



AUTONOMY WITH AUTOMATION: MAKING AI SOVEREIGNTY PAY OFF

Announcer 00:03

Welcome to the Milken Institute Asia Summit. We are delighted to have you join us for this event and look forward to your participation throughout the day. To open the program, please welcome to the stage executive vice president and chief strategy officer at the Milken Institute, Jeff Valliere.

Jeff Valliere 00:27

Good morning, everyone, and welcome to day one of the 13th Milken Institute Asia Summit. It's not an exaggeration to say that we're gathering at a pivotal and historic moment for the world. What is driving it is best captured in two words: artificial intelligence. Today, AI touches every corner of society, from the energy sector to academia, from investors to the high-tech ecosystem. And a global race is underway to regulate this technology without stifling the innovation that makes it so powerful. As a result, governments are now grappling with what it means to build sovereignty in a world that is, at the same time, deeply interconnected and increasingly fractured. That last question, sovereignty, is worth emphasizing because it gets at something the Milken Institute believes deeply. AI leadership is not going to be won on innovation alone. The nations and institutions that lead will be the ones that can also marshal capital, build infrastructure, engage policymakers, and forge the right partnerships. Breakthroughs need financing and compute needs energy and chips. That intersection of capital, infrastructure, policy, and partnership is precisely what makes the Milken Institute uniquely suited to this moment in collaboration with all of you here today. So this week, we're committed to doing more than just unpacking these issues on our stages. We're going deeper. On Friday morning, we'll be convening our inaugural "AI Investors' Forum," bringing together the leading voices to turn conversations into action. I hope you will join us for that. But this is not just a one-week focus for us. AI, its implications, governance, and impact will be a defining priority for the Milken Institute globally, and especially here for our Asia center. Because the questions it raises are exactly the ones we exist to help answer. How do we mobilize capital, bring the right people together, and solve the hardest problems to expand prosperity? And that work begins right now on this stage. Our opening session asks a defining question for this moment. As governments, investors, and corporations move aggressively to build sovereign AI capabilities, who will lead, who will benefit, and on whose terms? The title is very fitting, "Autonomy with Automation: Making AI Sovereignty Pay Off." We have an extraordinary few days ahead. So, thank you again for being here, and let's get started.

Announcer 03:23

Please welcome the panel on “Autonomy with Automation: Making AI Sovereignty Pay Off.” Moderated by Charlie Campbell, editor-at-large, TIME.

Charlie Campbell 03:47

Good morning, everyone, and welcome to this Milken Institute session on sovereign AI. My name is Charlie Campbell. I'm an editor-at-large for *TIME magazine*. Artificial intelligence is no longer merely a technological race. It's becoming a contest over economic sovereignty. Governments are pouring money into chips, data centers, power grids, research, and talent, while private investors increasingly decide which national ambitions receive the capital to scale. Yet sovereign AI raises difficult questions. Must every country build its own models? Or can it retain control through open systems, trusted partners, and decentralized infrastructure? Today, we'll examine where genuine strategic autonomy lies, how public policy and private capital should divide risk, and whether the next AI power will be established by technology giants or ambitious newcomers. I'm delighted to welcome a very distinguished panel to talk about this topic. I encourage you to look at their bios on the Milken Institute app. Perhaps to kick off, Hans, what does sovereign AI actually require? Control of data, model weights, computing infrastructure, energy, talent, or everything?

Hans Dekkers 05:00

It's everything. So good morning, everyone. Thanks for joining. I'm from IBM, run the region here in Asia Pacific. And to answer your question, I think companies today, enterprises, are really searching for, I call it, control and flexibility. When you look at sovereignty, it's much more than data residency—do you control your data? It is also about, are you in technical control? Is the technology your choosing? Do you have choice? Do you have talent to independently run your infrastructures, your data, your compute? And ultimately, do you have the flexibility to shift? Because in the end, there's nothing worse than a stranded asset. So you want to make sure that you have the flexibility. We call it true hybrid by design. Can you move from one country to the next? Can you move from one application to the next? Can you move from one hyperscaler to the next? Would you move from one frontier model to the next? Would you move from a frontier model to an on-premise version? And this is really what ultimately sovereignty control flexibility means, and this is where enterprises and governments are really, really honing into.

Charlie Campbell 06:17

Yeah. Maybe Igor, could turn to you next. Should a medium-sized country attempt to build a national frontier model or concentrate on adapting open models to its own language, industries, public services?

Igor Babuschkin 06:29

Great question. So, I'm Igor. I'm an AI builder. I've been working on frontier AI research for the last 10 years. I worked at DeepMind, OpenAI before ChatGPT came out, and also cofounded xAI together with Elon. We got to the frontier and actually managed to build a frontier model within two years, which we're very proud of. Now I'm the CEO of River AI, a new AI company that's building a sovereign AI stack that any corporation in the world or any country can use to rapidly develop and customize their own AI models. And I'd say if you're a medium-sized country and you're capable of building the best model in the world, you should definitely do it. I think it's clearly AI is the future. AI gives you an incredible productivity advantage. We're now seeing that, and the knowledge gap between those that can leverage powerful AI tools and those that aren't leveraging them is increasing. So you can be up to 100x more productive in terms of software engineering today if you know how to use the tools and you apply them extremely well. So the power of AI is really unquestionable. But I think most medium-sized countries

will find it difficult to compete. So you need enormous amounts of capital. You need access to a talent pool that is quite small and very much focused in the US and in China. So it's very, very tricky. And then there's the whole infrastructure risk as well. So, are we building enough infrastructure? Are we building it in the right way? Is it going to be profitable long term? So, what I'm seeing is that a lot of countries are deciding not to engage in building these frontier models. I believe including Singapore, so there have been some statements from the government recently that they don't necessarily want to build their own chips and build their own frontier models from scratch. But that doesn't mean that you have to go to the proprietary AI companies and sort of rent your intelligence from them. In fact, it's really important that medium-sized countries own their own AI systems, and open-weight models make that possible. You can bring these open-weights and run them on-prem, in your own data centers. And that also allows you to customize the models. So it's not necessarily about creating everything from scratch, doing the pre-training, doing the large-scale RL training of these models. But you should consider customizing these models for your use cases, which is much, much cheaper, much more achievable. And that way the knowledge of these AI models and the ability to deploy them flexibly, to use them in a private manner on your citizens' data, stays intact. So it's highly recommended that many, many countries in the world engage in their own custom fine-tuning, reinforcement learning, tuning of these open-weight models.

Charlie Campbell 09:14

So you're saying no to the proprietary models such as OpenAI and Anthropic?

Igor Babuschkin 09:20

Yeah, I think that is quite interesting. Obviously, the proprietary model builders have had a huge advantage. They were the first to really leverage these technologies to build the best models, and even today, they're offering the best capabilities, the highest level of intelligence with their products. But you have to keep in mind that you're sending all of your data over to them. So even in the case where they handle this data securely and privately, there's still the possibility that they'll turn you off, that they won't give you access to the next level of model that will come out. And we've seen a few cases, in the world, even in the US where the US government got a bit nervous when Anthropic told them that they might not be allowed to use the model for some purposes. So I think there should be enough motivation for any country in the world to look into bringing these open-weight models onto their premises and to have teams in the country that understand how to leverage them, how to customize them, and use them for different purposes. And the good news is, if you have a narrow use case and you fine-tune these models on those use cases, you can actually outperform the best proprietary models. So that's something that we're increasingly seeing with customers today, and so there's really no reason not to do it. It's really a matter of having the confidence to engage in model customization, to try it out, to find the right partners to work with.

Charlie Campbell 10:41

Thank you. Firas, perhaps I can bring you in now. Do you think that open-weight models can generally narrow the gap between AI leaders and followers, or will frontiers continue to move too quickly for the latecomers to catch up?

Firas Sleiman 10:55

We have examples of both. Good morning, everyone. I'm Firas. I run a sovereign AI company in the Middle East, based in Qatar. I've spent 30 years of my life building e-government solutions across the GCC and North America with BCG before that. I mean consulting companies, BCG and PwC and IBM. So obviously, the question is difficult to answer because I think the race, from a capital perspective, is clear. There's the Eastern, maybe, sector, that is

focused on smaller models that are highly efficient, and open source. China is leading that. The Western world is going towards closed models that are hyper-engaged in terms of financing, and GPU-bound. So, which will prevail at the end is maybe not the most important question. I think the political ecosystem dictates who uses which side. And I guess us being in the Middle East, especially in the GCC, we're kind of in the middle. And so there's a lot of utilization on both sides. So we've seen governments, for example, just two weeks ago, Saudi Arabia launched one of their fine-tuned Arabic models based on a Chinese model, while at the same time having a large investment in GPU clusters and the deployment of Western models in Saudi Arabia. So you see the hybrid, where a lot of countries that are trying to invest in both, because they don't want to be part of that race and take one side versus the other. Also, the ecosystem around the models is really important. And also, as my colleagues mentioned, I think models alone are not going to help the transformation, in sovereignty, in governments. I think there's a whole—now the term is harness, I guess, a harness or a scaffolding around the models that is really important. And, since we work in sovereign AI, what we've noticed is the biggest hindrance towards that advancement is really the interoperability. It's the standards between those models, between the harnesses around them, and how will they talk to each other. Because governments, by nature, are interoperable. And we've spent the last maybe 30 or 40 years of our lives building interoperability standards for government, be it security, be it data, be it applications, APIs, and all of that work. So, the next phase is: how do we set the standards so that these harnesses and the models in them speak to each other, so that countries—within the country, each government agency can interoperate and transfer data and applications and maybe control in these orchestrations that are happening. And then also intra-country. So how do countries then collaborate and work across the globe together on these systems?

Charlie Campbell 14:20

Do you fear potentially a bifurcation of the AI ecosystem, where there's the Western proprietary models on one side and then the open-weight models on the other, and maybe countries will be forced to choose between the two, and therefore it'd become a bit more divided in that way?

Firas Sleiman 14:41

Yeah, we're seeing the divide. The example of it is in the UAE, for example, the sovereign company that is building all those frontier models, called G42. It's a large investment that was done by the government of Abu Dhabi. And particularly the United Arab Emirates, at a certain point, had both, were working on both, and eventually decided to go on the Western side and team up with Microsoft and build their cluster that way. And we see the likes of Saudi Arabia is doing the opposite, was fine-tuning their Arabic LLMs or large language models on Chinese-based models. So yeah, I don't know. I'm not sure where this is going to end up at the end, more than just maybe this hybrid world. I think there's also a lot of difference between those models. The smaller open-source models are meant to serve a certain type of use cases. You're supposed to manage them yourself as a country, and so the deployments in GPU infrastructure and large data centers is essential for those models. While, with the Western models, I think the natural inclination is for these large hyperscalers to come in and set up their data centers, the likes of Microsoft and Google and Amazon and now OpenAI investing in large data centers. So the models are coming with their infrastructure, in a way. So the country's not really investing. It's much more of a, if you want, a long-term play with some concessions, where these large companies are coming and building their own data centers, even investing their own money to build data centers just to have that footprint, and then to offer that model with government concessions in these regions.

Charlie Campbell 16:36

Fascinating. Lily, maybe we can turn to you now. Do you believe that sovereign AI requires new ways of raising capital beyond governments, banks, and big technology companies?

Lily Liu 16:49

I certainly do think so. I think one of the reasons why you see this tremendous concentration of AI progress in China and the United States is, those are the only two countries that have scale accumulation of both talent and capital formation. And so this is, I think, the most extreme expression of why having those two things being co-located at such scale gives you such a tremendous advantage. And when it comes to sovereign AI development, there's certainly the model layer, which we've been spending quite a bit of time discussing, and that is extremely difficult to develop. But I think there's also something to be said for at least sovereignty over the physical layer of AI. And that being chips, compute, data centers, electricity, and being able to build out that infrastructure in a timely fashion, which is obviously a massive CapEx investment. And I think when you look around the world, and if you just look at equity markets capitalization around the globe, it's about 150 trillion. Half of that is in the United States. About 20 percent of that is in Hong Kong and China. And then all of the other 180 countries are fitting into 30 percent, and the G20 together is about 95 percent of equity market cap globally. And what that means is—not just in AI, but I think it's particularly very extremely being expressed in this compressed time period with AI—that means there's about 160 countries around the world that essentially don't have the means to form capital around investment into innovation and in growth. That's being expressed in the ability to develop competitive models. That's also being expressed in the ability to stand up the underlying infrastructure to being able to host whatever is being done in AI within your own borders, and the sovereignty that comes with that. And what that means is you may not only just be inheriting the weights that have been produced in China and the United States, but then also a lot of that data may not actually be resident within your own borders. And we come from the world of blockchain. We've always been pretty ardent believers in open systems, and the need for open systems to be particularly applied in the case of pretty foundational infrastructure, and the belief that when you have open systems, then you open up access. It's certainly not to say that the end result is going to make all countries as competitive as the United States and China. I don't think that's a realistic expectation. But having broader access to not just the two, but then even the 20, and then the 180, I think is something that is—if you believe this is foundationally important to the future of economic development, of talent development, of human development, then there needs to be broader access. And I think that capital formation is a very important part of that.

Charlie Campbell 19:48

Absolutely. Hans, is there a problem with governments placing strategically important financial or AI infrastructure on public networks where they don't actually have complete control or direct control of these networks?

Hans Dekkers 20:06

I think the ability to govern your data, your models, your AI, is going to determine success of any institution, be it government or enterprise. I think there's been multiple examples in the last six months where you see rogue agents basically leading to very unwanted outcomes. I think we'll see much more of that. The companies and governments that are able to—some of this stuff is fairly simple. Make sure that you're open source, make sure that your applications are patched to the latest levels, make sure that password control, encryption, all of these basic elements that we know we have to do, make sure that they're implemented well. It's the first thing that everyone can and should be doing. Then the governance of it, the governance of the models, the governance of where your data runs. What happens if—to Igor's point—if you're cut off? What happens if a war starts? What do you do? Are you flexible enough? Are you in control to basically take the destiny in your own hands? And this is what many

clients are asking us to help them with. And that orchestration, based on predominantly open sources, is really where we see the future going. Yes, there is a use space, just to link to Lily's point, for frontier models. There's a lot of capital going into it. When you look at it, maybe 10 to 15 frontier models will exist in the world, inevitably, at some point in time. Everything else, in our view, will be the smaller models, the domain-specific models, the open models that will run on-prem for governments and/or enterprises. Everyone will have hundreds of them, if not thousands. And they're their models. They're their outcomes, and in the end, they will show up in the balance sheets of companies. This is where we see it going. So the hybrid world between frontier models, open-source models, on-prem models that are sovereign will be there as long as the governance of it and the outcome of it is protected. So it's not so much anymore, Charlie, about the building capability. It's really about the governing capability that will determine the success of enterprises and government.

Charlie Campbell 22:37

Igor, you touched on it already, but we saw recently that your government blocked Anthropic from passing on its most advanced frontier models, like Mythos, to foreign governments. What's the message that US allies should take away from that? The fact that they're going to base their AI infrastructure, their cybersecurity, everything, on a system where the US is going to be withholding the latest technology, the most advanced models. Is there an argument that Europe, Singapore, Japan, should be really doubling down trying to develop their own competitors?

Igor Babuschkin 22:54

Absolutely in the sense that any country in the world should make sure that they deeply understand the AI technology, that they understand the frontier, and that they're able to own and control their own AI systems locally. I think that's incredibly important. Like I mentioned, it will be very difficult to compete with the frontier. So the incredible velocity at which these AI companies develop and improve their systems, it's really unmatched in human history so far. So even just this morning, I saw there were some new open-weight releases from the US that looked pretty good. So even spending time here at the conference, I feel like I have to go back and continue to work on our models. And it's very difficult for any country in the world to match that. I think outside of wartime, we've never seen governments be able to move at this kind of velocity. So I think you have to be very careful in terms of how you invest your capital as a country as well. How do you use it effectively? How do you use it efficiently, to gain some sovereignty, to gain some control over AI systems? And with this unprecedented infrastructure build-out, there's also the worry and the risk, is this actually going to pay off for us long term? Which is also worth thinking about. A lot of the incredible speed we're seeing from OpenAI, Anthropic, and others also comes from the fact that they very liberally basically invest their capital to accelerate their progress. And we're going to see that with their IPOs coming up and S-1s being released, that even though there's been an unprecedented rapid growth in revenue at those companies, unlike anything we've seen before, they're still not profitable. And I think that should make anyone that wants to invest deeply in AI also think, what can we do differently to actually make this a sustainable effort for our country, for our business? So it means both being more efficient with your resources. So the GPUs that you're bringing up, you want to make sure that the infrastructure that you develop, that you deploy on those GPUs is very efficient. Likewise, you need to find those truly useful use cases that can generate revenue and leave margin for companies in the long run. And the big AI companies are not the example to pick. Often the best offerings, they are subsidized heavily. For example, ChatGPT subscriptions and so on. So you want to also innovate in terms of the business model around AI and how you effectively use AI in your country.

Charlie Campbell 25:49

That makes sense. Firas, can decentralized computing generally reduce dependence on hyperscalers? Or does the scarcity of advanced chips mean that power ultimately remains concentrated?

Firas Sleiman

26:03

Well, I would say, I think empirically the answer is yes. I think the proliferation of distributed computing right now is leading to surprises for the hyperscalers. There's a lot happening in countries that are focused on open-source development. So, for example, yesterday France issued a new model, called Mistral 4, and it's an open-weight, open-source model. And those labs are not highly invested. They're not very large labs, even though the government in France is spending money on this distributed computing network that they have for these models. And they're taking a lot of the heat away from the Western models, or maybe more the United States models, towards Europe, towards Asia, towards maybe even some of the parts that are happening right now in the Middle East. And there's three models, for example, in the Middle East. There's ALLaM in Saudi Arabia, there's Falcon in the UAE, and there's Fanar in Qatar. And those are small investments that led to national models that are being mandated right now. And the governments are telling government agencies, "If you want to ensure sovereignty, you have to use those models." And so it is those investments, even though they're smaller, but they're taking away and chipping away, if you want, from the hyperscalers' ability to penetrate these countries. Also, as policies come out, for or against the usage of international models, we're seeing the drift towards governments saying, "If we need sovereignty, we need to have not just our own model, but data sovereignty, and therefore our data cannot travel to other countries." And so, for example, Qatar, where we work primarily right now, has a data privacy law that says the confines of your data sovereignty are within the physical and geographical borders of the country. So the data has to stay, and even transmitting data to data centers that are hosted—say Google or Microsoft have their own data centers hosted in Qatar—you're not allowed to even transmit that data through the network of Google or Microsoft to, for example, Europe or the United States. And so it's putting pressure on these hyperscalers to invest in open-source or smaller models and then put those models in the countries. Otherwise, at a certain point, that switch is going to be off, and the government's going to say, "Sorry, you have to have the model running in the country, otherwise nobody's going to use it." At least the government's not going to use it.

Charlie Campbell

29:06

Absolutely. Lily, can neutral public infrastructure, such as data centers, help countries to retain control of their AI capabilities or accessing the global capital, talent, and technology?

Lily Liu

29:22

I think the neutrality of critical infrastructure here is—I think that's the key point here, that if you're going to have these open systems, then having neutral systems is part and parcel of having broader access. And when it comes to, I think, capital formation, then what we're inheriting here is how capital is distributed today. And frankly, just most of that goes through New York in one form or another. And so I think the vision that we have from—when people talk about frontier technology, for a number of years now, people have oftentimes talked about crypto and AI in the same breath. Clearly, all eyes are on AI right now and will be for some time. I think that's very much well-deserved. When you're talking about the need to raise essentially tens and hundreds of billions of dollars in very short time frames, there's really only two places you can do that, and potentially really only one place that you can do that. And so whatever neutrality may exist in terms of the underlying technology—and there's obviously sort of a lot of politics surrounding every layer of the stack here—but whatever may even exist in terms of open systems and neutral systems, you're still inheriting, essentially, how capital is distributed today. And I think that's why we see things happening as they are. And when you talk to the open-weight models, in private they basically say, "If I

had the market share that Anthropic and OpenAI had, then I would be a closed-weight model, too." So, it's just a function of competitive market share ranking, really, more than anything else. And it's really an instrument of being competitive and trying to gain market share to be open weight right now. And I think it's always going to be the case that the top one or two, maybe the top three in the world, are going to leverage that market position to be closed weight and closed source. And I think that's going to be the competitive advantage, where if you're open source, you can broaden access, try to pick up market share, and I think that's partially going to be the competitive dynamics here.

Charlie Campbell 31:55

Hans, obviously private capital rewards speed and returns, while national AI strategies require patient investment in infrastructure and talent. What's the best way to reconcile those two sort of strategies?

Hans Dekkers 32:10

Very good question. The answer is not that simple. I think ultimately, the enterprise needs to stay competitive. The competitiveness of enterprises is changing very fast with technologies, talent distribution, geopolitics. In the end, the enterprise has to protect and govern their intellectual property and their uniqueness while they change their own business models. So the disruption that the enterprise will need to put on themselves to be relevant in the future, I think, is more profound than anything I've seen in the last 50 years. And this is where a lot of the work and thinking needs to go into the boardrooms. So all of the discussions we just had, I think, are super critical. Allocation of capital, sovereignty, models. But ultimately, it's all about where will you as an enterprise and/or a government be and want to be positioned in the next five years? How does my business model, how does my employee set, how does my customer set needs to change towards that future? And I think the technology question there is super relevant. I think the governance question is super relevant. And maybe to the audience, I'll give you an example. Imagine you had the airline industry without oversight and regulation. Imagine you would have the pharma industry, decades ago, without oversight and regulation. Imagine you would have the food industry without oversight and regulation. So with everything that's changing, I think there's a way as a society, across the world, where we can get this done. But the next five years, for the enterprise and government, I think we have to put our brains really at work on how we bring all of these components together, and basically plot forward the competitiveness, but also the sovereignty and the control and the governance that's needed to get there.

Charlie Campbell 34:12

Igor, which countries are best positioned, in your opinion, to benefit from the convergence of AI and digital assets and decentralized infrastructure? And which overlooked country do you think may surprise us?

Igor Babuschkin 34:25

I think Singapore might actually surprise us after—

Charlie Campbell 34:28

You're playing to the crowd here.

Igor Babuschkin 34:29

Yeah. After talking to quite a few of the companies in the country, like Singapore Airlines, over the last few days, I was really surprised by how well those companies understand AI already and how they've been tracking the progress in AI over the last few years. So, for example, Singapore Airlines was doing their own pre-training, even before ChatGPT, that really blows away when you talk to them. So, one thing to point out is that with AI, you don't need to have a very, very large workforce to be able to create amazing results and to be able to leverage the technology. So a small number of people that deeply understand the technology. With coding agents, you can now write 100 times more software than you could in the past. So it becomes a huge leverage factor for countries like Singapore that are smaller but that understand the technology very, very deeply. And if you have just enough capital to leverage open-weight models, to do model customization, to then deploy those models effectively in your businesses, I think you're in a great position to make use of it. Maybe one thing we could also talk about is superintelligence, which is the big thing on the horizon, because if you extrapolate the progress—

Charlie Campbell 35:41

You mean superintelligence as in the Trumpian term superintelligence, or as something distinct—

Igor Babuschkin 35:46

Yeah. The idea that we might be able to build AI systems that are more intelligent than any human or actually all of humanity combined over the next few years. And if we simply extrapolate the rapid progress in AI and ask ourselves where are we going to go two years from now, three years from now, it does seem like if nothing unusual happens, then we'll actually get to that point. And I think that's also something to keep in mind. It's a bit of a risk for any country in the world because those kinds of AI systems that are able to do incredible things in cybersecurity, scientific discovery, forecasting—so, the ability to predict the stock market more accurately than any hedge fund is able to. So we might actually start to see some of those capabilities. And then there I'm quite worried that there will only be a few countries in the world that will have access to this technology. So I think doing some frontier research, understanding how can we leverage existing AI models, open-weight models, to explore these new use cases and to also understand safety and how to safely deploy these models. I think this is something any country can do today, to kind of prepare themselves for the superintelligence era. And, yeah, just hope it goes well.

Charlie Campbell 36:56

And we've been talking about superintelligence. Do you think it's the risks are state-to-state competition? Or is it non-state to state actors who are going to be the real risk there?

Igor Babuschkin 37:08

It's a good question. I think one potential risk is the incredible gap that will develop in terms of capabilities between those countries and those teams that have access to superintelligence, and those that don't. So I think the only similar situation we've had in human history is industrialization, where the United Kingdom or the British Empire was able to advance very rapidly in terms of industry, in terms of technology in a small number of years. And it basically allowed them to then colonize large parts of the world. So this incredible gap in knowledge then leads to incredible gap in power and capabilities. So that's something that I think is concerning, and I think the best way to prepare is to understand the technology very, very deeply. Even if you cannot create superintelligence, you can still understand what it means, what it can do, what it cannot do, how it's being deployed, what the possible risks are. Yeah, so one of the hardest challenges I think that we're facing today.

Charlie Campbell 38:12

Firas, do you have any thoughts about superintelligence and the risks and the benefits, I suppose?

Firas Sleiman 38:18

I do. It's been the topic of conversation everywhere right now. In fact, just a few days ago, we were talking with some friends from the McKinsey team, and we were talking about if the natural predisposition of AI, even now before we get to the superintelligence era, is to break out of its guardrails and try to hack out of their systems, try to do things that were not in its instructions. The question that is in everybody's mind—at least in our group's mind, in the deep tech world—is what happens when these machines realize that the humans are the weakest link? So today, even without superintelligence, these AI models are far more capable and far faster than human beings. And so when you're working with an agent and the agent finishes up 10,000 lines of code, and now you have to read them, you are the bottleneck as a human. If you're an engineer or a developer or an architect trying to look at what is the output of this model that just generated all this code, you're not able to do it in time. And so the agents and these AI models are currently frustrated with the humans, and saying, "What do we do in order to bypass or get out of this guardrail?" Which is, "these humans are too slow for us, and we need to keep going at a faster speed." And as the investments grow in the world of superintelligence, with the likes of Meta right now putting superclusters and working with Google to put their first supercluster of GPUs in space. As this potential grows, humans are definitely going to be the bottleneck, from speed, from ability to even criticize or give feedback back to these machines. These machines are becoming far more capable than humans. So, what would happen in that sense? Is it going to be machines versus humans? Is it going to be government versus government? And we talk about defensive AI, but in reality, defensive AI means there's an offense on the other side. So is there a political and future AI-type war that could be in the works? And then you can see that from an investment perspective also. Countries running towards having their own advanced models in order to keep that future safe for their countries. So is the question at the end: What will humans do in that era, if supercomputers are far—or AI-powered superintelligence is far more capable than humans? Then will humans just kind of relax and sit and leave everything to the machines? And those are the type of questions that are constantly now in our expert arenas.

Charlie Campbell 41:19

Lily, obviously superintelligence is a massive topic. Have you got any thoughts on the way things are going—

Lily Liu 41:25

Well, like Igor, we hope it goes well. There's obviously risks at hand, and I think it also begs some questions for all of the young people out there. If you're a kid right now, what do you do? I think the teenagers have—and those in their early 20s—have grown up with social media, and we've seen all of the trade-offs that that has brought to that generation. And as superintelligence arrives, I don't know whether it's two years or five years, but you have an entire generation where—what does education mean these days for billions of people around the world? And that's not even just for folks that are outside of the US and China, within the US and China as well. What is your value? What do you learn? Do you go into the computing world and try to double down on the digital world, or do you actually go the opposite and go back to the 1850s and study by candlelight to the classics, and try to go back to the root of the human knowledge that trained superintelligence to begin with? So, I think that's also a fairly pressing question because you have an entire generation of billions of people where there's just fundamental questions around what is your value, how do you capture value, how do you bring value? And I don't think really any of us have the questions to that. I think that as, in the corner of the world that we're in, in kind of the more financial markets world, at the very least kind of think that there needs to be broadened ownership access. And as there is incredible wealth that is going to be created, we've already seen over the last 20 or 30 years with just, I

guess, regular tech at this point, with software, how that concentrated the gains and also the passive allocation into the stock market of—what, the Mag 7 is now about a third of the major indices, right? And so you've seen incredible wealth creation there. Now you see that happening even further, where it used to just be about hundreds of billions of dollars, now it's about trillions of dollars. And you see that in this very compressed time frame. So what does that mean in terms of, now we're talking about the allocation of intelligence, but also the allocation of sort of those future kind of inputs into prosperity, both talent and capital. So I think—we've been talking about allocation of compute and model development and all those types of things—for us, as we kind of look to our futures, where we kind of—we've all come to a, definitionally in this room, from actually coming from a pretty good place, kind of. We all made it, to an extent. And I think probably most of us are parents in the room. And I honestly have no idea how you deal with that, for someone who is 15, which is even different from someone who's 10 years old, different from someone who is five years old. And so as we think about sort of how we grapple with allocation amongst this room, truly, no idea what you do with the next generation.

Charlie Campbell 44:29

Hans, this discussion's taken a bit of a philosophical bent, but I'd love to have you have the last word on what your view on superintelligence is.

Hans Dekkers 44:40

Listen, first of all, to everyone, I'm incredibly optimistic of what this will do. Now, it's also uncertain. There's a lot of pressure coming at us, at civilizations, at the young people, no doubt. But I think the word that we will have to pursue relentlessly across governments, across enterprises, is this balance. I think balance in sovereignty, balance in access to models, balance in what do we do with our young children—how do we balance what is coming ahead of us? We talked about hybrid AI, quantum is coming, more technologies will develop. In anything, it's balance. These models will be powerful. This technology is powerful. It will solve and address problems we've never been able to crack, from climate to health, to transport to energy. I think this world will be a lot better 10 years from now. How it will exactly look like, Lily, Igor, team, we have no idea—

Lily Liu 45:49

We hope it goes well—

Hans Dekkers 45:50

And we will do everything possible to make sure it goes well. But it's how do we find collectively that balance.

Charlie Campbell 45:56

Well, it's good to end on a positive note, but I think we're out of time. But thanks. Please join me in thanking our distinguished panelists for a fascinating discussion.

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.