A NEW INNOVATION ECONOMY: A CONVERSATION WITH NVIDIA CEO JENSEN HUANG

Announcer 00:00

Please welcome President and CEO of Nvidia Jensen Huang in conversation with Institute Chairman Michael Milken.

Jensen Huang 00:09

Thank you. They're big fans of yours.

Michael Milken 00:30

I think that's for you.

Jensen Huang 00:33

I think the woo woo is for you.

Michael Milken 00:35

I think your outfit is for me. So welcome. So AI—is this the next industrial revolution?

Jensen Huang 00:50

Yes.

Michael Milken 00:51

Next manufacturing revolution?

Jensen Huang 00:53

Yes. Let me explain. So, all of us have been talking about the technology of AI, that it can perceive the world, it can generate content, it can translate, it can now even reason and solve problems, use tools, use the web browser, read PDFs, do research for you, and so we know what the technology is able to do, and that's very exciting in itself. It's transformative, completely in itself. We understand that the technology is unlike any IT technology of the past. Remember IT? Technology is a tool. You have to use it to make it effective. You have to sit in front of the computer and use it. But now AI has the ability to automate. And the concept of robotics, and robots are very well understood. And so imagine a physical robot, but, you know, we understand that. We can imagine that—imagine a digital robot, and it's in, you know, in the computer in your data center, doing work for you. And so, this is—it's exciting, because for the first time, it's no longer just replacing or the next generation of the IT technology that we know, but for the first time, it actually could augment and add to the digital workforce. So the part of the economy that it's part of is much larger than a trillion dollars—it's part of the 100 trillion dollars. And so that's the first layer. The second layer is, how do you generate this AI? Where does the AI come about? You know, whereas the last generation of computers was software written by hand and it runs on CPUs, what Nvidia took some 33 years to build, is this idea of a new type of computer that learns the machine, learns to write the software itself, and it runs on this processor, this computing platform we invented called accelerated computing and GPUs. So now the question is, how does the AI get produced? And it gets produced in essentially what people call data— AI data centers, but it's essentially a factory. It's unlike a data center. It doesn't look like a data center. It's quite large in scale. It does use energy, and it produces—you apply energy to it, and it produces these things called tokens, but they're basically numbers, and these tokens can be reformulated into numbers or words or images or pixels or videos or chemicals or protein combinations for drug discovery or even motor skills necessary to drive a robot or steering wheels to drive a self-driving car. And so that these tokens are being manufactured by this factory. And so what's interesting is that people are starting to understand that there's a whole new industry that has been created. This new industry has a factory, and this factory, there are AI factories. And how large can the these factories be? You know, they could be a giga—we're building ones that are about a gigawatt, and each gigawatt's about, you know, $50-$60 billion. And over the course of the next, you know, call it 10 years or so, I wouldn't be surprised to see tens of gigawatts of AI factories being built around the world. And so that's the second layer. The third layer, that's, that's probably even more profound is that, for the very first time, you have a capability, a technology that affects almost every industry, from financial services to health care to manufacturing, the logistics to, you know, e-tail to entertainment, you name it. And so this infrastructure, if you will, this AI factory now becomes an infrastructure for a whole bunch of other industries and just like the last generation, this infrastructure is kind of hard to understand. But the last generation, we had the information infrastructure, and the generation before that, we had the energy infrastructure. And now we have the intelligence infrastructure and the internet. The internet was—is the information into infrastructure, and artificial intelligence is this one. So now, you know, I

think when you look at AI from those different lenses, you could start to understand the impact of AI to the technology industry that we're in, to a new industry that every country wants to be part of—if anybody have excess energy, you're going to want to be part of this industry to the infrastructure that affects every industry.

Michael Milken 05:28

So let's step back for a couple minutes and talk about the skill sets needed to interact. We estimated a number of years ago, that if you took the most modern agricultural technology that— in the world, you might eliminate a half a billion jobs of what's going on in farming sub size, etc. There's significant questions today—who's going to be disintermediated. Now, I had a substantial advantage in the 1960s that I could calculate yields in my head, and then in 1970 they came out with the calculator, so I got disintermediated. Then—

Jensen Huang 06:20

It looks like he's done pretty well since the invention of the calculator.

Michael Milken 06:24

And I could remember millions of trades, and then computers started storing those trades. What do you see in the concept of work and the interaction with the technology you're going to provide?

Jensen Huang 06:40

Yeah, so we've—all of you have heard a lot about job displacement. Every job will be affected. Some jobs will be lost, some jobs will be created, but every job will be affected. And immediately it is unquestionable, you're not going to lose a job—your job to an AI, but you're going to lose your job to somebody who uses AI. But let me give you, and those are fairly common sense things to have observed. But let me give you the two extremes that you might want to consider as well. Computer technology, computer science has benefited about 30 million people. There are about 30 million people in the world who knows how to program and use this technology to its extreme, and it's really benefited all of us that have been in this industry the last 30 years, potentially one of the best and most, most wealth-creating industry you could have selected. I could have been a petroleum engineer. My dad was, and I could have been a doctor. My mom thinks that everybody should be a doctor. But I chose to go into computer engineering. And it turned out to have been quite a good choice. And however, there are about 30 million people like in this industry. And so we've created, Mike, in the last you know, 30-40, years, probably the greatest technology divide the world's ever seen. The instrument that we've invented, we know how to use, but, the other eight, seven and a half billion people don't. I'll put on the table that, in fact, artificial intelligence is the greatest opportunity for us to close the technology divide. And let me prove it to you. You know, if we just look in this room, it's very unlikely that more than a handful of people know how to program with C++, and an equal number know how to program in C. And yet, 100 percent of you know how to program in AI. And the reason for that is because the AI will speak whatever language you wanted to speak. You could draw a schematic and show it

to it. You could draw a picture and ask it what to do. You could, you know, obviously, talk to it in words. You could write a prompt. You could describe your prompt in a very explicit way. You could describe your prompt in a very implicit way. And if you don't know how to program that computer using AI, just tell the AI, I don't know how to program you. How do I program you? And the AI will tell you exactly how to program it. And so I think that— and the number of people who are using ChatGPT and Gemini Pro and these AIs kind of demonstrate that, in fact, this is one of the easiest to use technologies in history. And so now all—everybody could take advantage of this capability where there's a teacher or a student wanting a tutor, and every student should use it as a tutor. I use it as a tutor every day. And so I think the ability for us to now use artificial intelligence to close the technology gap is incredible. So that's one extreme. The other extreme that I will say is that, remember, we're—we have a shortage of labor. We have a shortage of workers. We don't have an abundance of workers. We have a shortage of and for the very first time in history, we actually have—we can imagine the opportunity to close that gap to put 30-40 million workers back into the workforce that otherwise the world doesn't have. And so you could argue that artificial intelligence is probably our best way to increase the GDP, the global GDP, and so those are two other ways to look at it. In the meantime, I would recommend 100 percent of everybody you know take advantage of AI and don't be that person who ignores this technology and this result loses [inaudible].

Michael Milken 10:27

So let's talk for a moment. They're going to walk out of this room, and Thursday, after six days of the conference, they're going to want to learn more about AI. Do they ask their computer to teach them about AI?

Jensen Huang 10:40

Yeah, excellent way to do it. Just pick up your phone, get yourself, you know, Perplexity is pretty good. ChatGPT is really excellent. Gemini Pro is excellent. I use all three of them, and I just ask it whatever you want to ask it about AI. And it'll tell you as deep as you like it to be. And I, sometimes I—in areas that are fairly new to me, I might say, start by explaining it to me like I'm a 12-year-old, and then work your way up into, you know, into a doctorate level over time and so you could all do the same.

Michael Milken 11:14

Let's look at it from another side. Jensen, your family came from Taiwan. You went to Washington, then eventually your parents moved to Oregon. And I've had a chance to finance many other entrepreneurs. Bill McGowan comes to mind at MCI. So he had a company that had 99 percent market share in AT&T that he wanted to take on. In those early years that you often talk about, you didn't know if you're going to make it or not. He often was wondering where the payroll was coming from every month. What did the other companies miss who had more access to capital than you did at the time? What did they not see that you saw?

Jensen Huang 12:10

Gosh—

Michael Milken 12:10

In other words, what did Intel not see in the market? What did they not recognize?

Jensen Huang 12:18

Yeah, the reason why I paused it to say it is because, from the very beginning, we imagined—so what we were trying to do as a company was to build—to invent a new way of doing computing that solves problems that normal computers can't. In fact, if you just wrote that mission statement out to do something that normal things can't, it's like—to go—I would like to build a new car to go places where normal cars can't. Well, usually, what happens if normal cars can't go there? Those places also aren't paved with roads, or they're not that desirable to go to anyways. And so we came up with this mission statement to solve problems that normal computers can't and several problems with that mission statement—turned out it took us 33 years to do and we succeeded at it. But the first thing is that the whole economy, the whole industry, the whole ecosystem, wants to go where problems can be solved. Nobody wants to go where problems can't be solved, and so where we are, was rather lonely. You know, there aren't other people solving this problem because it's hard to solve. There aren't many customers because they tend not to choose problems like that. They want to have the problems be solvable, not unsolvable. And then the other thing is that Intel watching us the whole time—I had the benefit—and you said they had greater source of cap—access to capital, and that's completely true because they were so successful doing what they were doing, they kind of rejected what we were doing. And that's, in fact, the good news. Over time, the reason why it took us so long is because it's hard, and the reason why we're here alone is because people left us alone for a long time, and there was a book that was written recently, and I picked it up and skimmed it. Peter Thiel's Zero to One book, in a lot of ways, is kind of a story by Nvidia too. You know, we chose to do something that nobody thought was possible or very hard to do and very unlikely to succeed. But to us it was very common sense. And so I think simultaneously because it was hard to do, and also because they were so successful doing what they were already doing, they kind of rejected the idea until everything came together.

Michael Milken 12:41

Well—and you're also trying to make sure your company doesn't go in the direction of Intel also. So you're the leader today. How do you get that culture of constant innovation? And if I wanted to talk about Captain Kirk, to go where no one has gone before.

Jensen Huang 15:08

I think partly, first of all, there's just no guarantees but, we have several things about our company that's really quite extraordinary, and I appreciate it as a person. I wish upon my kids and people I love to have the same history, same experience, which is that long suffering that comes with struggle, and you never take anything for granted.

You're super, super efficient. You're trying to save everything you can. Save every penny you can. Because you don't know how long the struggle is going to last. You have incredible resilience because it took a long time to do it. And so the company has that in its character. Almost everything we undertake these days, even these days are, you know, 5-10 year endeavors. We're probably the deepest in this new area called physical AI, which translates to robotics in the world and the fundamental technology necessary for the next generation of AI. We're probably the furthest along, the deepest of anybody, and so I think, those characteristics of dreaming big on the one hand, and having the resilience and the character to suffer until you see it happen. I think that's very good. I think the other part that's good is you're always going out of business. For us, for 30 years, we're always in a perpetual state of going out of business. And so you don't take anything for granted. And I don't—you know when there's a setback, it doesn't trouble me too much, when we make mistakes, it doesn't surprise me too much, when we have success, I don't take it for granted, and we don't celebrate it too much. And we really stay, you know, stay focused on doing our work. And so I think part of that is, just how long it took to build a company.

Michael Milken 17:13

Let's talk for most laypersons, how do you make a chip? What is required to make a chip? So we'd all like to go out there and make chips. We have no idea how to go about it, but we'd like to, and as you remember, the US passed an act, we're going to invest $62 billion, and then they discovered, six months later, there's no one in the United States that knows how to build a factory, and we need to get 7,000 people from Taiwan here.

Jensen Huang 17:45

Well, with all things, I think craft and artistry matters. If you want to learn how to build a chip, you know, I would start with YouTube and then—so it turns out we started—we're very good at building chips, and the reason for that is because we build—not since IBM in the 60s, has a company like us existed where we come up with a blank sheet of paper design, a brand new architecture, create the chips, create the systems, create the networking, create the infrastructure, write all the software, take that software to the market. Have the world's developers and ecosystem develop for that computer, kind of like we develop for iPhones, and you develop for Windows, you develop for Nvidia. And so not since the 60s and 70s, when IBM built everything from the ground up, has a company like us existed. We build the chips, but we build the entire system. And we're really an AI infrastructure company today. If you look at the systems, we build, each one of our chips, if you will, it's a ton and a half chip. It's $3 million each. We build these things in very high volume. We manufacture it, assemble it, and then we test it. We use a supercomputer to test a supercomputer, because you have to be smart to test if the computer you make is smart. And so we test—everything is liquid cooled. And then we test everything, assemble everything. We disassemble everything, put it on a plane, ship it to wherever the data center is, assemble it again outside their door, put it inside their data center. This entire process has 200 manufacturers and suppliers working with us around the world. We build a couple —100 billion dollars of it a year. At this moment, we're the largest technology company—chip company in the world, I guess, if you will. And so the—it's incredibly hard to do. Our R&D budget, you know, per generation, is probably about 20-30 billion dollars. And so these are—this is a giant game, but we're working into an industry, you know, Mike, that you know that the intelligence industry will likely be measured in trillions of dollars, and so the amount of investment that we're making is warranted for the opportunities ahead.

Michael Milken 20:18

So we've all had a chance to read about potential restrictions on where you're going to be able to sell your chips. There's pros and cons that people put forth in this debate. Lay out the issues as you see them.

Jensen Huang 20:37

Nvidia's technology is oftentimes described as a national treasure. And the technology, obviously, is important to this new industry called artificial intelligence. And so one extreme—in one side, we want to make sure that this technology is available only to the friends of our nation. We want to ensure that access of this technology doesn't fall into the hands of people who might use it for military reasons. And so those are the arguments for limiting the access for economic reasons, for national security reasons, and the fallacy of that is no government, especially the government of our adversaries, are limited by the available capacity of computing in their country for their military reasons. Our country is not. No countries are. If they need it for military advance, they'll just secure whatever computing resources that they already have. And there are millions and millions of Nvidia chips in just about every country already and so it's not going to limit shipping additional GPUs—Nvidia technology into whatever country—is not going to limit their military. I think the reason why—the reason for leaning into the export of this technology is we want to build the world's AI. Whereas American standards are being adopted around the world, the ecosystem of artificial intelligence will build on top of our standard versus somebody else's standard. And we are not alone. Nvidia, of course, is the world leader today, but in our absence, if we don't serve a particular market, if we leave a market altogether, there's no question, somebody else would step in. Huawei, for example, is very formidable, one of the most formidable technology companies in the world. They'll step in and so the the reason why it would make sense is to win in the marketplace. Make the American standard the global standard. Have AI be built on top of American technology and, of course, very importantly, it's a giant market. You know, when we were limited to ship to China, the Chinese market in a couple years is probably about $50 billion, the the market we've left behind is utterly gigantic. $50 billion—so you have a feeling for that number—$50 billion is like Boeing, not a plane, the entire company. And so that's the business opportunity that I think we could enjoy, bring back tax dollars, create jobs, advance our technology even further.

Michael Milken 23:38

Well, also your interaction with the customer, you're missing that chance to learn from that interaction customer.

Jensen Huang 23:46

The most important thing for any business is the interaction with a customer.

Michael Milken 23:50

And what have your customer—what have you learned from your customers over the last few years, and their demand for one, their demand for chips, but their use of chips, and how is that feedback to you and Nvidia?

Jensen Huang 24:04

Well, we work with just about every AI developer in the world and we learn everything from how our architecture, how the nature of our technology is used, is optimal or not optimal for the future of AI. And so when we understand what an AI researcher would like to do with, for example, the AI model for a virtual cell, you know, we've now made very good progress with virtual proteins. And so now we're working on virtual cells. And if we can understand how cells interact and how its pathways can be expressed, and understand basically the meaning of cells and its dynamics, and we could do that with AI. So that AI model is very different than AI models for large language models, and so understanding how people want to use that helps us change our architecture in the future to be better suited.

Michael Milken 25:02

So as I mentioned to you Jensen, you know, many years ago, I went to IBM and tried to convince them, at what they called their super chip at the time, that we should use it for medical research. And then they wrote me a Dear John letter—Thanks, Mike. Great presentation, but we're going with video games. So where does the demand you see in the future? Is—are certain industries going to play a larger role, such like bioscience, and what you could do in that area? Where do you see the demand from different industries?

Jensen Huang 25:41

Well, if you look at where AI is today, as large as Nvidia has been already, and as large as the AI industry is already, we're serving basically the consumer internet market. If you just kind of take a step for a second, that's the only tiny little sliver of the global economy that we're serving. Above this are health-care industries, life sciences, manufacturing, the actual manufacturing. In the future, the factory will be one gigantic robot, orchestrating a whole bunch of robots inside, working with people to build products that are robotics. So you have robots building robots, building robots. And this nested layer of technology is nearly upon us. And so that application for manufacturing, industrialization, plants and factories and that entire area needs this new AI called physical AI. If we can solve that, we're talking about trillions and trillions of dollars of industries.

Michael Milken 26:48

So last question—we might have a lot of people after listening thinking, how do I get a job at Nvidia? Okay, what are the skill sets you're looking for today?

Jensen Huang 27:00

Well, if you can't design a chip—if you tell me that you learn how to design chips on YouTube, I think that's gonna tell me a lot. And so, no, no—look, you know, we're—Nvidia is the world's first chip company, of course, that we, as I mentioned, we build the entire AI infrastructure stack, but we have digital biologists, we have quantum chemists, we have computer graphics engineers, we have roboticists, we have language experts. We have expertise across a very large domain of science, and we have a very large domain of industries. And so, you know, we serve health-care industry, we serve financial services industry. So if you have domain expertise, we love that. We love people with domain expertise. And so we also love people with general basic intelligence. And if you love hard work, and especially if you love to suffer, you know exactly who to call.

Michael Milken 27:27

We have seen over the years those—I don't know if my microphone is working or not. Can you hear me? A famous professor friend of mine tells me, when he gets a new class, he tries to figure out who got in because of their ability and who got in because of their family relationships. And that is—so I was saying, when you talk about what built Nvidia— hard work, a lot of challenges, tough days. I was saying this professor, friend of mine, points out that he tries to figure out when the class comes in, who got in on their ability and who got in because of relationships. That's easy, he tells me, within a week. What's harder is, how long will it be before the person that got in on relationships is working for the person with ability. That takes a long time, he needs to get to know the students. Those hard days have paid off, and we look forward to see what you can accomplish in the future.

Jensen Huang 29:26

Thank you. Thank you very much, Mike.

Disclaimer: This transcript was generated by AI and has been reviewed by individuals for accuracy. However, it may still contain errors or omissions. Please verify any critical information independently.