

URL: <https://ecorner.stanford.edu/videos/the-near-future-of-ai-entire-talk/>

Andrew Ng is the managing general partner at AI Fund, a startup studio building new AI companies from the ground up; the founder of DeepLearning.AI; and an adjunct professor of computer science at Stanford University. In this presentation and conversation with Stanford adjunct lecturer Ravi Belani, Ng talks about upcoming opportunities for entrepreneurs exploring AI technology and addresses questions of responsibility and risk.



Transcript

(inspirational music) - Welcome everybody to the Entrepreneurial 00:00:19,200 Thought Leaders Seminar, a Stanford seminar for Aspiring Entrepreneurs. ETL is presented by STVP, the Engineering Center, the Engineering Entrepreneurship Center here at Stanford, and BASIS, the Business Association of Stanford Entrepreneurial Students. I'm Ravi Belani, a lecturer in the Management Science and Engineering department at Stanford, and the director of Alchemist and Accelerator for Enterprise Startups. Today we are thrilled to welcome Andrew Ng to ETL. How many people know Andrew? Okay, so Andrew really doesn't need an introduction, but we will give one anyways for those who don't. Andrew is truly a child of the world. He was born in the UK to parents who immigrated to the UK from Hong Kong, but was raised in Hong Kong and Singapore, went to Carnegie Mellon and very early on signaled that he was no ordinary student. He got three bachelor's degrees at Carnegie Mellon in computer science, statistics and economics, and graduated at the top of his class. Then went on to MIT where he got a master's in electrical engineering and computer science. Then came over to the left coast and got a PhD in Berkeley in computer science with a focus on artificial intelligence and reinforcement learning.

Andrew is generally viewed as one of the preeminent thought leaders on AI today. He was the founding... He's the co-founder and head of Google Brain, and the former chief scientist at Baidu, where he built the company's AI group into several thousands, into thousands of people, several thousand people. But he's as passionate about AI as he is also about the development of you, of students around the world. (audience cheering) And I know there's a lot of love for Andrew. (audience clapping) He's a former associate's professor and director of the Stanford AI Lab, and currently an adjunct professor in computer science at Stanford. How many people have taken one of Andrew's classes or wanna take one of Andrew's classes? (Ravi chuckles) And he's a beloved professor here at Stanford, but he is also viewed as a beloved teacher to millions outside of Stanford. He's the co-founder and chairman of Coursera, the world's largest MOOC platform, and through his online education work and his online AI education work, he's reached over 7 million people. He was listed as one of the world's hundred most influential people by Time Magazine in 2013. And today, Andrew's the managing director and partner at the AI Fund, which is a startup studio building new AI companies from the ground up, and is also the founder of deeplearning.ai.

He focuses his time primarily on his entrepreneurial ventures, looking for the best ways to accelerate responsible AI practices in the larger global economy. There's fantastic content already online that Andrew has given, including a longer version of today's talk that you can find on YouTube. And so instead of reduplicating that, Andrew's gonna give a teaser talk, a

10-minute discussion followed by, we'll do a quick fireside chat and then we're gonna open it up for really interactive Q and A with you. So start thinking about your questions now, because the time is gonna fly by, but without further ado, please welcome Andrew. (audience cheering and clapping) - Thanks a lot, Ravi. 00:03:22,710 Thanks, good to see everyone here. Can everyone in the back hear me okay? Cool, awesome. You know, I've taught CS-239, my machine learning class in this room many years, but all these years I've taught in this room. I've never seen my face that big before. What I'd like to do today is chat with you about opportunities in AI.

So one of the difficult things to understand about AI is, is a general purpose technology similar to electricity, meaning it's not useful just for one thing, it's useful for a lot of different applications. If I were to ask you, what is electricity good for? It's almost hard to answer that because it's useful for so many different things, and AI is like that too. So one of the major trends that we've seen in the last year, few years is that prompting is revolutionizing AI application development. And I wanna just dive a little bit deeper into this, because I know this is an engineering class. I know many of you may be from an engineering background. I'm gonna just go a little bit deeper than this into this than than I might otherwise. But if you were to say, want to build a AI system for many years, the typical approach this use supervised learning. So let's say I want to build a system to rate restaurant reviews as positive or negative sentiment. You know, then you would collect data. Maybe that takes me a month.

I would train the AI model. Maybe that takes me a few months. Find a cloud service to deploy my AI model, maybe that takes, you know, a few months. And so for the past, most of the past decade, a realistic timeline to build and deploy, a value AI system was, you know, maybe six to 12 months. But with prompting, the timeline is now very different. You can specify a prompt in minutes or hours, and then deploy a system to production in just hours or days. And I know that probably maybe most, maybe all of you will have played with large language models as a consumer too, like ChatGPT and Bot and Bing Chat. I think that in terms of startup opportunities, I'm excited about the use of large language models, not as a consumer tool, which is fantastic and exciting, I use Chat GPT and Bot regularly, but instead the application of large language models as a developer tool. Because this is allowing a lot more AI applications to be built and dramatically lowering the barrier to building many applications. So I know that in this talk you don't normally have speakers write code, but this is an engineering class.

So let me actually show you, you know, exactly what I mean by that, right? It turns out that if I want to build an AI system today, this is all the code I need. And this means that, you know, if you take CS-106 or something, we're going to code, right? In a CS class with just a little bit of code, and put Open AI tools, load my key. I don't know, the what? STVP lectures are great, you know? Also made many friends. I've never written that before. And so hopefully this, okay, thank goodness got that right. And so this is positive sentiment and just in seconds, you know, that's all the code it takes now to build an AI system in code to look at a piece of text and process it, to look a piece of email and route it, or to start to build the beginnings of a chat bot. So over the last don't know half year, one of my teams, deeplearning.ai has been working with many of the AI tool builders to create short courses on how to use tools like I just showed you. Because there are many AI applications that used to take me six months to build, that I think any of you will now be able to build in one or two days. And this opens up the set of things that you could do, and the set of the prototypes you can build. And in fact, from a startup perspective, when it took us six months to build something, you know, what we do is have a product manager study it, do the user studies, make sure it's the right thing to build, then go build it.

And after all that investment is like, boy, let's hope it works. But what I'm seeing with these very fast development times is if it takes you a couple days to build something, I'm seeing a lot more startups as well as big companies say, "You know what? I have 10 ideas for features, I'm gonna build all 10 things and then just ship them all, and then we'll see how users use them or don't use them, and we'll just keep what sticks." And this is very different prototyping, much leaner methodology than I've seen startups use, you know, before prompting. With one important caveat, which is that responsible AI is important. So don't do this, don't ship, you know, things that could cause harm. But we have a lot of applications like inspecting bits of metal in factories, you know, where there is really no harm, no risk of bias, where I think there's very fast shipping methodology, that just innovates very quickly in AI. So where are the opportunities? So the size of these circles shows what I think is the value of different AI technologies today. Supervised learning started to work really well about a decade ago at labeling things. Such as label this ad as, is this something you're likely to click on or not? Or label this X-ray with, you know, what's the medical diagnosis? And supervised learning for a single company like Google is worth more than a hundred billion dollars a year. And there are millions of developers working on it. And it might even grow in the next three years to double say, so massive momentum, lots of applications to be figured out.

And then generative AI is a new entrance, where frankly the revenue, the value of the revenue from generative AI today is much smaller. But given the amounts of interest and excitement and commercial interest, I think it will much more than double in the next three years. And three years is an artificially short time horizon. I think you were to look out six years, if it continues to compound at this rate, maybe the value from generative AI will, you know, even start to approach that with supervised learning. But all that room for growth, the light-shaded region for supervised learning or generative AI, which are probably the two most important tools today, are where there are a lot of opportunities for any of us to identify and build to concrete use cases. And what I hope to take away from this talk is AI technologies are general purpose technologies, meaning that they're useful for many different tasks. When supervised learning started to work well about a decade ago, it actually took us a long time, it took us annoyingly long over the last decade. And it will be, take us annoyingly long over the next decade to figure out use cases for generative AI. But still want to use this to make ships more efficient or for medical

diagnosis, or for education product recommendations or something else. We're still figuring out concrete use cases for supervised learning.

And even though we're not yet done doing that, we have another fantastic YouTube generative AI that even further expands other things we now do with AI. And, you know, one important caveat, which is there will be fast along the way. How many you remember Lensa? Raise your hand if you do. Wow, almost no one, that's fascinating. So Lensa's revenues took off like that through last December. It was this app that could let you upload a few pictures of yourself and draw a cool picture of you as an astronaut or a scientist or something. And it was a really good, really hot product until last December, after which its revenues did that. And I think that's because Lensa was one of what will probably turn out to be multiple thin software layers built on top of someone else's very powerful API. That was a good idea, people liked it, but it wasn't a long-term defensible business. And when I think about generative AI as a developer platform, I'm reminded of when, you know, Steve Jobs gave us this phone, right? And shortly after someone wrote an app that I paid 1.99 for to do this, to turn the phone into a flashlight.

And this was also a good idea, it was a great product, but it just was not a defensible business either, because there's a very thin software layer built on top of someone else's very powerful development platform. But in the same way after we got the iPhone, after we got the smartphone, someone else figured out how to build Uber, Airbnb and Tinder, much longer-term defensible, very valuable businesses that are still standing the test of time. And I think what those opportunities as well to build long-term valuable franchises of businesses on top of generative AI. So where are the opportunities? So I felt, yeah, I felt years ago, but even more strongly now, that because of emerging AI technology, there are a lot of projects that are now possible that were not possible, You know, one or a handful of years ago. And I wound up starting AI Fund, which is a venture studio that sequentially we're a set of entrepreneurs to start companies. We actually average about one startup a month now, because I felt, I previously, as Ravi mentioned previously, I had led AI teams in Google and Baidu. And having led AI teams in big tech, I couldn't see how it could possibly operate a team in a big tech company to pursue the very diverse, very different sets of opportunities that I saw and wanted to pursue. And starting different startups to pursue those valuable projects seem more efficient than having one company, even the big tech company go after such a large set of resources. But having said that, I think AI and generative AI also offers a lot of opportunities for incumbent companies, which often have a distribution advantage, right? Where exactly are the opportunities? So this is what I think of as the AI stack. At the lowest level is the hardware layer.

Very valuable but also very capital intensive, needs a lot of resources to build and very concentrated. So I'm sure there'll be valuable startups built there, but I personally don't play there because of how capital intensive and how concentrated it is. There's a cloud infrastructure layer, also very capital intensive, very concentrated. Very valuable, but at least when I build startups I tend not to play there. The other layer that's interesting is a developer tooling layer. So what you just saw me do was use Open AI as a developer tool. And I see this space as hyper competitive. Look at all the startups chasing Open AI, but there will be some mega winners. So whereas incumbents have a startup kind of a distribution advantage, I think for many startups, having a technology advantage may give you a best shot at doing something meaningful there. So I personally tend to play it here only when we think we have a technology advantage, 'cause that buys us a better chance to become one of the huge winners.

And then with most ways the technology innovation, lot of the media attention, social media, what people tend to talk about is the tooling, the technology layer. There's one other layer that I think has got to be even more valuable, and that's the application layer. Because in many ways the technology for the end front and tooling layer to be successful, applications need to be built on top of them. They generate even more revenue so that they can afford to play the infrastructure layer. And what I'm seeing is that there are a lot of opportunities at the application layer where the intensity of competition is not, frankly not nearly as high. Maybe just one example, I've been chatting a lot with the CEO of Meeno, which is a startup that applies AI to romantic relationship coaching, right? And, you know, I'm an AI guy, I feel like I don't know anything about romance. (audience chuckling) And if you don't believe me, you can ask my wife, she will confirm that I don't know anything about romance. But when we decided, when we had conviction that AI could be applied to relationships, we wound up partnering with Renate Nyborg, who's the former CEO of Tinder. And because she ran Tinder, she understands relationships in a very systematic way, more so than anyone else I know. And so with my team providing AI expertise and her providing relationship expertise, we're able to build, you know, a pretty unique relationship mentoring application that we just announced a few weeks ago.

And this you may not probably know, but Renate actually occasionally stops by Stanford campus and talks to Stanford students as part of her user product research. So it's possible you've seen her around. Just one last thing and I'll love to go with Q and A. Over the last few years, AI fund, we've been tuning our process for building start ups. I'll just share that with you. So we often start off with a lot of ideas, right? And one example of another startup we built was Bearing AI, which uses AI for smart routing, a very large oceangoing vessel. So if you're a ship captain should just sail at 20 knots or 22 knots, it's like, who knows? Most ship captains just make some decision. But because we are able to get global weather and ocean current data, we can make recommendations to ship captains for how to get there on time and use about 10% less fuel. But this idea was suggested to me by Mitsui, which is a major shareholder in a major shipping line that operates very large ocean going vessels. And this one of those things, I would never have thought of this idea myself, 'cause, you know, like I've been on a boat, but what do I know about global maritime shipping? But Mitsui suggested this idea to me.

And we then validate the idea, make sure there's a technical feasibility and a market need. Recruit a CEO. We are fortunate to find Dylan Keil, who's a fantastic CEO with one successful exit before. And then we spent three months in our current process building a technical prototype with the CEO and doing deep customer validation. If it survives, two thirds chance of surviving, one third chance of not surviving, we then write a check in that allows the company to hire executives, build an MVP, and off it goes to raise additional rounds of capital. And I think this is what we- And so Bearing AI, well, now it's actually, there are now hundreds of ships on the high seas guided by Bearing AI. Ships guided by Bearing AI have 75 million miles, which is the equivalent of going 3,000 times around the planet and save about half a million dollars in fuel costs per ship per year. in addition to significant carbon emissions. I think we've save about, I wanna say about a million tons of CO2 emissions so far. But this kind of idea that like I would never come with this idea myself, but I've learned that my swim lane is AI, but when I work with experts in other sectors, there are often these exciting opportunities that are very valuable, but frankly, how many groups in the world are experts in AI and shipping? Or expert in AI and relationships? I find that the competition intensity at the application layer is often much lower.

And then just one last thing kind of, you know, just full disclosure, something I hope all of you will do too. My teams only work on projects that we think move humanity forward. Responsible AI is important and on multiple occasions we've killed and I will continue to kill projects that we may assess to be financially sound, but based on ethical grounds. So lots of exciting opportunities I think at Stanford, the lots of great costs can take in engineering and elsewhere to learn about that AI tech. And then when you find, you know, applications or go play at the infra tooling there too, I think there are lots of opportunities, but I think there are, what I'm seeing is, frankly my team at AI Fund, we have so many startup ideas. We use a task management software, we use Asana to track this huge list of ideas, and it's actually quite clear to me there are a lot more good ideas for AI businesses than, you know, people with the skill to work on them at this moment in time. So hopefully there'll be more than enough projects for everyone, for all of you, all of us to work on. All right, thank you. (audience cheering and clapping) - I wanted to just start off with that closing statement 00:19:00,360 that you made about how there's more opportunity than there are students with skills or people with skills to pursue them. And given that we have an audience full of students, I wanted to start off by mapping out advice for students that are entering into the university regarding AI.

So if you wanna pursue a career in AI right now, and let's say your child was entering Stanford, what advice would you give them in terms of how to spend their time? - Yeah, so, you know, there's one thing that's actually 00:19:28,048 really worth doing when you're assigned to students, which is take classes. Because it turns out that I feel like there's actually one pattern I see for both undergraduates and graduate students, including PhD students, which is, there's so much exciting stuff to do, you just wanna jump in and do it, right? In fact, I've seen undergrads in their freshman year, you know, try to join a research lab and start doing work in AI. That's okay, nothing wrong with that. But it turns out that while project work is one way to learn coursework is I think an even more efficient way to learn, especially when it comes to mastering the fundamentals. Because professors will put a lot of work to organize the material in a way that's efficient to learn and digest. So I would say, you know, take classes in AI technology or in entrepreneurship and gain those skills. I've seen students jump in and then if you are trying to work in a research lab without strong skills, you end up, you know, like labeling data or something, which is fine. You learn some things, but you actually learn a lot from taking courses. And then in addition to that, after you start to master the foundational skills, after you know how to use AI technology or, you know, then as you start to practice, find exciting use cases across campus. I do a lot of work, you know, over with people over in climate science or in healthcare to take my AI expertise and then marry it with a different discipline that I'm not expert in to find exciting applications.

And hopefully that type of a practice will help many of you find exciting projects to work on as well. - Do you need to take technical classes? 00:21:01,230 Do you think you need to take computer science classes if you want to pursue a career in AI? - Need is too strong, 00:21:08,580 but I definitely encourage you to take technical classes. I think we're moving toward a world where, frankly I, at some future point, I think everyone should learn to code, or rather, I think it'll be useful for everyone to learn how to code for a couple reasons. Everyone has access to data, right? This is different than the world used to be even a few years ago. And especially with generative AI, your ability to get something to work is much higher than ever before. The barrier to entry is much lower than ever before. And so if you learn just a little bit of coding, the amount that you really accomplish is significantly greater than if you don't know how to code at all. - And are there any skills that separate out 00:21:48,900 the great AI founders? I know AI right now is like, it's a sea that's rising all boats, but if you separate out the great ones from the good ones, are there any salient skills that you notice that the great AI CEOs or founders have that the good ones don't? - Maybe since you say AI, 00:22:06,705 I would say is often technical depth. - Okay. 00:22:11,067 if you say great founders around great AI founders, but I feel like AI is evolving rapidly, and we definitely have lots of entrepreneurs that, you know, pitch the VCs without really knowing what they're doing, and the smart VCs can stiffen up quite quickly.

And it makes a huge difference. I think the technology unfortunately, you know, is like somewhat complicated for a lot of applications. So a team that actually knows what they're doing will execute an AI project 10 times faster, you know, than a team that doesn't. And 10 times is not a made up number. I literally see people take a year to do something. I go, boy, I know that other team would've done, performed this level in two weeks or maybe a month. So for many AI startups, application starters, infra startups, you kind of have to know what you're doing. So doesn't have to be you. If you're a technical co-founder, maybe that's okay. And then second thing I see among many of the great founders is speed.

I find that as a startup you'd be surprised, we've found that the great founders, the sheer speed of decision making. And,

you know, I sometimes talk to people from big companies and they'll say, "Oh, we move so fast." But when I kind of sit them side by side, how does it take you to make the decision? Then I talk to great founder, how does it take this decision? Maybe here's one story. I was chatting with the Meeno CEO, Renate Nyborg, former CEO of Tinder. I was on the phone with her one day and she was making a major architecture decisions. Because the architect thing, you know, there are basically two architect- Two major software architectures under consideration. And the team had laid it out, list out some pros and cons. So this team with me and some of my friends and said, "These are pros and cons." And then one of my- My AI Fund and I said, "You know, we're not sure, but here has some reason that we prefer architecture A." And then Renate said, "Okay, guys done, decision made, go and implement, you know, architecture A." And after I thought, well, did Renate just make a massive engine decision in basically 30 seconds? And she did. And I realized after it, I don't think there was a better way to make it, because it's not as if, you know, if the company waited another week would've been a high quality decision, and if it was wrong, I'm sure they would fix it, you know, the next week. But until you've lived through the speed of a great business, most people, I know so many people that think their organizations are fast, but we stack them up to the real speed of a fast moving CEO, they have never actually seen speed in their life. One important caveat, do be responsible.

I know that move fast and break things sometimes, you know, is the wrong approach, so tremendous speed when you are not being callous with people's lives and livelihood and things that could cause real harm. But so long as there's an important caveat of responsible AI, many of the great CEOs move faster than most people realize people can move. - And so let's just double click on that, 00:25:00,750 on this theme of responsible AI. Just because I know this is a hot topic that maybe people aren't thinking about, which is you are clearly on the side of AI for good, for responsible AI. Many of your brethren like Jeffrey Hinton and other famous leaders in the AI space have come out and are concerned that the pace of AI development will become an existential threat to humanity. So much so that famously there was a petition signed by Elon Musk and Steve Wozniak and many thought leaders asking for the halting of the foundational- The deepest foundational models of AI for society to sort of catch up. You did not sign that pledge. Can you share a little bit more detail about that? Was that a difficult decision for you to make? And can you share more details about why you didn't join them, and what your philosophical view is regarding if AI poses an existential threat? - So I honestly don't see how AI poses 00:25:57,330 any existential threat to the human race. We know AI can run 'em up, you know, self-driving costs have trashed leading to tragic loss of life, automated tradings trash the stock market. So we know poorly designed software systems can have a dramatic impact and responsible AI is important.

But recently I sought out, you know, people like Jeff and others that were concerned about the question of AI extinction. And I tried to understand why they thought this way. Some were worried about bad actor using AI to create a bio weapon. Others were worried about AI evolving in a way that inadvertently leads to human extinction, similar to how we as humans have led to the extinction of many species through simple lack of awareness sometimes that our actions could lead to that outcome. But when I tried to assess how realistic these arguments were, I found them to be vague and nonspecific about how AI could cause all. And I think that, I found frustratingly frankly, that trying to prove AI couldn't is akin to proving a negative. And I can't prove that super intelligent AI won't be dangerous, but I can't seem to find anyone that really knows exactly how it could be. And... But I do know that humanity has ample experience controlling many things far more powerful than any one of us like corporations and nation states. And there are many things that we can't fully control that are nonetheless safe and valuable, like airplanes.

You know, no one can control an airplane, it's baffled around by winds, and the pilot may make a mistake. But in the early days of aviation, airplanes killed many people. So we learn from those experiences, build safer aircraft, devise rules by which to operate them, and today most of us can step into airplane without fearing for our lives. And I think it will be like that too for AI. So I think the AI extinction I find to be very unfortunate. What I'm seeing because, you know, doing some work in K-12 education as well, what I'm seeing is that kind of really unfortunately, I see high school students now considering working in AI and some will say, "AI seems exciting, but I heard it could lead to human extinction, and I just don't wanna be a part of that." And so I find that the over hyped AI extinction narrative is doing real harm. So I'm very concerned about that. - Thank you, Andrew. 00:28:13,290 One more question then I'm gonna open it up, which is I loved the detail on the low-hanging fruit opportunities. I know that's on everybody's, all the entrepreneur's minds of what to pursue.

And so I appreciated the attention and the presentation on that. I wanted to ask about what's gonna be the next big technology shift in AI, because things are changing so rapidly, especially as the models now are getting smaller and open sourced, it feels like we've already conquered language. Visual AI is getting very, very good. What's next? What are you seeing that's around the corner that others might not be aware of? - Yeah, you know, I've done it. 00:28:43,680 About six, several months ago, I was predicting visual AI is coming next, but now it's like, all right, visual, so I guess now got to come with something new. But in all seriousness, I think visual AI would be much more about the analysis of images, rather than just generation of images. But I think we're at like the GPT 2 moment for visual AI. Is not yet working, but I think it'll work much better and this will impact self-driving cars, for example, when we can finally, you know, solve problems in the long tail. And then I think, actually one of the thing that I wrote about just today in a newsletter called The Batch is I think one thing that many people find controversial, but I think is coming is the rise of edge AI. And I know this is controversial, many of us who trained to write SAS software, you know, lends itself nice to subscription business model.

How do you even find people? Like how do you hire engineers to write desktop applications? Like who even does that anymore? But I think that because of, for various causes including privacy, I think that in the next few years we'll see more AI

applications running at the edge, meaning on your laptop or on your cell phone. So I think that'll be coming. And then I think there'll just be a lot of work coming in the application there as well. - Okay. I wanna open it up to the students. 00:29:56,177 You're the reason why Andrew's here. Student You mentioned that the... 00:30:00,210 On your slides you put the potential from reinforcement learning or the general value is a dot relative to the potential for unsupervised learning. Do you think there is still potential for generalist agents like Gato and other reinforcement learning models in society and in your AI stack for startups? - So, okay, 00:30:22,590 so technically large learning models are trained using reinforcement learning and unsupervised learning and supervised learning. But leaving that aside, I feel like- I'm not convinced that reinforcement learning is near a breakthrough moment, at least in the next small number of years.

A lot of excitement about what we could do in reinforcement learning applied to robotics. A lot about, you know, CS faculty, right? Chelsea Finn, Emma Brunso, many others are doing exciting research there. But we do have a data problem. So it turns out that texts on the internet sounds a lot like texts on your documents. So we can learn from lots of texts on the internet to do really well on your text documents. And images on the internet look a little bit like images that you care about. So we have a lot of data, but because every robot is different, I'm struggling, many people are struggling to see how to get enough data to have to use your recipe of scaling up data and compute where for reinforcement learning. And people are working on it. Over the weekend at the CS faculty retreat, you know, there was a talk I think- Who gave the talk? Shoot, I think it was on how to do this, early ideas of how to do this. But I think we're still a few years away from those breakthroughs in reinforcement learning.

But it's a great research topic by the way. Just because, you know, just because it's not working right now doesn't mean you shouldn't do research on it. So I think it's a great research topic. Student So I just want to know your thoughts 00:31:43,980 about what are the security concerns which is coming up by you abusing that like LLM models, like all these new attacks like prompt injections, data leakage, jailbreaking. So what's your thought around that? Like... - Yeah, so- 00:32:01,170 Because it's just starting up this new technology, so I'm assuming there's more things which will come up. - Yes. 00:32:09,206 So I think that for the near future there'll be a bit of, a little bit of a cat and mouse thing going on. So I think... I'm seeing different companies approach this with different tools to watch out for prompt injections, for data leakage.

Actually deeplearning.ai is actually working with a partner on some things that they hopefully will announce very, you know, soon on a portfolio of tools. By the way, those of you that have not yet, you know, done it, go and fool around with prompt injections, see if you can get an LLM, you know, to do something. Well, don't do something actually harmful. But I actually find it kind of intellectually interesting whenever I use an LLM to prompt it to see how robust the safeguards actually are. And it actually looked the older language models a year ago, it was super easy to get the older models, you know, frankly to give you detailed directions to do things that they should not give anyone detailed directions to do. But the more modern language models are much smarter, but it's still sometimes possible. Sorry. But what I'm seeing as well for a lot of corporations, a lot of corporations because of these worries will ship internal facing product first. Because presumably, you know, if it says the wrong thing to your own employee, more understanding, less likely with scandal, and test products internally for quite a long time, or even build capabilities for safe internal use before turning out to external use. But I do see different companies, yeah, different tools for trying to reduce this.

Ravi Terrific. Next question? 00:33:43,620 Student Thank you so much. 00:33:45,000 Hi there, my name is Chinette and I'm an international student from Hong Kong. I'm curious to ask, 'cause you know, I'm hoping that after I graduate I can hopefully go back home to work closer with family, but at the same time I feel like by going back I'm closing a lot of doors behind me because for example, in Hong Kong for example, you can't access ChatGPT without a US number, which makes access to some of these resources really difficult. So I'm curious to see what are your thoughts about navigating this complex modern landscape? - Yeah, I don't wanna comment on, I don't know... 00:34:15,540 Complicated. That's actually one thing I'm seeing. I've been to quite a few places, you know, in Asia recently. And what I'm seeing is that many countries are developing surprisingly good capabilities for building large language model applications. The concentration of talent for generative AI deep tech is very concentrated in the San Francisco Bay area.

I think because there are basically two teams that did a lot of the early groundbreaking work. You know, Google Brain, my former team and Open AI, and subsequently people left and started a lot of companies here in California Bay area. So I think that concentrated talent is very high. And it's interesting, even when I'm in, you know, Seattle, great city, love the city, on weekends, you know, I hang out with friends but the conversations are not about generative AI. Whereas here, you kind of like, if you go to coffee shop, actually one of my friends was visiting from Taiwan, so he was hanging out with us of a week. They went back and he said, "Yeah, I went to coffee shop and, you know, there was no one talking about AI, that's so weird." (audience laughing) So at least at this moment in time there's really heavy concentration, but I see less the deep tech layer, but the application layer, I see that skillset developing quite quickly globally as well. Oh, and I think the opportunities in a lot of places would be local opportunities. So the shipping company that we built, we built for a Japanese company that happens to operate global lines of shipping. So I think a lot of the businesses will be, you know, playing locally where that country or that geography is strong. Those businesses will be more efficient to build in places other than Silicon Valley.

'Cause where do I go to find a large seaport, you know, here to do that type of work? Student Hi Andrew, thank you so much for your time. 00:36:01,830 My name is Komel. I wanted to ask you if you think we'll ever reach a threshold on human dependence for AI? Or if you think it'll just continue to grow exponentially. - So I think we are already really, really dependent

00:36:15,210 on tech, right? Imagine if, you know, if the internet were to shut down, I think people would die. I don't think that's exaggerated. I mean, but seriously think about, you know, how we get food supply chain, healthcare. If the internet were to shut down, I think that will lead directly to, you know... What will happen to our water system, right? Healthcare system. So, and I think that technology is very useful, and so long as the supplies remain reliable, I feel like it's okay to depend on technology. I mean, heck, I wish- I don't know, without dependence on agriculture system how many of us would build a farm and hunt enough food to keep ourselves alive? Maybe, maybe we could do it, but it's pretty challenging.

So I think dependence on tech seems it's gonna keep on growing for a while. - But do you think there'll be a moment 00:37:06,932 where there's a difference in that relationship, not just in degree, but in kind? You know, the famous singularity point where we don't even know what we don't even know about how technology's developing? Do you think that will occur? - Yeah, you know, the technological singularity 00:37:20,850 is one of those hypey things that I don't even know what it means. So it is one of those, it is exciting science fiction, but as an engineer and scientist, I don't know how to talk about it. It turns out there are a few terms in AI that are vague and undefined, but there are a lot of emotions, a lot of excitement about it. And I don't really know how to think about those things in a systematic, rational way. But I think our tech- There's actually one thing. I think that our technology, our relationship technology is changing rapidly. Today, you know, I probably use Chat, you know, GPT for or Bing or Bot pretty much every day now. And so the workflow of many people have changed, right? Yeah. I think keep on changing.

- And do you have a view, 00:38:07,710 I know this is also might be more of one of these sort of hot topics that's not substantive, but on the consciousness of AI? That AI will become conscious? - Yeah, so the thing about consciousness is, 00:38:18,870 is important philosophical question, but I don't know of any test for whether something is conscious or not. So I think it's important philosophy, and philosophy is important, but as an engineering scientist, I don't know, there is no definition for what is conscious or not. And thus we can kind of debate it, you know, at length. And there's actually one other formula for hype, which is if someone comes up with a very simple definition for consciousness, so someone says, "Oh, if you can recognize yourself in the mirror, you're conscious," I made that up. It's not a good definition of consciousness, but you're aware of yourself, see yourself in the mirror, then it's actually pretty easy to get a robot to recognize yourself in the mirror. And then you can generate newspaper headlines saying, "AI has achieved consciousness." What it did for your kind of, you know, silly little- For your very small definition of consciousness, but that gets misinterpreted by the broader public for a grander statement than it is. So I see some of that hype in AI as well. Ravi Terrific, thank you. 00:39:13,053 Next question. Student Hi.

00:39:18,870 Earlier you outlined the AI stack and recently we've seen a lot of cool things coming out of like Nvidia, Intel and other like tech companies. I'm curious on what your thoughts are on what companies like AWS and Google, like in the infrastructure layer need to do in order to make like AI and enterprises and business really effective and possible? - Sure. Boy, so there's a lot going on in that space. 00:39:40,440 By the way, you mentioned Intel and Nvidia. I wouldn't- I think I'm actually seeing really exciting work from AMD as well. I've been pretty impressed by the MI200, MI250, and I'm excited about MI300 GPUs coming out as well. And I think the ROCm stack is becoming, you know, not parity of CUDA but better than most people given them credit for. But in terms of AWS and Google. So it turns out that if you were to use a lot of the LMM startup tools, the switching cost is actually pretty low. So we were to, you know, start with one LLM-API call, if you want to switch to a different LMM provider, the number of lines of code is actually pretty low.

So they're low-switching costs. But it turns out that zero and Google Cloud and AWS are fantastic businesses, because once you build on any of these clouds, you know, the switching costs tend to be very high because you have so many API host integrations. So that's why I think that a lot of the startups selling API calls still have, you know, some work to do to find a business model that may be somewhat more defensible. I think that Open AI's Chat GPT enterprise, that feels like a, you know, more defensible business than just selling API calls. By the way, Sam was actually a Stanford undergrad. He actually interned in my lab, so a lot of Stanford roots, but he's a smart guy. I'm sure he'll figure out- Confident he'll figure out some good directions. Yeah. And then I think AWS and GCP and zero are all racing to continue to develop LMM capabilities and make it easier to use and bring more customers and yeah, it's a very dynamic space. - And as AI gets democratized, 00:41:19,770 it feels like things are shifting more towards compute and data as predictors of success.

If that's the case, do you think the locus of innovation shifts from academia to industry, where the companies are gonna really be dominating at the forefront of AI? - Yeah, so what I'm seeing now is that there's a subset, 00:41:36,990 but there's a small subset of things that are easier to do in the big tech company, which are the ones that require massive compute resources. And I do think people's perceptions are distorted, because frankly, I've been on the big tech companies before, right? So I understand, you know, marketing and big tech companies, but standard big tech company marketing is, look, you need the data, you need to compute, only we have it. Why don't you just give up and don't compete with us, right? Or come apply for a job and come work with us. That is, this has been the explicit PR strategy of at least one big tech company because I know, you know, what was discussed internally exactly at that big tech company. So I would say don't buy into that marketing message. It is true that there is a subset of work that requires massive capital, training in very large foundation models. That is much easier to do in a big tech company than in academia like Stanford. But that's a small subset of all the happenings in AI, and there's plenty of work at Stanford at the application layer. It turns out because of scaling laws, we're actually pretty good at predicting what will happen for very large models by training on more moderate-sized models. So very good scientific work can be done at much smaller models.

And then also, you know, I routinely run kind of, you know, models on my laptop for inference. Like, I don't know, when I'm on an airplane, you can run like the 7 billion LLaMA model on your laptop, right? And so there's actually a lot of stuff that you could run on your own personal computer. Ravi Terrific. Thank you. 00:43:01,263 Next question. Student Thank you Andrew. 00:43:08,640 So from AI expert and also the investors' perspective, so what AI driven healthcare applications do you see have the great potentials to have the breakthrough in the future? And what challenge and obstacles should we, we be aware of? Thank you. - Yeah, so boy, there's a lot of complexity that question. 00:43:26,340 So I feel like a lot of healthcare people tend to focus on the diagnostics and the treatment. So I think lots of opportunity is there.

I think that the revenue model is to be sorted out. So we've seen, you know, Pair and (indistinct) struggle in the public markets, or kind of bankruptcy kind of levels almost. So I think prescriptive digital therapeutics is definitely going through challenges, but what's the recipe for shipping AI products and, you know, in the payer-provide ecosystem, what will payers be willing to pay for? I think that many businesses are sorting that out. I think that will work. And it's actually one other huge set of options in healthcare that I think tend to be underappreciated, which is operations. Instead of the medical stuff things like scheduling, you know, who should scheduling the MRI machine or doing kind of a patient management systems. I think those type of healthcare operations have fewer regulatory hurdles, and I think is also rich set of options. And then lastly, does the go to market question of do you wanna go to the market in the US or in other countries where the regulatory hurdle could be very different depending on? The US fortunately doesn't have as great a shortage of doctors as some other places. And those, therefore other places that are more amenable to, you know, responsible but still easier adoption of AI than the United States. Student Okay.

I have a super quick question. 00:44:47,190 You mentioned that your team at the AI Fund has so many ideas for AI applications that you have a whole asana of them. What exactly is your process for generating these ideas? - Oh, thanks, Jose. 00:45:01,110 that deeply understand the domain. It turns out that there are a lot of people in the world, you know, including like CEOs of Fortune 500 companies, but really a lot of people that really understand the domain have thought deeply about something for months or even a couple years. And when we get together with them, they're sometimes very happy to share their idea with us because they've been looking for someone to validate or falsify it, and also to help them build it. So we actually get a lot of ideas, some internal with a lot from subject matter experts that just not yet had a AI building partner. - Terrific. That's fantastic. 00:45:29,611 Thank you Andrew so much for sharing your insights.

(audience clapping) Lots of love. Thank you for sharing your insights with Stanford's ETL course MS&E 472, and the students all around the world. Everybody, next week we're gonna be joined by Stanford professor Kathleen Eisenhart here at ETL physically in person. Professor Eisenhart is also the author of Simple Rules. You can find that event and other future events in this ETL series on the Stanford eCorner YouTube channel, and you'll find even more of our videos, podcasts, and articles about entrepreneurship and innovation at Stanford eCorner. That's ecorner.stanford.edu. Thank you everybody, thank you Andrew. - Thanks, Ravi. 00:46:12,532 (audience clapping) (soft music)..