MIT CSAIL Alliances | Thompson_Project_6(1)

Welcome to *MIT'sComputerScienceandArtificialIntelligenceLabsAlliances* podcast. I'm Kara Miller.

[UPBEAT MUSIC]

On today's show, a researcher who spends his time thinking about when business leaders win big by using AI and when they don't.

So much energy right now is on the hype. People are like, I'm going to do everything, and you have to say, Well, maybe it doesn't work for everything.

Neil Thompson talks about who the first movers are and why they're shelling out money to be first, plus tuna, ice cream, bananas, and what they should teach businesses about adopting AI.

There are some predictions that you're going to make that are just not that valuable, and not worth it. So that can of tuna it can sit on the shelf for an extra two days and it doesn't matter, it doesn't cost you very much.

And importantly, you don't have to spend a whole lot to figure out exactly how many cans of tuna to stock. Today, we'll look at how you make those kinds of predictions and when it's worth making them. That's coming up next.

It's important to remember that bigger is not always better. The biggest building is not necessarily the prettiest, the biggest company sometimes loses its way, but this is a story about how bigger is better, at least when it comes to deep learning. And to make the point, we're going to talk for just a minute about proteins in the body.

So protein folding comes because when the body produces the proteins that make up, basically, our entire body, when our bodies produce those, they actually produce it as sort of one long line. So think of it as a sentence with a whole bunch of words, and what's weird is it folds into this three dimensional shape. And so it has to fold up in this useful way that makes it sort of a three dimensional thing that's in the body.

Neil Thompson is the director of the FutureTech Research Project at MIT, and he also has appointments at CSAIL and MIT's Initiative on the Digital Economy, and he says, Both people and computers have tackled the problem of how do you fold this protein for a long time, and computers weren't very good at it, until that is, they were.

In 2020 the journal *Nature* announced that DeepMind, an AI company that had been acquired by Google a few years before, had made what *Nature* called, A Gargantuan Leap, and that leap was that DeepMind knew almost perfectly how proteins would fold in the body.

What's really I find amazing about this example is it shows just how important it is to have this very, very large amount of computing power. And what I mean by that is harnessing more and more computers than we've ever had to solve this particular problem.

Thompson notes that a bunch of folks in academia were trying to use deep learning to address this very problem, of how proteins would fold, but he says, The people at DeepMind just had a lot more power at their disposal, 1,000 times more compute, and honestly, that made all the difference. So here's the takeaway.

Smart people are great and you want smart people working on your problem, but even so, they are going to be limited by the amount of compute available to them.

So if you have two people, one of whom is at an organization that's got a lot of resources, has a lot of compute, and you have someone equally smart in an organization that doesn't have it, the person who doesn't have it is not going to be able to compete.

What does that mean for the future? Well, not surprisingly, it's going to have profound effects.

Only people with a lot of resources are the ones who can really be at the frontier, and that means that if you're at Google, it's pretty easy to do these things, and if you're at a small university that doesn't have much, it's very hard for you to do this, and we can already see these effects in some of my research. We can see that even in AI research, so the stuff that academics specialize in, we already see the influence of industry going up because they're the ones who have the resources to do the really big problems.

Within industry, it seems like it would also be an issue because you've got behemoths with very, very deep pockets and a lot of access to compute power, and then obviously, there's always like startups or small companies and they want to really get far in deep learning too. When you just look at that business landscape, what does that say to you about what that looks like?

Yeah, so let me maybe take just a brief moment of digression before I answer your question, because people might say, Well, wait a minute, you're talking about only big companies doing this, but actually, if you think about something like ChatGPT, anybody can sign up for ChatGPT, and it costs them $20 a month or something like that. Doesn't that mean that everybody can use it?

Yes, it's true that it's open and available like that, but the question is, how many of these big models are there going to be, and who's going to control them? Because, of course, if one person controls them, they can also control the price. Economists worry about that a lot about the monopoly effects of these things.

And so I think we absolutely can say, that because these costs are so high these big firms are going to be dominating the creation of these big models and that's going to give them a lot of control about how these models perform, what they're optimized for, and ultimately who can use them at what price.

Do you do you worry about that, who controls it, and how much these things will cost?

Absolutely, and let me give two versions of that worry that get to me. So one is sort of from a broader environmental and social point of view, which is just that these models, we're now running them for so long. So we can be running them for-- I think training a single model that you might use you might be running on tens of thousands of chips for months and months of time.

So that's taking a lot of energy, producing a lot of carbon. So there's a worry there at a broader society level as we escalate like this we're going to be using a lot of resources.

I also worry about it because we want to make sure that these models as we're building them really reflect the priorities across society, not just a profit oriented one that a company might have. And I don't want to demean that too much, I don't want to say that the profit motive is bad. Often the profit motive pushes firms in the same direction as society, that sort of White capitalism works, but not always. Right?

Sometimes they can be at odds or sometimes we can just say, There are public benefits we want, which are not reflected there. We want to make sure that other people can also be building these models so that we can have good performance on those things as well.

Do you feel like something fundamental has shifted here because I can imagine somebody thinking like, Hasn't cutting edge research always been expensive? Don't big companies always have? I mean, they always have deeper pockets than some scrappy startup?

Sure, absolutely. So I think here it's important to say, just as you're hinting at here, we often get this division of academia does some things and industry does some other things, and often that division is basic research versus applied research. So basic research saying, Oh, we're going to look at these underlying properties of the world and figure them out and then someone else is going to apply it. And applied research tends to be a lot more expensive, and so that's usually a good divide that exists.

But one of the challenges there is with AI is that so much of the progress with these models comes from making them bigger that you can't have that sort of neat divide of, Well, academia is going to do this stuff, the basic research, because actually so much of that basic research involves what happens when the model is 100 times as big, or how do we make it 100 times as big. And so you still need the resources to get there, and so that division doesn't exist as cleanly as it can exist elsewhere.

OK. So given that, do you have a vision for what you'd like to see in terms of the development of large language models?

So what I'd like to see is to make sure that there is still capacity in other groups other than just industry to be able to build these really cutting edge models. For example, with these large language models that everyone is talking about like ChatGPT, what I want to make sure is that-- for example, academic speech training such a model, and maybe there's only one academic model or two academic models and people have to collaborate across universities or something like that--

But I want to make sure that there is this capacity. And so people may be familiar there's this new National AI Research Resource called NAIRR, which is being built precisely to provide computing power to academics, and so that they can actually be building models, but right now that-- so we've seen an important first step in them doing, that but the scale of it is still not quite where we need them to get in order for academics to be really doing all of the things that are exciting in AI right now.

So let me ask kind of a related question. I know that a lot of people at MIT and a lot of people at other places, they're working with businesses who are trying to think about, OK, AI seems very important. How do we incorporate it into our business? What's the most effective way?

And I wonder in terms of just what you've seen, you've got older established companies with deeper pockets, probably, trying to incorporate it. You've got younger scrappier companies, which, I guess, would have the advantage of their old ways are less entrenched and so maybe they can just start fresh being like, Well, let's just do this with AI, we've never done it any other way.

Give me a sense of the landscape you've seen in terms of implementation of AI into business processes, because it's hard to change the way people do things?

Absolutely. So as you're saying, I think very small companies there is this sort of like AI-born companies that are just starting out and the assumption is, they're going to use AI for a lot of things, they're not going to have that many people, they're going to try and harness it, and so we definitely see a lot of that.

I think, though what we see is, if you look across the economy some work that has been done by the Census Bureau and some of the folks working with them shows that it is still overwhelmingly big firms that are the ones that are really adopting AI, and so it may be surprising given the amount of coverage that these things get in the news, but it's only something like 6% of firms that are actually adopting AI at this point, according to the Census.

And that is overwhelmingly those big firms because of the resources, because it's so expensive to hire the people and do stuff. Now you can say, Well, OK, those big firms are doing it, but how successful are they? And this was an interesting case where we actually did some detailed work on a specific example of an implementation with a grocery store, and what we found was that there was a lot of promise, but also a lot of cost in actually doing this.

So I want to get to the grocery store example, but first, it's interesting, 6%. So that's a pretty small amount of implementation so far, but I also wonder if within the category of big companies which applies to everything from oil and gas exploration, to retail, to all sorts of things, if you see a breakdown between--

Like Amazon's a big company, but so is Walmart, so is ExxonMobils, so are lots of other places. Do you feel like the 6% that are using it are highly concentrated in certain areas of big companies?

Mhm, good question. So we certainly know anecdotally areas that are having a lot of influence like this. We are actually doing some work in the lab, it's not quite done yet, but where we're trying to actually figure out what part of the processes are going on. But, I think, if you look broadly across what people are doing, I think what you see is, first of all, there's a real concentration in places where there's a lot of data or a lot of interactions.

So you see a lot of customer service things because that's very expensive to have the people doing it, and so as much as you can interact, there's a lot of data there, so people do that. There are lots of prediction work. So the Amazons and stuff of the world, or the Googles, it's worth a lot to them to be able to predict what product to show you, what ad to show you, what result to show you. And so all of those cases, you get a lot of AI implementations.

So that's sort of on the one extreme, and then on the other extreme, you get what I would traditionally call scientific computing. So scientific computing might be the big oil companies trying to figure out where to drill, or people doing big scientific simulations of weather, or something like that. There also, you see a lot of AI adoption going on at that site.

Well, it seems like you're saying, the places where the stakes are high. If you can find another place to drill it's worth some investment, but it also sounds like some of the places that are already doing AI in-house, like I mean, I assume, Apple, Amazon, Google, all these places, this is a thing they're doing anyway. They might as well try to use it.

Oh, absolutely.

Yeah.

Yeah.

OK, so let's talk about the grocery store because I think this is a really interesting, like this just goes to the question of, OK, you're a big company, how do you implement AI or do you is it even worth doing? And just on that question, do you feel like a lot of people come to you and are like, Is this even worth my time and my energy, because it is time and energy and money?

So, I would say, perhaps not entirely for the good. I would say that so much energy right now is on the hype. I think it's more people are like, I'm going to do AI everything, and you have to say, Well, maybe it doesn't work for everything. And I think what's nice about the supermarket example is it shows that it can be valuable, but you really have to understand that it can also be high-cost, and there's a trade here. So let me lay out a little bit more of the situation.

Tell us a little about supermarket

Yeah. So this is a European supermarket chain and they have the problem that all supermarkets have, which is they need to predict how much to order, and this is a little hard. You don't know how much ice cream you should order if suddenly it's going to be a really beautiful day, and suddenly people are going to buy you out.

Something I didn't realize until I did this case, for example, that there's a huge spike in nacho consumption. If you happen to have a football game with your local team on that Sunday, for example--

Ah, didn't think about it but it makes sense.

But it makes sense, right?

Yeah.

Exactly. So there are all of these things that if you are the manager of one of these grocery stores you need to be thinking about. And so historically, we've had systems that have basically said, Well, OK, last year, how much did people buy about this time? And then and then the manager will make a little adjustment. So it might be like, Oh, last year we did people bought 20 cans of tuna.

Right.

I'm going to order that, but I've got a few more people in the area, and now I'm going to order 22 or something.

OK.

So they might make these little adjustments, but it's worth a lot of money to them to get this right because if they order too much, they got just extra stuff sitting on the shelf, and even worse, some things like bananas or something like that will just go bad, so that's totally wasted money.

You don't want a lot of shrimp sitting around for too long.

Exactly.

No.

And conversely, if you were to-- well, obviously, there's some sales you didn't make, your customers are unhappy. So this doesn't work and either way if you get it wrong. And so what this supermarket said is, Well, one of the things that AI is typically known for being good at is integrating a bunch of different types of data and synthesizing them into predictions, and that's a particularly strong point. So, OK, well, let's try that.

And so indeed they did try it, and they found a remarkable success. So the errors that they were making when they ordered 10 and actually only five was consumed, that getting it wrong portion, they dropped that by a third, which is a big deal in this area.

Right.

That's a lot of improvement, but it was also the system was so expensive that they didn't want to adopt it.

OK.

Now, what has happened since this sort of initial decision is they've gone back now and they've started to look at it in more detail, and in fact, when we did this case study, we also went in and did that analytics in more detail. So we did the economics of it, and what we saw was that, they definitely should have adopted it, but only some places. And the intuition behind this is because you do a lot of compute, because that involves a bunch of costs, there are some predictions that you're going to make that are just not that valuable and not worth it.

So that can of tuna, actually, it can sit on the shelf for an extra two days and it doesn't matter, it doesn't cost you very much.

Right.

And so getting that prediction right, you don't want to spend a lot of money to get that prediction right because having one extra on the shelf is the swing, the spare capacity that you need to do it. Conversely, the shrimps you were just talking about, or the fruit and like the bananas and stuff like that, all of that, it's definitely worth doing. And so what to me this example really showed was, first of all, the power of these AI systems to do this kind of work, which is really exciting.

But also that unlike traditional computing where we don't pay much attention to how much compute it uses because most of the stuff we do does not take very much, you actually do have to pay attention to it with these AI systems and think about the cost and the trade-offs there.

So do you think everybody is starting to learn in business that you have to think about, Wow, this could save us $50 million in lost groceries, but then on the other side, this will cost us $50 million.

[LAUGHTER]

Like maybe not, maybe that doesn't make any sense actually.

Exactly, and I think this is going to be most salient to the people when you might need to design your own system. For example, so maybe you look at all the stuff that the big tech companies are producing and you say, That's all well and good, but my problem is much more specific, and it might still be very valuable, but now you really have to build a big system on your own and that can be very expensive.

And just to put that in context, so they estimated that GPT-4 cost OpenAI more than $100 million to train. So these really are very big numbers.

So when you step back and companies across sectors come to you and let's say they're like, We're very excited. AI it's very important, our CEO said we should get on this right away. Do you sort of step back and think with some companies there may be no problem where the cost of AI is justified in solving your particular problem, or do you feel like, No, no, it should be everywhere, at least somewhere across the economy.

So I think for most firms there will be applications, and for many of them quite a few applications that they can do. I think the question is going to be more in two different directions. So one is going to be like, Is it every problem? And I think what our research says is that, If that problem is hard enough, if it's specific enough, then you actually have to think very carefully about whether you want to do it and really, really do the math.

I think the second area where firms need to think about this and why they're often going to have at least some adoption is that we can get these partial adoption cases. So we did another case study that exemplified this kind of-- well, and I thought it was very interesting.

So this was an insurance company, a car insurance company, and they had a system whereby one of their clients got in an accident, they would take a picture with their cell phone and just send it in. So nobody actually goes to see it, they just send in the pictures and an assessor looks at it and says, OK, you need to replace the hood, you need to replace the bumper, so on.

And they wanted to develop a system that could do that in a fully automated way and found that they couldn't at a reasonable cost. Again, they had that situation, but what they found they could do is produce a system that was largely accurate but not enough for their whole thing, but they could use that system to then pre-fill a form that the assessor was doing.

So now the human assessor would look at it and be like, Yep, the hood. Yep, this, this. And that was all filled in for them already and they're like, Oh, no, no, no. The tire doesn't need to be replaced, this other thing can happen, or something like that. And so it made the assessors much more productive.

And so that sort of partial automation or sometimes called augmentation, I think, can be a very powerful thing that a lot of firms are going to have to think about, and the reason why that ends up being very possible, even if the full automation isn't possible is because for these deep learning systems there's this remarkable thing that getting that last bit of performance is incredibly expensive.

And this is sometimes discussed in the context of what's called an AI scaling law, but it's this idea that you have some system and maybe it has 30% error, and if you wanted to have that, so if you want to get down to 15%, You need to expand your computing and you might need to expand your computing by like 4,000 times. OK. And then if you want to go from 15 to 7 and 1/2, you have to expand it another 4,000 times.

Right, right. You might think you just have to double it but, No, no, no, the stakes are high.

Exactly. Exactly. And so consequentially what that means is that very high performance, those very high performance systems, are incredibly expensive.

So I know you think some about the future of jobs and organizations, and I wonder then if you think that the integration of AI into the business world, is that going to be like an earthquake for not necessarily jobs existing or not though, you can say, Yes, it's going to be an earthquake for that, but also like the nature of jobs, or is it going to be like a tremor?

[LAUGHTER]

We're on the Richter scale.

Yeah, Yeah. So this is definitely a big change, but I also think that some predictions are often saying, Well, like what could be done, and that's not quite the same thing as, What will be done. And this is a big actually a big area of work that my lab is researching at the moment because we're trying to understand, well, just because a human uses a vision to do something, could a computer vision system do that same thing? And the answer is like-- Sometimes the answer is, Yes, but the cost would be too high.

And so we're actually going into some of the data that the Bureau of Labor statistics and some of the other parts of the US government have to understand actually these tasks and how much value there is in replacing them, and what we find is that many, many tasks, individual firms are not going to want to automate them themselves. And so that means that either those things won't get automated, or we're going to have to have a big shift in the overall economy.

So it's going to have to be-- someone's going to have to build a big cross-cutting system that works for everybody. So if you imagine-- let me give you a very concrete example.

OK.

So think about a bakery, so if you look at the tasks that bakers do, one of the things they do is they look at their ingredients and they say, Are all of these still good? Should I put them in my food or should I not? And that's about 6% of what they do. So it turns out that if you think about 6% of a baker's time, and you maybe have five bakers in a little bakery, it might be worth about $14,000 worth of time in order to replace that with an automated task. Building your own system to do that is way, way more expensive than $14,000.

Right.

And so you don't that individual bakery doesn't want to do it, and so that's the part where you say, Oh, actually some tasks may not be adopted right away. But there's also the other side where you say, Well, if some bakery was able to figure it out and now offer a system where everybody has a little camera and they're doing it, maybe you could automate all bakeries at the same time. And so, I think, that to me is a really interesting change that the economy is going to have to go through in order to actually get the benefits of automation.

Well, I mean, does that bakery take credit cards, but they did not build the system there's no way. It'd be way beyond them. They're OK with being like being told, Here's a machine, you run the credit card through this or whatever, people tap, but they're not building that system.

Exactly. And so in just in the way you said, I think people will in fact adopt these platform systems or these AIs or service systems, but it's also we're saying, It took a long time for stores to adopt these credit card things. Many took a long time, and so that will slow down the adoption of AI versus some of these predictions where people are just saying like, Oh, my goodness, there's just a huge number of tasks that will be replaced right away. Some will be replaced right away, some of them you're going to have this diffusion process.

It also, just to go back to what you had said at the beginning that about 6% of companies, I think, are using AI, there can be a big difference between first movers and everybody else. Just because there were word processors, did not mean that all medical records got switched from paper to digital in a day. In fact, decades and decades, and I would venture to say there are probably some still in paper.

Yeah, right. Absolutely.

And so it sounds like-- I wonder if you also think there's going to be a lot of companies that are not even close to being the first movers here, that are going to take decades to change.

Absolutely. And I think the question is just going to be like, What does that mean for the overall economy? So I think if you have a little corner store bakery and they make your favorite apple pie or something like that, it's not going to be a huge incentive for them to adopt that in any short period of time. The incremental value is not that high, but I think what we should expect is that big firms, the ones that have scale, you should expect them to be moving pretty quickly on this.

Let me ask you a big picture question about global competitiveness in computing. There's a great quote you told *Politico,* here it is. "A huge proportion of the algorithms that have pushed computing forward have come out in the United States. Many of the biggest supercomputers have been here that overflows into all these other areas of society and gives them benefits. We're actually really losing that lead." Really?

Really. Unfortunately, yes. And let me say for a moment here, I want to speak more broadly than just, because I think people are very aware that there's a ton of exciting work going on in the United States, although also in places like China, but let me set that aside and just think of computing more broadly because AI is just one part of it.

So if you go back to the early days of computing in the 40s and 50s, people were starting to develop these systems, and as they did it, they realized, Oh, I have to build these algorithms, and you think of these as like the recipes that the computer is going to use to solve these particular problems. And what's amazing is a huge number of those important first algorithms were developed in the United States and many of the improvements on them.

So one of the big things that computer scientists work on, particularly theoretical computer scientists, is this question of, How can I do this thing more efficiently over time? And so it turns out that improvement is actually very, very big. So I think when people think about the progress in computing they almost always think of Moore's Law and progress in the hardware. In fact, algorithm progress has also been incredibly rapid and has also been very important for this.

But for the United States, what we saw was that in those initial days when computing was primarily in the United States, there were many people who were just working in the United States, who had been born in the United States, who were doing these things. But in the coming decades what happened was, the United States brought more and more excellent computer scientists and mathematicians from elsewhere in the world to the United States to do these things.

And so that was very exciting and produced a lot of benefits that we had, and as you mentioned also, this is also true in supercomputers. We also built some of the biggest supercomputers, the Department of Energy did it, for example, to safeguard some of the nuclear stockpile and things like that. So there was these very big ones, but in all of those areas, we see that China is investing a huge amount.

They have, obviously, more people who can work on these things, but they're also just investing a lot, and we're not investing enough as a country in order to maintain a lead in these things. And so China is continuing to close this gap, and we see that in some important algorithmic awards they've started to win in terms of some of these lists of the biggest supercomputers called the Top 500 List.

You can see that their presence on that is in some dimensions actually eclipsing the United States already, and so we should really worry about that, because when we make these cool discoveries that overflows into the companies that want to use these things. We discover a new way to model air flow, and the people who make jetliners, or the people who build long distance trucks, or something like that, adjust what they do to become more efficient.

Right.

And that happens in many, many different dimensions, and we're just not investing enough to keep the lead and being able to have those things.

And for people who are like, But, I mean, look at Silicon Valley, aren't we kind of running the world when it comes to tech? Are you saying like, Yeah, but that stuff is based on a previous generation of America being ahead, that does not predict the future.

Yeah. What I'm saying that it makes sense that Silicon Valley is all that it has been because we have companies that have been there for decades, that have been benefited from these advantages and so there's been a huge lead, but it's also the case that if China makes some of the key discoveries that are the next generation of computing, that's going to give them a huge advantage in those areas.

And so let me actually be a little bit more concrete. So people may be familiar with Moore's Law, is this idea that as we miniaturize the parts of our computer chips, we can put more parts on them and that makes the chips more powerful, and that that process has been going for decades and decades and made computers much, much more powerful.

But that ability to miniaturize we're running out of our ability to do it, and so as we run out, the question is like, Well, what is the successor? And so parts of it, it might be quantum computing, for part of it it might be optical computing. There are other technologies like neuromorphic that people are thinking about, and they're also just a lot of other materials. So people may have heard of spintronics and stuff like that. So there are a lot of possibilities.

Probably no single technology is going to be the successor for everything, and so the question is like, Well, whoever gets there first in some of these technologies is going to be able to shape what that is, and there are companies are going to have a big advantage in the computing that they can do in the same way that we see today with the people who can get GPUs and who can run their stuff faster.

And the problem is that right now in a lot of areas China is leading in these next generation technologies, and so that means that next set of things, if we don't get there first, actually could matter a lot for our advantage.

So it sounds like your advice would be the government should invest a lot more into computing in some ways in the way China is doing.

Yeah, absolutely. And I mean, I think we already see some good first steps in this with the CHIPS Act recently where we see an investment in making sure that we have chips and can secure a supply, that's definitely very important, but we are way underinvesting in these next generation systems.

I think it's a little hard because for so many decades this sort of model of Silicon just miniaturizing it has worked so well. It's hard to be willing to invest billions or tens of billions of dollars, but even tens of billions would be low compared to the benefit that society would get if we can actually get them.

A final question about quantum computing, which I know there's some promise, but there's some maybe skepticism that should be applied. I wonder if people should be thinking, this is coming right around the corner, that will be implementing this all over the place, or what?

Yeah. So I don't think it's right around the corner, and I don't think it's going to be very general, which is going to be important here. So what I mean by very general is like when you buy a computer you upgrade to the new one, you just assume that the new computer will do all the stuff that the old one can do. Right? It isn't going to be that. You're not going to be able to upgrade from the classical to the quantum and be able to do everything you could do before. It's only going to be a subset of problems you can do.

And that really arises because counter to I think most people's expectations about quantum computing, quantum computing is much slower than classical computing.

OK.

And you say, Well, if it's a lot slower, how the heck does it do better? And the answer that is those algorithms we were talking about earlier, that the recipe for doing it. And it turns out that quantum computers, because of their quantum nature, have access to a broader set of algorithms than a classical computer, and sometimes those new algorithms they get are dramatically faster.

And that means that even if the classical computer has a speed advantage, the quantum computer has an algorithm advantage, and which one of those wins, is basically a race. And we call this the classical hare versus the quantum tortoise.

OK.

[LAUGHTER]

And what we see is that indeed in some cases, if that algorithm benefit is big enough, you can get it, but what we see is that it's not at all a general thing. It's going to be specific problems that have these good algorithms, and even for those problems, those problems are going to need to be big enough that there's enough of an advantage that comes from the algorithm.

And so it's complicated?

It is a little complicated, but bottom line, for most of the things that we do, in the near-term quantum isn't going to get us there. There probably will be a small number of things that if they can sort out some of the engineering issues, will be very, very promising, and the key is going to be understanding which category different problems are in and we're working on that right now.

Neil Thompson is the Director of the FutureTech Research Project at MIT. He also has appointments at CSAIL and MIT's Initiative on the Digital Economy. This has been a super interesting conversation. Thanks.

It's my pleasure.

[DRAMATIC MUSIC]

And if you're interested in knowing more about the CSAIL Alliances Program and the latest research at CSAIL, please visit our website at cap.mit.edu. You can grab this podcast on Spotify, on Apple Music, or wherever you get your podcasts, plus, if you want to learn about our new digital course coming spring 2024, it's called Driving Innovation with Generative AI, you can find details on our website. I'm Kara Miller, the show is produced by Matt Purdy and Nate Caldwell with help from Audrey Woods. Tune in next month for a brand new edition of the *CSAIL Alliances Podcast,* and stay ahead of the curve.

[DRAMATIC MUSIC]