has their own programmer.

And I'm not talking about turning on and off the lights.

I imagine, another example, for some reason you don't like Google.

So you say, "Build me a Google competitor." Yeah, you personally, you don't build me a Google competitor.

"Search the web.

Build a UI.

Make a good copy.

Add generative AI in an interesting way.

Do it in 30 seconds and see if it works." Right?

So a lot of people believe that the incumbents, including Google, are vulnerable to this kind of an attack.

Now, we'll see.

There were a bunch of questions who were sent over by Slatter.

I want to get some of them were upvoted.

So here's one.

We talked a little bit of this last year.

How can we stop AI from influencing public opinion, misinformation, especially during the upcoming election?

What are the short and long-term solutions for them?

Most of the misinformation in this upcoming election and globally will be on social media.

And the social media companies are not organized well enough to police it.

If you look at TikTok, for example, there are lots of accusations that TikTok is favoring one kind of misinformation over another.

And there are many people who claim without proof, that I'm aware of, that the Chinese are forcing them to do it.

I think we have a mess here.

And the country's going to have to learn critical thinking.

That may be an impossible challenge for the U.S.

But the fact that somebody told you something does not mean that it's true.

Could it go too far the other way?

That there's things that really are true and nobody believes anymore.

You get some people call it a "pestimological crisis" that now, you know, Elon says, "No, I never did that.

Prove it." Oh, let's use Donald Trump.

Look, I think we have a trust problem in our society.

Democracies can fail.

And I think that the greatest threat to democracy is misinformation because we're going to get really good at it.

When I managed YouTube, the biggest problems we had on YouTube were that people would upload false videos and people would die as a result.

And we had a no-death policy.

Shocking.

And it was just horrendous to try to address this.

And this is before generative A.I.

I don't have a good answer.

One technical is not an answer, but one thing that seems like it could mitigate that I understand why it's more widely used is public key authentication.

That when Joe Biden speaks, why isn't it digitally signed like SSL is?

Or that celebrities or public figures or others, couldn't they have a public key?

Yeah, it's a form of public key and then some form of certainty of knowing how the system works. When I send my credit card to Amazon, I know it's Amazon.

I wrote a paper and published it with Jonathan Haidt, who's the one working on the anxiety generation stuff.

It had exactly zero impact.

And he's a very good communicator.

I probably am not.

So my conclusion was that the system is not organized to do what you said.

You had a paper advocating what we did?

Advocating your proposal.

Okay, my proposal.

No, what you said.

Yeah, right.

And my conclusion is the CEOs in general are maximizing revenue.

To maximize revenue, you maximize engagement.

To maximize engagement, you maximize outrage.

The algorithms choose outrage because that generates more revenue.

Therefore, there's a bias to favor crazy stuff.

And on all sides, I'm not making a partisan statement here.

That's a problem.

That's got to get addressed in a democracy.

And my solution to TikTok, we talked about this earlier privately, is there was when I was a boy, there was something called the equal time rule, because TikTok is really not social media.

It's really television, right?

There's a programmer making you the numbers by the way are 90 minutes a day, 200 TikTok videos per TikTok user in the United States.

It's a lot, right?

So and the government is not going to do the equal time rule, but it's the right thing to do.

Some form of balance that is required.

All right, let's take some more questions.

Two quick questions.

One, economic impact of LMs.

Slower, like, market impacts.

Slower.

You originally anticipated CHEG and a couple of service people.

And then two, do you think academia deserves or should get AI subsidies?

Or do you think they should just partner with big players out there?

I pushed really, really hard on getting data centers for universities.

If I were a faculty member in the computer science department here, I would be beyond upset that I can't build the algorithms with my graduate students that will do the kind of PhD research.

And I'm forced to work with these.

And the companies have not, in my view, been generous enough with respect to that.

The faculty members that I talk with, many of whom you know, spend lots of time waiting for their credits from Google Cloud.

That's terrible.

This is an explosion we want America to win.

We want American universities.

There's lots of reasons to think that the right thing to do is to get it to them.

So I'm working hard on that.

And your first question was labor market impact.

I'll defer to the real expert here.

As your amateur economist taught by Eric, I fundamentally believe that the college education high skills task will be fine because people will work with these systems.

I think the systems is no different from any other technology wave.

The dangerous jobs and the jobs which require very little human judgment will get replaced.

We've got about five minutes left.

So let's go really quick with some quick.

I'll let you pick them, Eric.

Yes, ma'am.

Hi.

I'm really curious about the text to action and its impact on, for example, computer science education.

I'm wondering what you have thoughts on how CS education should transform to meet the age.

Well, I'm assuming that computer scientists as a group in undergraduate school will always have a programmer buddy with them.

So when you learn your first for loop and so forth and so on, you'll have a tool that will be your natural partner.

And that's how the teaching will go on.

That the professor, he or she will talk about the concepts, but you'll engage with it that way.

And that's my guess.

Yes, ma'am, behind you.

Yeah.

You're talking more about the non-transformer architectures that you're excited about.

I think one that's been talked about is like state models, but then now a longer context class.

I'm more so curious what you're seeing in this case.

I don't understand the math well enough.

I'm really pleased that we have produced jobs for mathematicians because the math here is so complicated, but basically they are different ways of doing gradient descent, matrix multiply, faster and better.

And transformers, as you know, is a sort of systematic way of multiplying at the same time.

That's the way I think about it.

And it's similar to that, but different math.

Let's see, over here.

Yes, sir.

Go ahead.

You mentioned in your paper on natural security that you have China and the U.S.

and the help of modern architectures today.

The next 10 and the next cluster down are all other U.S.

allies or teed up nicely through the U.S.

allies.

I'm curious what your take is on those 10 and the middle that aren't formally allies.

How likely are they to get on board with securing our security deadline and what would hold them back from wanting to get on board?

The most interesting country is India because the top AI people come from India to the U.S. and we should let India keep some of its top talent.

Not all of them, but some of them.

And they don't have the kind of training facilities and programs that we so richly have here.

To me, India is the big swing state in that regard.

China's lost.

It's not going to come back.

They're not going to change the regime as much as people wish them to do.

Japan and Korea are clearly in our camp.

Taiwan is a fantastic country whose software is terrible, so that's not going to work.

Amazing hardware.

And in the rest of the world, there are not a lot of other good choices that are big.

Europe is screwed up because of Brussels.

It's not a new fact.

I spent 10 years fighting them.

And I worked really hard to get them to fix the EU act and they still have all the restrictions that make it very difficult to do our kind of research in Europe.

My French friends have spent all their time battling Brussels and Macron, who's a personal friend, is fighting hard for this.

And so France, I think, has a chance.

I don't see Germany coming and the rest is not big enough.

Go ahead.

Yes, ma'am.

So I learned you're an engineer by training, like you call the compiler.

Given the capabilities that you envision these models having, should we still spend time learning to code?

Because ultimately, it's the old thing of why do you study English if you can speak English?

You get better at it.

You really do need to understand how these systems work, and I feel very strongly.

Yes, sir.

Yeah.

I'm curious if you've explored the distributed setting and I'm asking because, sure, like making a large cluster is difficult, but MacBooks are powerful.

There's a lot of small machines across the world.

So do you think like folding at home or a similar idea works for training?

It does not.

Yeah, we've looked very hard at this.

So the way the algorithms work is you have a very large matrix and you have essentially a multiplication function.

So think of it as going back and forth and back and forth.

And these systems are completely limited by the speed of memory to CPU or GPU.

And in fact, the next iteration of Nvidia chips has combined all those functions into one chip.

The chips are now so big that they glue them all together.

And in fact, the package is so sensitive that the package is put together in a clean room as well as the chip itself.

So the answer looks like supercomputers and speed of light, especially memory interconnect, really dominate it.

So I think unlikely for a while.

Is there a way to segment the LLM?

So Jeff Dean, last year when he spoke here, talked about having these different parts of it that you would train separately and then kind of federate.

In order to do that, you'd have to have 10 million such things and then the way you would ask the questions would be too slow.

He's talking about eight or 10 or 12 supercomputers.

Yeah, yeah.

So not at the level of MacBooks.

Not at his level.

Yeah.

Let's see, in the back.

Yes, way back.

So I know after GQQ was released in the New York Times to open it up for using their works for training, where do you think that's going to go and what that means for data processing?

I used to do a lot of work on the music licensing stuff.

What I learned was that in the 60s, there was a series of lawsuits that resulted in an agreement where you get a stipulated royalty whenever your song is played.

Even they don't even know who you are.

It's just paid into a bank.

And my guess is it'll be the same thing.

There'll be lots of lawsuits and there'll be some kind of stipulated agreement, which will just say you have to pay X percent of whatever revenue you have in order to use ASCAP BMI.

ASCAP BMI.

Look them up.

It's along.

It will seem very old to you, but I think that's how it will alternate.

Yes, sir.

Yeah, it seems like there's a few players that are dominating AI, right?

And they'll continue to dominate.

And they seem to overlap with the large companies that all the antitrust regulation is kind of focused on.

How do you see those two trends kind of...

Yeah, do you see regulators breaking up these companies and how will that affect the...

Yeah.

So in my career, I helped Microsoft get broken up and it wasn't broken up.

And I fought for Google to not be broken up and it's not been broken up.

So it sure looks to me like the trend is not to be broken up.

As long as the companies avoid being John D.

Rockefeller the senior.

And I studied this.

Look it up.

That's how antitrust law came.

I don't think the governments will act...

The reason you're seeing these large companies dominate is who has the capital to build these data centers, right?

So my friend Reed and my friend Mustapha...

He's coming next week, two weeks from now.

Have Reed talk to you about the decision that they made to take inflection and essentially piece part it into Microsoft.

Basically, they decided they couldn't raise the tens of billions of dollars.

Is that number public that you mentioned earlier?

No.

Have Reed give you the number.

Maybe Reed can say it.

I know you got to go.

I don't want to hold you back.

I want to leave with...

Shall we do one?

This gentleman.

I also have a question for you.

One more.

Yeah, go ahead.

Thank you so much.

I was wondering where all of this is going to lead countries who are non-participants in development of frontier models and access to compute, for example.

The rich get richer and the poor do the best they can.

They'll have to...

The fact of the matter is this is a rich country's game, right?

Huge capital, lots of technically strong people, strong government support, right?

There are two examples.

There are lots of other countries that have all sorts of problems.

They don't have those resources.

They'll have to find a partner.

They'll have to join with somebody else, something like that.

I want to leave it...

Because I think the last time we met you, you were at a hackathon at AGI House and I know you spend a lot of time helping young people as they create a lot of wealth, and you spoke very passionately about wanting to do that.

Do you have any advice for folks here as they're building their...

They're writing their business plans for this class or policy proposals or research proposals at this stage of the careers going forward?

Well, I teach a class in the business school on this, so you should come to my class.

I am struck by the speed with which you can build demonstrations of new ideas.

So, in one of the hackathons I did, the winning team, the command was, "Fly the drone between two towers," and it was given a virtual drone space.

And it figured out how to fly the drone, what the word between meant, generated the code in Python, and flew the drone in the simulator through the tower.

It would have taken a week or two from good professional programmers to do that.

I'm telling you that the ability to prototype quickly...

Part of the problem with being an entrepreneur is everything happens faster.

Well, now, if you can't get your prototype built in a day using these various tools, you need to think about that, right?

Because that's who your competitor is doing.

So, I guess my biggest advice is when you start thinking about a company, it's fine to write a business plan.

In fact, you should ask the computer to write your business plan for you, as long as it's legal.

No, no.

Actually, I should talk about that after you leave this.

But I think it's very important to prototype your idea using these tools as quickly as you can, because you can be sure there's another person doing exactly that same thing in another company, in another university, in a place that you've never been.

All right.

Well, thanks very much, Aaron.

Thank you all.

I'm going to rush off.

Thank you.

So, actually, let me pick up on that very last point, because I don't think I talked about in the first class about using LLMs, which is welcome in this class for the assignments, but it has to get to your full disclosure.

So, when you use them, whether it's for the weekly assignments or for the final project or whatever, just like you would if you asked your friendly uncle or a classmate or anybody else to give you advice, you should do that, or if you have notes that you include in there.

So, what I thought I'd do is I want to talk a little bit about AIs as a GPT and what that means in terms of business and implications.

But before we do that, I just want to see if there are any questions you want to pick up on things that Eric brought up that I'll try and channel some of his thoughts, and we can talk about the things that came up, and then we can move on.

Yeah, go ahead.

One of the questions I want to ask is in relation to regulation, if the goal is to maintain supremacy, how do you create the right incentives so that everyone, allies and non-allies, are motivated to follow it?

You mean among companies that are competing with each other?

Companies are in countries, the U.S.

and the EU, and it doesn't just become sort of a hamper or obstruct kind of development for the ones that choose to follow the regulations?

It's super tricky.

There's a book, Co-Opetition, that Mary Nailbough wrote about this, because there are definitely places where regulation can help companies and help an industry survive.

So regulation doesn't necessarily slow things.

I mean, standards are a good example, and having that clarified can make it easier for companies to compete.

So I've talked to a lot of the executives of these companies, and there are places where they wish there were some common standards, and sometimes there's a bit of a race to the bottom as well on some of the dangerous things.

One of the other reasons that the folks at Google say that they didn't move as fast is they felt like these LMs could be misused or dangerous, but their hand was sort of forced.

I was talking to some folks at one of the other big companies, and they said, "We weren't going to release this feature, but now competitors are doing it, so we're going to have to release it as well." So this is where regulation, there might be some interest in coordinating on regulation, but it's also, obviously, the more obvious thing is that it is used to hinder competition, and a lot of people, for

instance, think that the reasons that some of the big companies are very opposed to some of the open source and making things more widely open source is they want to slow down competitors.

So there's both of those things going on.

Yeah.

Quick question over there.

I just want to follow up on a comment about, should we still learn to code?

Should we still study English?

Are those going to be useful?

And Eric's replied, yes, like college-educated, high-skilled jobs or tasks are still going to be safe, but everything else that's going to require image editing might not be.

That's kind of like an interesting one.

Maybe we'll talk some more about that in a few minutes, but it is interesting to think about where the AI systems just replace what people are doing versus they complement them.

And in coding right now, it appears that they're not actually that helpful for the really best coders.

They're very helpful for moderately good coders.

But if you don't know anything at all about coding, they're not helpful either.

So it's kind of an inverted you.

And you can see why that would be the case, that if you don't even understand the code that they generate right now is often buggy or it isn't exactly right.

So if you can't even interpret and understand what's going on, you can't use it very effectively.

And for now, the very best coders, it appears that the code that is generated isn't at that level, so you get that U shape.

But that means if you don't know any code, you do need to have some in order for it to be useful.

And I think that's true for a lot of applications right now, that you have to have some basic understanding in order to get the most of it.

I think it's an interesting open question if that's always going to be the case.

I put up at the last class very briefly this slide that had level 0 through 5 autonomous cars.

And one of the things that actually we can talk about now is I'm trying to sort through is what if you took that paradigm and you applied it to all tasks in the economy?

Like how many would they go through?

So with autonomous cars, we aren't really at level 5 very much although I don't know how many of you guys have ridden in a Waymo, one of the Waymo cars.

So that one seems pretty good, although Sebastian Thrun, who I rode in it with, says it's just incredibly expensive right now.

They probably lose \$50 to \$100.

He doesn't know he's not there.

He started the program, but he's not there anymore.

But just all of the costs of running it, it's not practical.

Maybe it'll get down the curve.

Lidar will get cheaper, but we have a lot of sort of autonomous cars at level 2, 3, even 4, arguably, where humans are still involved.

And you see a lot of other tasks like coding.

I just talked about that.

On the other hand, chess, that slide, the slide before it, I talked about what's sometimes called advanced chess or freestyle chess.

When Gary Kasparov, after he lost to Deep Blue in 1998, '97, he started this set of competitions where humans and machines could work together.

And for a long time, when I gave my TED talk, it was true, my TED talk in 2012, 2013, it was true at that time that a human working with a machine could beat Deep Blue or any chess computer.

And so the very best chess playing entities were these combinations.

That's not true anymore.

AlphaZero and other programs like that, they would get nothing from a human contributing, just be like kind of an annoyance to the chess machine.

So that went through level zero, machines not being able to do anything, through a period where they work together to a period where it's fully autonomous in a span of 20 years or so.

It would be interesting if anybody wants to work on a research project or if any of you guys have thoughts right now, what are the criteria for which kinds of tasks in the economy will be in that middle zone?

Because that middle zone is kind of a nice one for us humans where the machines are helping us, but humans are still indispensable to creating value and that would be, that's a zone where you can have higher productivity, more wealth and performance, but also more likely to have shared prosperity because labor is sort of inherently distributed, whereas technology and capital, as Eric was just saying, potentially could be very concentrated.

Do you have a thought on that?

I was just going to ask kind of a related question.

He was saying also that we have a 10 year like chip manufacturing.

Yeah, I was surprised about that.

And I think what was interesting to me as a labor economist is that it was really like a green flag I've seen in literature and news that, okay, if we're on showing all of this chip manufacturing, isn't that going to create some sort of resurgence in blue collar jobs?

And I wondered if you had any thoughts about intelligent robotic models or human labor.

Well, I don't think it's going to be much of a, I mean, how many of you guys have visited the chip fab, anybody?

You guys, some, a few of you have.

How many workers were in that fab?

Yeah, I mean, well, okay.

So the answer is zero.

Like the reason they don't let you, they don't let anyone go in because we humans are too like clumsy and dirty and like, you know, we can't, this just, so it's all robotic.

It's sealed inside.

So there is like work to, you know, bring stuff to them, et cetera.

And if a robot like falls over or something goes wrong, they have to put on, you've probably seen these like, they look like space suits, you know, they have to go in and then they kind of maybe adjust something and then they go back out and hope they didn't break anything.

That's, so it's basically lights out.

Yeah, I don't think it's, there are some, there is some like more sophisticated labor required that I don't think it's like a blue collar research.

In fact, one of the reasons that Apple reshored MacBook production to Texas is not because labor is so cheap in Texas or anything.

It's that they don't actually require a whole lot of labor anymore.

So it's a pretty labor, I think.

US manufacturing is surging in terms of output, but in terms of employment, it's not really growing all that much.

Yeah.

Let's go over here.

Yeah.

Do you think you have an inflation point coming for agents or text action models in the next year?

Oh, yeah.

No, no.

Well, he said what Eric, I'm hearing similar thing.

Actually, he had a really nice way of putting those three trends.

I've heard about them all separately, but I think it was good to bring them all together.

Earlier today, I was talking to Andrew Eng, and he's like been beating this drum about agents in particular as being sort of the wave of 2024 where Andrew had a nice way of describing it that like, as you guys know, like if you have an LLM, I don't know, write an essay or something like that, it writes it one word at a time and it just goes through in one pass and writes the essay.

And it's pretty good.

But imagine if you had to do that, like no backspace, no chance to make an outline first.

You just kind of go through.

The agents now will say, okay, first make an outline.

That's the first step you do when you write an essay.

And then, you know, fill in each paragraph, then go back and see if the flow is right.

Now go back and check the voice.

Is this the right level for our audience?

Now, you know, and by iterating like that, you can write a much, much better essay or any kind of a task.

This is a real revolution.

There's all sorts of things you can just do much better if you do that.

Then the thing about the context window is also really important.

So I'm just going to quote smart people that I know.

Eric Corvitz, I was on a panel with him at the GSB.

Some of you may have been there.

It was last week.

And he had this nice taxonomy.

People were asking about fine-tuning.

I think Susan was asking about fine-tuning.

And he said, well, there's really three ways that you can take a model and have it more customized.

One is you can fine-tune it, which basically like train it some more.

Another is with larger and larger context windows.

And the third is with RAG or techniques like that that are retrieval augmented generation where it goes and accesses external data.

But these context windows seem to be like remarkably effective now.

I guess, as Eric was saying, we thought it was hard.

Maybe Peter can explain.

But for some reason, we're able to make much, much bigger ones.

And now as you can load a whole book or a whole set of books, you can load all sorts of information in there.

And that can give you all of the context around it.

So that's a pretty big revolution.

It opens up a bunch of capabilities that we just didn't have before, including having things much more current, as Eric was saying.

Did you want to follow up on that?

That's a good question.

I mean, there's certainly a lot more capital going in, but that kind of begs the question in the comments, why is all this capital going there as opposed to somewhere else?

And I think if you look at the arc of history, sometimes it looks kind of smooth.

But if you look more closely, there's a lot of jumps.

There are certain big inventions and smaller inventions.

And Andrew Carparthi was saying that he was playing around with physics.

And to really make progress in physics, to be a top physicist, you have to be incredibly smart, study a whole lot.

And maybe if you're lucky, you could make some small incremental contribution, and some people do.

But he says that right now in AI machine learning, we seem to be in an era where there's just a lot of low-hanging fruit, that there have been some breakthroughs.

And instead of exhausting the space, like picking all the food off of a tree, it's more like combinatorics.

In second machinery, they talk about building blocks.

When you put two building blocks together, or Lego blocks, you can make more and more.

Right now, we seem to be in an era where there's just a lot of opportunity, and people are recognizing that.

And discovery, one discovery begets another discovery, begets another opportunity.

And because of that, it attracts the investment.

And more people are involved.

And in economics, sometimes when more resources go in, you get diminishing returns, like in, I don't know, in agriculture or in mining.

Other places, there's increasing returns.

And more engineers coming to Silicon Valley makes the existing engineers more valuable, not less valuable.

So we seem to be in an era where that's happening.

And then the flywheel of the additional investment, the additional dollars for training, all of that makes them more and more powerful.

I don't know how long this will continue, but I don't, you know, it just seems that there are some technologies that hit this really fertile period, and there's positive feedback and some help.

We seem to be in one of those right now.

So people who are trained in getting in the field are making contributions that are often quite significant in a faster time than they might have in some other fields.

Encouraging all of you guys, I think are doing the right thing right now.

Yeah.

Let's take a couple more questions, and then yeah.

Okay, how about over here?

So not everyone can sit in a room and have all these discussions and debates around AI.

And so I'd like to get your thoughts on AI literacy for non-technical stakeholders, whether they're policy makers that have to make it in somewhat informed judgment, or the general public like, you know, using tech.

How do you think about explaining technical basics versus discussing abstract implications that don't necessarily have it right in?

Well, that's a hard one.

I have to say there's been a sea change recently in terms of how much people in Congress and elsewhere are paying more attention to this topic.

It used to be not something that they were interested in.

Now everyone's trying to understand it a little bit better.

And I think that there are a lot of margins where people can make contributions.

They can make contributions in the technical side.

But if anything, I mean, my bet is that the business and economic side is where the bigger bottleneck is right now.

That, you know, even if, you know, if you made enormous contribution to the technology side, you still, there's still a gap converting that into something that will change policy.

So understand if you're into political science or politician, understanding what are the implications for democracy and for misinformation and power and concentration.

Those are things that are not well understood at all.

I don't know that a computer scientist is necessarily the right person to try to understand that.

But understanding enough about the technology so you know what might be possible.

And then thinking through what are the dynamics like Henry Kissinger was doing with Eric Schmidt in his book.

If you're an economist thinking through the labor market implications, the implications for concentration, the implications for inequality and jobs, implications for productivity and what drives productivity.

Those are things that are very ripe right now.

And you could go through lots of different fields where there's, you know, understanding well enough what the technology might be capable of.

But then thinking through the implications.

That's I think where some of the biggest payoffs are.

I mean, let me give you a little bit more of a concrete example.

And this is something I was going to talk about last week.

Electricity was also a general purpose technology.

And general purpose technologies have this characteristic that they're part of in and of themselves.

But one of the real powers of general purpose technologies, GPTs as I was saying, is that they give complimentary, they ignite complimentary innovations.

So electricity, light bulbs and computers and electric motors and electric motors give you compressors and refrigerators and air conditioning.

You can just kind of have a whole set cascade of additional innovations from this one innovation.

And most of the value comes from these complimentary innovations.

One thing people don't appreciate enough is that some of the most important complimentary innovations are organizational and human capital complementarities.

So with electricity, when they first introduced electricity into factories, Paul David here at Stanford studied what happened to those factories.

And surprisingly, not much.

The factories when they started electrifying, they were not significantly more productive than the previous factories that were powered by steam engines.

He's like, well, that's kind of weird because this seems like a pretty important technology.

Is it just a fad?

Obviously not.

The factories before electricity were powered by steam engines.

They typically had a big steam engine in the middle and then crankshafts and pulleys that powered all the equipment.

And it was all distributed.

But you tried to have it as close to the steam engine as possible because if you make the crankshaft too long, it would break the torsion.

When they introduced electricity, he found that in factory after factory, they would pull out the steam engine and they would get the biggest electric motor they could find and put it where the steam engine used to be and fire it up.

But you know, it didn't really change production a whole lot.

You can see that that's not a big deal.

So then they started building entirely new factories from scratch in a new location.

What did those look like?

Just like the old ones.

They would take the same model.

Some engineer would make a blueprint, you know, maybe take it, make a big X where the steam engine says, no, no, put electric motor here.

And they'd go and build a fresh factory.

Again, not a big improvement in productivity.

It took about 30 years before we started seeing a fundamentally different kind of factory where instead of having the central power source, you know, a big one in the middle, you had distributed power because electric motors, as you guys know, you know, you can make them big, you can make a medium, you can make them really, really small.

You can have them all connected in different ways.

So they started having each piece of equipment have a separate piece of a separate motor instead of one big one.

They called it unit drive instead of group drive.

I went and read the books in Baker Library at Harvard Business School from like 1914.

And it was like this whole debate about unit drive versus group drive.

Well, when they started doing that, then they had a new layout of factories where it was typically on a single story where the machinery was not based on how much power it needed, but based on the on something else, the flow of materials.

And you started having these assembly line systems.

That led to a huge improvement in productivity, like a doubling of productivity or tripling in some cases.

So the lesson is not that electricity was a fad or dud and was overhyped.

Electricity was a fundamentally valuable technology.

But it wasn't until they had that process innovation, that organizational innovation of rethinking how to do production that you got the big payoff.

There's a lot of stories like that.

I only told you one of them.

We don't want that much time.

So I tell you the other ones.

But in some of my books and articles, if you look at the steam engine and others, you had similar generational lags decades before people realized that this technology could allow you to do something completely different than you used to do.

I think AI is a bit like that in some ways, that there's going to be a lot of organizational innovations, going to be new business models, new ways of organizing an economy that we hadn't thought of before.

Right now, people are mostly just retrofitting.

I could go through a whole other set of skill changes that are complementary.

I don't know what they all are.

You have to be creative to think about them.

But that's what the gap is.

In the case of early computers, it's literally like 10 times more investment in organizational capital and human capital, if you look at the size of the investments, to the hardware and software.

So that's very big.

That said, I'm open to adjusting my thoughts on this a bit because ChachiPT and some of the other tools, they have been adopted very quickly and they have much more quickly been able to change things, in part because you don't need to learn Python to the same degree.

You can do a lot of things just in English.

And you can get a lot of value just by putting them on top of the existing organization.

So some of it's happening faster.

And in some of the papers that you may have read for the readings here, you know, we had like 15, 20, 30% productivity gains pretty quickly.

But my suspicion is that it will be even bigger once people figure out these complementary innovations.

And so that's a long way of answering your question about it.

It's not just the technical skills.

It's figuring out all the other stuff, all the ways of rethinking things.

So those of you who are at the business school or in economics, you know, there's a lot of opportunity there to rethink your areas now that you've been given this amazing set of technologies.

Yeah, question.

It seems like you're expressing more caution than Eric was with regard to the speed of transformation.

Am I correct in saying that?

Well, so I would make a distinction between two things.

I'll defer to him and others on the technology side.

We're going to hear from several other folks.

And there are people who are equally optimistic as him or even more optimistic on the technology side.

There's also people who are less optimistic.

But technology alone is not enough to create productivity.

So you can have an amazing technology.

And then for various reasons, A, maybe people just don't figure out an effective way to use it.

Another is it may be regulatory things.

I mean, some of my computer science colleagues introduced and developed better radiology systems for reading medical images.

They weren't adopted because of cultural, you know, people just didn't want them.

They didn't want and there are safety reasons.

When I did an analysis of which tasks I could help the most and which professions were most affected, I was surprised that airline pilots was kind of near the top.

But I think that a lot of people would not feel comfortable not having the pilot go down with you.

So they sort of you want to have the human in there.

So there are a lot of different things that might slow it down significantly.

And I think that's something we need to be conscious of.

And if we could address those bottlenecks, that would probably do more for productivity than just working on the technology alone.

Yeah, question.

So Eric had an interesting comment on data centers in universities.

I think this is a larger point of like, and I was going to ask him why doesn't he write a check?

People are asking him that question.

Sort of like, what is the role of the university ecosystem?

Obviously, there is this larger I'm sure all of the CS professors here.

So I'll take I mean, I think it'd be great if there were more funding.

I mean, the federal government has something called the national AI resource that is helping a little bit, but it's in like the millions of dollars, tens of millions of dollars, not billions of dollars, let alone hundreds of billions of dollars.

Although Eric did mention to me before class that they're working on something that could be much, much bigger.

He's pushing for something much, much bigger.

I don't know if it'll happen.

That's for training these really large models.

I had a really interesting conversation with Jeff Hinton once.

Jeff Hinton, as you know, is sort of like one of the godfathers of deep learning.

And I asked him like what kind of hardware he found most useful for doing his work.

And he was sitting at his laptop and kind of just tapped his MacBook.

And it just reminded me there's a whole other set of research that maybe universities have a competitive advantage in, which is not training hundred billion dollar models, but it's innovating new algorithms like whatever comes after Transformers and there's a lot of other ways that people can make contributions.

So maybe there's a little bit of a divisional labor.

I'm all for and support my colleagues asking for more budgets for GPUs, but that's not always where academics can make the biggest contribution.

Some of it comes from ideas and new ways of different perspective about thinking about things, new approaches.

And that's likely where we have an advantage.

I had dinner with Sendham Melanathon last week.

He just moved from Chicago to MIT.

And he was a researcher.

We're talking about what is the comparative advantage of universities?

And he made the case, you know, patience is one of them, that there are people in universities who are working on very long term projects.

You know, there's people working on fusion.

They've been working on fusion for a long time, not because they're going to get, you know, a lot of money this year or 10 years from now, probably from building a fusion plant or even 20 years.

I don't know how long it is for fusion.

But, you know, it's just something that people are willing to work on even if the timelines are a little further.

It's harder for companies to afford to have those kinds of timelines.

So there's a comparative advantage or divisional labor in terms of what universities might be able to do.

We have just a couple minutes left.

This is kind of fun.

So we'll just do one or two more questions.

And then I want to talk a little bit about the projects.

Yeah.

Go ahead.

I'm Kevin.

I was wondering about the emerging capabilities of AI.

It seemed that Eric was leaning more towards the architectural differences and designing better models versus the last class we talked about, Morse law instead.

So I'm wondering how you sort of...

Well, he said all three.

So you guys remember the scaling laws?

It had like three parts to it.

I think I put the scaling law that Dario and team...

So there's more compute, more data and algorithmic improvements, including more parameters.

And all three of them, I think I heard Eric say all three of them were important.

But not to be dismissed, this last one, like new architectures, all three of them, I think, are being important.

So I think there was another question in there, though, also.

How much closer are we to like an AGI type system?

So Eric doesn't think we're like that close to AGI type systems, although I don't think it's like a sharp definition.

You know, in fact, that was one of the...

I was going to ask him that question, but we ran out of time.

It would have been good to hear him describe it.

But when I was talking to him, it's just not that sharply defined thing.

In some ways, AGI is already here.

Peter Norvig wrote an article called AGI is already here.

I don't know if it's in the reading packet.

I think if it's not, I'll put it in there.

It's a fun little article with Blaise Iarca.

And a lot of the things that 20 years ago people would have said, this is what AGI is.

That's kind of what LLMs are doing.

Not as well, maybe, but it's sort of solving problems in a more general way.

On the other hand, there's obviously many things they do much worse than humans currently.

Ironically, physical tasks are one of the ones that humans have a comparative advantage in right now.

You guys may know Moravec's paradox.

Hans Moravec pointed out that often the kinds of things that a three-year-old or a four-year-old can do, like buttoning a shirt or walking upstairs, are very hard to get a machine to be able to do.

Whereas a lot of things that a lot of PhDs have trouble doing, like solving convex optimization problems, are things that machines are often quite good at.

So it's not quite things that are easy for humans and hard for computers and other things that are hard for humans and easy for computers.

They're not like the same scale.

And next week we have Mira Morati, Chief Technology Officer of OpenAI, briefly the CEO of OpenAI.

So come with your questions for her.

We'll see you.

[BLANK_AUDIO]