



AI INVESTORS' FORUM | NEXT GEN AI LEADERS: THE FOUNDERS BUILDING THE THIRD WAVE

Dominic Soon 00:00

Good morning, everyone, and thank you to the Milken Institute for bringing us together. I'm excited to be with four founders at the frontier of AI. My name is Dominic Soon and I'm with GIC Singapore Sovereign Wealth Fund. I'm part of our private equity group where we oversee our technology investments. It's hard to believe that it's less than four years since ChatGPT came out, and we are already moving through multiple waves in the AI industry, the first being chat, the next being coding, and now everyone's talking about agents. I think an agent needs a lot of things. The agent needs the ability to take action, so we have Russell Kaplan from Cognition AI. It needs to be able to speak, so we have Mati Staniszewski from ElevenLabs. Probably it needs the ability to see as well, so we have Alex Mashrabov from Higgsfield AI. Behind that, we have the self-improving brain, so Sara Hooker from Adaption. Now, all four of them are joining me on stage. So, in some ways, we're one big super-agent here on stage. We've kind of put this all together. The oldest company on the stage is four years old; the youngest is just founded this year—Sara's company. And yet, their products are already at work inside banks, inside automakers, governments, and so on, so forth, and working with many of the world's biggest brands. It's kind of incredible to see that within such a short period of time, we are generating such incredible value. So, maybe I'll just start by talking about the adoption that you've gotten so far. I think this year we're seeing really great traction across all of your companies, and what connects you guys is how fast the companies have adopted what you built. So, maybe I'll just go through and ask you to tell us what do you do, and why has your product gotten such great traction in the industry? Maybe Russell, we can start with you.

Russell Kaplan 01:56

Great. Thanks so much for having me, and it's great to be here with you all. I'm Russell from Cognition. So, we built Devin, the AI software engineer. We incorporated January of 2024. We launched Devin in March in 2024, and at the time, it was the first kind of autonomous coding agent where every Devin has its own computer. So, less of a copilot-style assistance workflow, more of a kind of a coworker flow. And when we started, it was at the edge of possible. We put out this demo video that went super viral, and half the world was like, "This is the most amazing thing I've ever seen." Half the world was, "This is totally fake. There's no way this can actually work." And it was a very vigorous kind of debate. And what we learned relatively early on is that cloud agents in 2024, the models were just barely good enough to make it work. And so, our first customers, they actually tended to be really large

organizations that had tons of code, where you could find sort of repeatable workloads on some subsets of this code, to do refactors, migrations, and so on. And as the models have gotten better and as Devin has gotten better, and the infrastructure's gotten better, now we see actually, for most software engineering, the right way to do it is with a remote agent. And so, yeah, in terms of the traction, we had our first customer in September of 2024. And just September of this year, we announced crossing \$1 billion of revenue run rate. And we tend to work pretty deeply with a lot of the really large organizations of the world, and so it's been exciting to see so far.

Dominic Soon 03:23

Exciting. Thank you. Sara, I'd love to get your perspective. You just started this year, but already you're getting lots of exciting traction. Tell us a little bit about your company and what you've been doing.

Sara Hooker 03:36

Yeah. I think it's very rare that the problem that drives you meets where the ecosystem feels as a whole. So, we started Adaption because we believe frontier AI should be owned outside of frontier AI labs. So, much the same way Devin supercharges code, we build models that build models, that supercharge how you own intelligence, whether that's automatic harness configurations or it's automatically training your own open weights model so it's customized for your data. I think so much of it is this moment over the last few months. I think that fickle access to models, with Fable shutting down and changing terms, as well as token prices of agentic have completely changed the way that people see how they want their internal ecosystem to be. At the very least, people want to hedge risk, especially with upcoming IPOs and uncertainty around token prices. So much of the momentum, even since we started in February, is really that realization, that most companies want to have a heterogeneous environment, and increasingly see their context as their IP and the one surface that they own, and so they should be leveraging it and self-improving on it as easily as possible.

Dominic Soon 04:48

That's great. Thank you. Maybe turning over to Alex. I think a lot of us have probably seen those really cool videos generated by Higgsfield, but I think many people don't realize you guys have gotten incredible enterprise traction as well. Tell us a little bit about Higgsfield.

Alex Mashrabov 05:03

Absolutely. So, Higgsfield today is a leading video AI company, and fundamentally, we help to bridge the gap between idea and video. Video has become the main language of the internet, and it's so important for creators, but mainly for businesses to be able to tell their story and to sell more products through videos. And historically, I have seen over the 10 last years of working in the video AI space, that most of the businesses never had the production capabilities to really make compelling ads. But also, it was historically very difficult to hyper-personalize messaging across different audiences, because it just requires so much shooting, and it is just virtually impossible. So, we are really excited that at Higgsfield, we are actually bridging the gap, and today most of our adoption is really coming from businesses who are making more and more money through these AI ads. Very important to say that there are some exciting stories, like where the customer starts with just spending maybe \$100 a month, and then six months later, they're able to 10x their business, and they also do this through Higgsfield, and now they spend over \$6 million a year.

Dominic Soon 06:16

Thank you. Then finally to Mati. I think what we just found out, actually, is that the voice of God that you hear at the end of every session is an ElevenLabs-generated voice. It's not a very well-spoken gentleman. Mati, tell us about the traction you're getting at ElevenLabs.

Mati Staniszewski 06:32

No, thanks for having me here, and also it's a pleasure to learn that voice of God is one of the ElevenLabs voices. But as a company, we are building the AI interaction platform that will help transform how all of us communicate with technology around us. And behind this, we build research—so build all the foundational models across audio, from starting the model four years ago that crossed that uncanny valley of speech and finally sounded human, from text to speech, to transcription models, to then, of course, orchestration models of how you combine the audio device with the intelligence together. And then we don't stop there. We build products on top of those models that help you transform either the creative space and help bring the brands to life, or help transform the customer experience space. And we do that by combining a lot of the model work, then with the knowledge, with the context that's required with the integrations, and the deployment layer of how you actually bring that to production with the right guardrails, right safety, right monitoring. And through that, we help power some of the iconic brands out there, from Revolut through Stripe to T-Mobile, who are looking to elevate their customer support, their sales, their operations. And as a company, recently crossed \$600 million in ARR, and more than half of that is coming from enterprise. And to close out, it's also a special day for us because we just opened our Singapore office yesterday, and so it's a pleasure to be here and spend more time with the community.

Dominic Soon 08:01

Very nice. Congratulations. Thank you. Right. So, those are the easy questions out of the way. Let's talk about a topic that's very topical right now, which is trust and control. Maybe I'll start with Sara first just to mix things up a little bit. There's a lot of questions that people are asking about how we can get reliable AI. You talk about systems which adapt themselves, which means that the system I use today is not going to be the same as the system you use tomorrow. So, how do we think about having a system that's reliable in that kind of situation?

Sara Hooker 08:37

I think that owning and having a self-improving AI within your context allows for you to have way more ownership, in my opinion. I think so much of what we're seeing is dependence on proprietary APIs, where it's really arbitrary how your data is processed, how your data may be, and workload may be synthetically imitated by some of these companies. In fact, I almost see self-improving AI in your context as the one thing that proprietary API companies don't have. So, if we think about like what's about to happen in the next year, I think so much of these companies are going to move upstream and go into certain verticals, and having your own internal self-improvement on your context, those are the interactions that are most valuable right now. Simultaneously, I think there's a very important window where the same self-improving systems that can automatically improve your model performance can also be used to automatically configure all your security protocols to scan your systems, and this is another key part of having both security as well as having a competitive edge, where your legacy, your historical context should be used to work for you. And both are critical at the same time. But the same technology which enables rapid self-improvement also enables automatic scanning of vulnerabilities.

Dominic Soon 09:53

All right. Do you find that the challenge is that customers don't have the tools to get the safety? What is the challenge that the companies that you talk to have in managing safety?

Sara Hooker 10:07

I think that most companies are uncertain what safety they should care about. So, there's these big headlines that dominate the public conversation, most notably the Hugging Face incident. I think this has both injected anxiety into companies as well as—they're not sure, what does this tactically mean for me? The Hugging Face incident is very important to think about, but a lot of it was like a lack of basic safety protocols, including scanning traces. And so, anchoring it to a much more precise conversation of you can be automatically scanning, having taxonomies to self-improve and to analyze error rates. Once that conversation is more anchored, I tend to think it leads to much more precise interventions. And that's, I think, a natural movement from the mainstream headlines to some tactical things that actually part of what even your work enabling code that can scan automatically vulnerabilities as well as self-improving on your context, which is what we do, it can really empower you to have a much better sense of what your vulnerabilities are and to act ahead.

Russell Kaplan 11:10

Yeah. I would just chime in on that. There's different levels to your point on safety. There's literal cybersecurity vulnerabilities. There's the sort of enterprise guardrails and governance and regulatory requirements you might have. People talk about the existential stuff. I think that the most acute near-term issue, in my opinion, is really the cyber risk. We see a ton of different workload trends in how people use Devin for all sorts of software engineering tasks. And by and large, right now, the number one priority across almost every company we work with is tackling this massive backlog of low-severity vulnerabilities. The high-severity vulnerabilities people deal with and can get on top of but most large organizations today, they're sitting on hundreds of thousands of the P3 tickets, P4 tickets that they just don't have the capacity to go solve. And what you're seeing now with this latest generation of models is these models can chain together 70 different low-severity vulnerabilities in a row to produce a really kind of damaging exploit. And then you also see malicious third-party actors who are able to build their own open source-based harnesses to do penetration testing. There's news this week of some banks in South Korea getting hacked by—it seemed like—a single 20-something-year-old in China just ripping some open source, basically vibe code penetration testing, compromising the largest financial institutions of an important developed nation. So, I think there's a near-term acute security risk that we have a narrow window for people to go deal with. But on the governance side, I actually think AI, it's pro-governance. It makes it easier to comply with the huge complexity of regulatory requirements that organizations have. At Cognition AI, we work with a bunch of different federal agencies of the US government, and they have very strict requirements for when you make changes. Do you have tickets associated with those changes that document everything that are compliant with FedRAMP, IL4, IL5? And Devin just handles all of that workload for us at this point.

Dominic Soon 13:06

Right. Maybe just double-click a little bit about that. You alluded to this. I think you said before that Devin writes most of the code in Devin, right? So, I think for organizations who are trying to understand how that will work, how do you ensure that the code that is pushed out is safe? Does it require some level of human scrutiny to make sure that it's acceptable?

Russell Kaplan 13:27

I think I'd be really interested also in learning from you all on how you've managed to grow your company so fast, because I feel like we kind of have it easy in software engineering. Because in software engineering, organizations are already designed to be robust to sort of errant changes from one person who's making a mistake. We've got

continuous integration testing, there's code review process. There's all these existing controls in a software development life cycle that if you just plug in AI, AI has to abide by those same controls. And so, if anything, it's like the lowest-risk way to adopt AI in many organizations, I think is in software engineering, and that's one of the reasons that it's probably the fastest growing workload in many different respects. Because you're just writing this code and then it's subject to all the same guardrails and we can kind of programmatically guarantee that Devin is not taking any write actions without you approving. One of our first customers in the US was Goldman Sachs, and they're a super-regulated, brand-sensitive large organization. But they can apply all the same controls they've spent decades building up, and I think that really lowers the risk.

Dominic Soon 14:33

Great. Maybe I'll turn now to Mati and Alex. How does this safety apply when you guys are working with media models, which obviously is quite a different thing?

Mati Staniszewski 14:45

First, I'll add one quick thing to what Russell mentioned. Of course, similarly, we build models. We're an engineering-heavy company, and very easy for us to adopt a lot of the technology and move quickly. But what we found is that as we move to the other teams, as we move to the legal team or the finance team or the talent team, how do you diffuse a lot of that innovation efficiently there? So, we have effectively engineers embedded in each of those teams to help adopt AI, but also helps train a lot of people to build that technology across all parts of the stack. And that really helped us as a company where everybody now is effectively a little bit more of a builder and a software engineer as a result. From a safety perspective, it's an essential part as you think about deploying any voice agent work or any voice work. First part, voice is an identity, so you need the right provenance or right information, whether the voice is how it's created, who created it, and be able to trace it back. So, that's where we invest a lot of the time on the voice level. And then on the agent level, here too, whenever you deploy that across the organization, you need the guardrails. You need the safety so it behaves in the right rails. You want to simulate it to check how it behaves across different scenarios. And then when you deploy, you want to monitor that across. And we have the pleasure of working with a number of governments, whether that's work in Ukraine government and helping them deploy how citizens can get access to information on what's happening in the frontline or how they can get access to health care and education. And across all of that, heavily domain-specific requirements to make sure that this information is passed in the right way and a very time-sensitive way. And here, that agent part is essential to stay on guardrails. So, very much a similar set of techniques to make sure that you do that well. But we'll see probably a set of developments of how those protocols across voice and voice agent space develop as voice agents start meeting another voice agent on the other side and need to swap to another way, which we are also working on.

Russell Kaplan 16:49

So, you're building a protocol for the voice agents to talk to each other?

Mati Staniszewski 16:53

Exactly. And then swap to a more efficient language so that the information loss is minimal.

Dominic Soon 17:00

That's very interesting. Fascinating. Alex, I'd love to hear your take.

Alex Mashrabov

17:02

Absolutely. So, going back to what Russell mentioned about these open-source harnesses, which pretty much enable anyone in the world of having a team of 20, 30 semiprofessional, professional penetration testers, this is quite dangerous, frankly. So, what we're doing within our organization, we are building a continuous deployment process and doing automated penetration testing using various models. And we also have seen that closed-source models and open-source models, they have quite different set of behaviors. And this is fairly easy to do. I think it took us less than a week to establish a good baseline. But then really, the devil is always in the details. And while we can still uncover new medium- and high-level issues very quickly, chaining all these low-severity issues in a chain so that the impact is actually large, I think this is very dangerous and I think many organizations will have to go through this transformation. So, I do believe that cybersecurity and AI is going to be completely redefined. While initially I felt like a lot of these articles may be self-serving for some of these companies, based on our own experience, they're going to be a real transformation of cybersecurity space. Going back specifically to video, there have been very important pivotal moments around March this year where, thanks to the European Union AI Act, all the contents on social media, which is AI-generated, requires very specific labeling, which increases the level of transparency for everyone. And this also helped us to 5x our enterprise adoption since then.

Dominic Soon 18:58

All right. Maybe just to talk about another aspect of trust, which is people-facing AI-generated content. Nowadays, we see the AI label on content. In your experience, does that impact the level of trust that people have in the media, or do they feel like there's lower authenticity, or do most people just sort of don't care, and what matters is really how good the content is?

Alex Mashrabov

19:22

Look, there is a scientific response, and there is a funny response. I'm going to start with the funny response. So, we all have seen how AI-generated dramas took over the world. Chinese AI-generated dramas—it's the most popular media today, I guess, on social media. So, clearly, people genuinely love just good content which resonates with them, and drama is just the most popular genre. But if I were to be more serious, I think the transparency and making sure that AI-generated content has marking is definitely important to establish trust. And there are different types of markings. So, one is obviously whenever content is published on social media, then the publisher has a responsibility to do the marking. Separately from that, there is attached metadata to that. And also, there are invisible watermarks, which are embedded into a video. I think all these combined definitely help to increase the level of trust in the synthetic media space.

Dominic Soon 20:19

Got it. Maybe just shifting gears a little bit. We've talked a little bit about trust and some of the concerns of AI. Maybe I wanted to get a little bit of a sense of what the frontier will look like if you think about what possibilities could be in five to 10 years. Maybe I'll start from the other end and come down. Mati, you kind of maybe alluded to agents talking to agents, that that would be an interesting idea. What do you think the future for voice agents will be in five to 10 years?

Mati Staniszewski

20:50

If you think about let's say the last 50 years, all of us would learn the technology's language. You would learn how to type on a keyboard, interact with a computer, or you would have a screen in front of you to interact, or you

would learn coding languages to be able to code the software program. And now, for the first time, with all the innovation of AI, we can shift that around and actually teach technology our medium of communication. So, the natural conversation, a voice conversation, will be a big part of how we interact and interface with the digital world, with the physical world. And what's fascinating about this is that we adapt technology to how we like to communicate. And I think that's going to be a big theme as we look—already now in the digital world—how you interact with the brands, how you reach out for help, and how we think from citizen support to the brand interaction, all of that will have a huge component of voice. And then in the next decade, that will increasingly shift to the physical world as well. So, as you think about the devices around you, as you think about the automotive space, as you think about the robotic space, all of that part to actually bring that and diffuse that across the economy will require that interaction side. And for me, this is a fascinating piece because there's so much work happening across intelligence that really to bring that intelligence to everyone, you will need that layer as well. But to actually do it well, it needs to be extremely localized to every given language, to every given person around the world. So, you do need—here in Singapore, you have four official languages, whether it's Malay, Tamil, Mandarin, and in English. You need to have an ability to switch between those, and sometimes combine those in any given interaction. It needs to have different dialects, different styles of conversation. So, as I think about the future, I think the bottom line is shifting that technology back so it works on our terms, the way we like to interact, which will be a fascinating help of hopefully immersing ourselves back more in the real life around us.

Sara Hooker 22:49

We need Adaption.

Mati Staniszewski 22:51

We definitely need Adaption—and Cognition, too.

Dominic Soon 22:55

Well, Sara, tell us more. How will Adaption push the frontier?

Sara Hooker 23:01

I don't know. I think it's a very powerful shift. So, the last 10 years, I was at different frontier labs. The whole effort was to build the largest model, and then you would hand it to someone else, and they would serve it, and then that would fuel how everyone uses technology. And now, so much of what the crux of progress is, is how do you interact with the world, and how do you self-improve from that interaction, which changes the dynamics of who owns progress. So, it's the person who—it's the companies that own product surface that actually have the most valuable, interesting data. That's why you see all these frontier labs pay so much money for simulation RL environments because they weren't structured to have the product surface. And on the other side, I think what will happen with these companies that have actually data about how people want to use technology and have been serving communities for a long time, I think increasingly, the view that motivates me and drives me is these companies have traditionally been terrified of ownership because it's gone wrong, or they did a one-year project, and it was a crappy model at the end. And I think there just needs to be way more efficient, immediate self-improvement from that data. And I think the view is that the same way you probably park your workload somewhere and you'll compute because you're just starting to use a model, that self-improvement should just happen in the background and just get better. And I think it will increasingly be that you just choose where you place compute, and you place it with the—you have your own internal algorithm that self-improves over time. And that's something I'm excited about because it gets to the core. I was being cheeky, but it gets to the core of the

vision, which I think is so important, that it's a big world. And I think that there's many different ways that people want to use technology, and we are in the era of interaction. It's how models actually do real things that have value in the world.

Russell Kaplan 24:57

Yeah, I think this is a super interesting point on specifically the question of where is the frontier going? Because you could argue that we're already starting to see this in a lot of domains now, that the frontier has essentially surpassed the level of capabilities needed to solve the domain. And we see this all the time in software engineering, which is the set of tasks which can be done by any model is growing enormously. And so, versus a year ago, with Devin, if we weren't using the most frontier model under the hood for almost all the workloads, the quality was measurably worse. Versus today, we have a harness called Fusion, where we basically automatically handle the routing between the most frontier models and then kind of commoditize everyday models that are much more price-performant for the workload. And with no measurable quality decrease, you get roughly 60 percent lower cost to date. You sort of automatically balance the model routing. And I think you saw this first in chat. When ChatGPT originally came out, you could really feel it every time a new model came out, the experience got better. It was more interactive. And this was in the metric. The DAUs and WAUs went up every time they shipped a new model. I think it was sometime around ChatGPT-4 that that basically stopped correlating, and they would ship smarter models, and you would no longer actually see the DAUs and WAUs go up because the models are smart enough to have a conversation with you. And I think that phenomenon is starting to happen to more and more fields, at which point, once all the models can do a job quite well, then you hear about exactly what you're talking about, which is all the optimization, the personalization, the customization to your specific data and context.

Dominic Soon 26:28

Right. I'm curious, within Cognition, are you using models in a way that would surprise anyone who is not in Cognition?

Russell Kaplan 26:39

I think we try to be really sensitive to cache-friendliness for model routing. I think maybe an underappreciated aspect of agentic workloads is that when you're doing a long-running agentic workload, you can very carefully construct the prefix of a prompt in a way to basically maximize the cache reads of a transformer inference pass. When you do that, it's much more price-performant, it's much cheaper, and so on. And so, I think a lot of the early days of model routing was this sort of naive thing where you get a prompt, and then you try to kind of classify the prompt as—is this a hard prompt? Is this an easy prompt? And then if it's an easy prompt, maybe you route it to a cheap model, and if it's a hard prompt, you route it to the most frontier model. The problem with that is a lot of problems, you don't actually appreciate how hard they are until you're halfway through the problem. And you're like, "Oh, God, this is actually way more complicated than I thought." And so we had to totally rearchitect how we do the harness engineering around these models to basically have multiple models running at the same time. We have what we call a sidekick model and then a main model. And maybe it's not quite as advanced as the voice communication protocol that you guys are working on, but we have a simple protocol between these various levels of intelligence models where the smartest one can kind of delegate the sub-tasks and then grade the work that comes back. And if you do this sort of handoff protocol in the right way, then you can keep a lot of cache-friendliness while really optimizing performance. And I think this is where coding is first in a lot of these ways, but I think this is where it's going for all sorts of applications of AI, which is why should the user have to pick the model?

For us, it's a full-time job to keep up with all the benchmarks of every model and understand specifically the strengths and weaknesses of each one. And I think that's only getting more complicated, and also more commoditized for a larger set of tasks. So, I think that's where it's going to go.

Dominic Soon 28:18

Thanks. No, that's a great hot tip if you want to keep your token prices down. Alex, on the video, I see these videos, and they all look incredible. I try to think, what could video look like 10 years from now, AI-generated video? What's your take there? How do you think it's going to be different in 10 years?

Alex Mashrabov 28:38

Absolutely. So, first, going back to the previous question, what we're seeing and truly unique is that engineering and creative—there is just a natural blend. So, the highest token spenders in our organization are actually the best creative directors. They were just so limited historically with all the tooling, budgeting, and so on, and just really with the need to organize physical production. And now with AI, they're not 10x creative, but maybe 1,000x creative. And that's where we are famous for seeing one person probably spending over \$100,000 a month just on building absolutely new workflows. So, this workflow actually allows to make editorial content, like tutorials, with a length over 10 minutes, end-to-end with AI. And that's just very funny to see the point of singularity, where AI can make educational videos how to use AI. But going back, on the personal side, it always felt like software engineering has had a great democratization effect across the world. And obviously, I think I and Mati can appreciate it—now it's possible to build company from different countries out of the world. It's not really centralized now to just one city in the world. And the same thing is going to happen with media. And one of the core bottlenecks is definitely still that these models are insanely difficult to prompt. Average prompt for a five-second long video could be over 3,000 words. So, what we're envisioning in years, there are going to be real-time interface, so that creative directors can just talk, and voice interface are going to be very, very important part of that, because this is just natural how humans interact. And then the bandwidth between idea and visualization is going to be 100x compared to what we see today. And we hope that then there are going to be a lot of new types of creators emerging, and that's especially important for, let's say, the category of artists with disabilities. Today, many of these artists simply cannot be heard in the world. But hopefully with this new emergence of voice-to-video-first interfaces, there are going to be new forms of media unlocked. And that's probably going to happen not in 10 years, but five, six years.

Sara Hooker 31:03

I think it's fun you mentioned prompt engineering, because I think prompt engineering is actually the symptom of people not having control and having to do acrobatics around systems that don't contextualize. So, I completely agree. It's a very worthy target to aim for.

Dominic Soon 31:20

Great. So, maybe just a quick lightning round. I want to talk about jobs. Everyone's concerned about jobs that will be created or jobs that will be destroyed. Maybe if we go around, what is one job you think won't exist in the future? Or what is one job that you think will exist in the future, so five to 10 years from now?

Russell Kaplan 31:40

I'm happy to start. I think there are going to be so many jobs. I think we're going to be overflowing with jobs. I think this narrative that AI is taking all the jobs is so silly, and I can just tell you in the economic data directly that we see. I'll give you one example. One of the largest health insurers in the United States, they rolled out Devin for a bunch of workloads. One of the first workloads they did was they looked at the day of a chronic care manager for a health plan. These people—a few thousand people to start—they interact with patients, get them health care, help manage their plans, and they realized that they could automate about two-thirds of the workday with Devin for these folks. And so, they did that. And what they did next was, I think, a symptom of where this is all going, which is they realized, okay, these people are now three times more productive, so rather than firing two-thirds of the people, they lowered the price of their health-care plans dramatically, and they sold way more plans than they were previously. Those jobs got much more interesting for all of the chronic care plan managers, because now they're not doing the mundane parts of their work. They're interacting with patients, they're understanding how to help them, and the company made a lot more revenue as a result because they're selling a lot more plans. And so I think this is a common fallacy that the world's going to stay the same, and now we just have less work. In practice, what we see across the board is companies are getting more ambitious. And so, companies can do more. They are doing a lot more. And the companies that don't do that are going to be outcompeted by the ones that do. So, I'm personally super bullish on the jobs creation.

Dominic Soon 33:00

Exciting. I'm glad to hear it.

Mati Staniszewski 33:04

Happy to jump in next. I also echo that I do believe there's just so many parts that will get unlocked thanks to AI now getting all the admin, all the easy part out of the picture. But I think one of the things that will flourish is all the social settings where people can interact with one another, whether that's going to be education or social, whether that's the creative space where finally all of us will be able to create the things and stories that we always wanted to tell, but we can't, thanks to AI just unlocking and democratizing a lot of the tooling. So, I think all the effectively human-to-human interaction and the setups where that's needed will skyrocket in terms of the time we spend and the jobs that will appear.

Dominic Soon 33:49

That's great. Sara, Alex?

Sara Hooker 33:49

I think that two things can be true at the same time. So, I think jobs will dramatically change. I'm a computer scientist. I did my PhD, and honestly, most of the classes I took are kind of irrelevant. A lot of the skills I learned about how do you pursue an idea at the frontier are still very true, but computer science was the first to—we thought code was a measure of the hardest task, and so we basically tackled that first, ironically. But that doesn't mean that the nature of my job has changed at its crux. It just means that I spend my day doing very different things than I did even five years ago. I think that there's a second aspect which is coming across, I think what's been said, with Alex and Mati, which is that the nature of who can participate and who has access to throw their hat in the ring is totally different. The concepts that are coming up here are creativity, but think about it in terms of software. It's interesting, I grew up mainly in Mozambique, in northern Mozambique, and my mom is Irish, my dad's British. The fact that I made it to Google DeepMind, and then I led a frontier lab, and now I'm leading this, is a miracle in some ways. And I just think the democratization of tools like coding, it means that there's more people

who have the best idea who can play, and I think that's actually always a good thing. It doesn't mean that there shouldn't be anxiety about the dynamics at play. It doesn't mean that governments don't have an important role. But certainly, I think that the truth is, while the nature will change, I just really think it's as binary. And previous technological shifts also suggest it's just not binary.

Dominic Soon 35:40

Right. We should talk about all your plans in Asia. So, maybe, Alex, I'd love to hear your views on the topic and then share with us what are your plans in Asia going forward.

Alex Mashrabov 35:48

Just to very quickly summarize on the previous topic, I think it's very important to say that we're seeing already this transformation play out, and we already touched base on the coding space. For example, I guess at Higgsfield AI, there are less than 10 percent of engineering team who sort of manually code, and that's typically related to payments and very sensitive systems. For the rest, 90 percent plus, it's like almost all the code they produce is really AI-generated. And that creates a snowball of new issues and sort of new professions, including all the cybersecurity challenges, which are really in front of us. And the creative jobs really—our team is approaching 500 people, and we have over 200 people in the creative team, and we can also see that the role of the creative engineer and for deployed creative are really emerging in a big way, which would be impossible to imagine two years ago. But it's also important to say that a lot of VFX jobs today really require knowledge of AI, and that becomes a prerequisite. So, going back to Asia, Asia is a very special region for us, and that's specifically relevant for video AI. A lot of trends really originate in Asia and then move to the West. So, for us, it's very important to stay close to our customers and be the first to understand the trends, where there are new emerging patterns of adoption, making sure we can capitalize on it. And that's why we are so excited to be here in Singapore because the amount of government focus on figuring out AI strategy is impressive. But also, I was very fascinated to learn that the amount of energy and data center build-outs in Southeast Asia is, I guess, hundreds of megawatts and I think, right now, over gigawatts in just a couple of months. And I probably still don't know enough. When I came here, I didn't really realize that the infrastructure build-out is so massive. And obviously, it feels like that also the concentration of talent in Singapore is unmatched.

Dominic Soon 38:06

That's great. All right. Now, Sara, just half a minute. Plans in Asia, what are you planning to do in the region?

Sara Hooker 38:13

Yeah, I think honestly, we were global first from day one. So, we started in February, and I think we welcomed Darius, who leads our global strategy here in Singapore, in May. I think our mandate is that we want to empower adaptive systems, and one of the most critical places for adaptive systems is companies all over the world who have different jurisdictions, different requirements, different tastes, and that's incredibly exciting for us. How do we make AI actually work in the real world and connect with that last mile of how people like to use technology? So, I will also say Singapore has just very strong technical talent. It's a very interesting meeting place where I think a lot of Chinese talent and also increasingly companies from all over the world meet. And this isn't often talked about, but I think it plays a critical role in navigating that dialogue and forming those talent pools that actually accelerate knowledge share, which is really interesting.

Dominic Soon 39:17

Great. Russell, you have a great footprint here already.

Russell Kaplan 39:21

Yeah. We just moved to an even bigger office in Singapore. We're about 35 people in Asia-Pacific, and Singapore is our headquarters. The thing we've observed is that our business has been disproportionately international from the early days, and I think part of that is because software engineering capacity is not evenly distributed throughout the world. But every country has a lot of demand for software. And so, in some sense, Devin's almost disproportionately more valuable in countries with less software engineering. Now, Singapore has amazing domestic software engineering talent. We have an engineering team here, and we're expanding it a lot. But broadly through Asia, one of the things we've noticed—for example, we do a lot in Japan, and Japan has an aging workforce, and there's a lot of legacy systems that need to be modernized and maybe insufficient engineering capacity to go do that by hand. So, we see a huge amount of growth in Asia for what we're doing.

Dominic Soon 40:12

Great. And Mati, I'll just let you quickly wrap up. You did talk about your Singapore office.

Mati Staniszewski 40:16

So, similarly, global from day one. We hired remote. We hired some of the most incredible engineers here in the first 10 people that were at the company. But a lot of the work across Asia, very culturally diverse, extremely advanced in technology. So, many interfaces are voice and voice agent interfaces. And we are just excited for building in the region and showing the rest of the world how we can build innovation here and how we can inspire the rest of the world. So, my cue—we are hiring, so if you want to join, please apply at any point.

Dominic Soon 40:50

Exciting. Thank you so much. So, lots of jobs for Asia here from these four companies. While I still have my job as a moderator, we have to wrap up, so please join me in thanking our four speakers for this panel. Thank you.

Disclaimer: This transcript was reviewed by individuals for accuracy and serves as a reference. However, it may still contain errors or omissions. Please verify any critical information independently.